Progressivity in Consumption or Income Taxes? General Equilibrium Analysis and Distributional Effects of a Progressivity Rebalancing in a Heterogeneous-Agent Framework

Andrea Maestri*

May, 2025

Under the supervision of Prof. Axelle Ferriere

Thesis Committee: Prof. Xavier Ragot

Abstract

Income and consumption taxes are major sources of fiscal revenue for all advanced economies. While progressivity is a well-established principle in the design of income taxes, consumption taxes are typically implemented through a single flat rate or a limited set of differentiated rates based on the nature of goods and services. These differentiated schedules introduce implicit progressivity in many consumption tax systems by varying tax burdens across households with different consumption profiles. Using a tractable heterogeneous-agent model with incomplete markets, this study investigates the macroeconomic and distributional effects of allowing both labor income and consumption taxes to exhibit progressivity in the same fiscal regime. To do so, I simulate a series of revenue-neutral fiscal regimes in which progressivity is shifted between the two tax instruments, in a novel exercise which I refer to as progressivity rebalancing.

The analysis reveals three key findings. First, the redistributive capacity of a tax depends critically on the characteristics of its base: labor income taxation, acting on a more unequal distribution, has stronger redistributive potential but higher distortionary effects. Second, welfare effects operate through both redistribution and insurance channels, with lower-income households particularly sensitive to the erosion of consumption smoothing when labor tax progressivity is reduced. Third, progressive consumption taxation emerges as a flexible second-best tool that can offset equity losses from flatter labor taxation, preserving redistribution and insurance while reducing efficiency losses. These findings highlight the relevance of progressivity in consumption tax design as a core component of modern fiscal policy, especially in contexts where labor tax progressivity is constrained.

*I gratefully acnknowledge invaluable support from my supervisor, Axelle Ferriere. Her technical expertise, human qualities and personal dedication were crucial to the final outcome and strongly contributed to make this journey rewarding and enriching. All remaining errors are my own. I wholeheartedly thank the long list of friends that walked together along the way, offering support, insightful conversations and human warmth. I dedicate this to my father. Per aspera ad astra!

Contents

1 Introduction

Consumption taxes represent a cornerstone of tax systems in advanced economies. Excluding social security contributions, they often rank as the second most important source of government revenue after income taxes. In 2022, consumption taxes accounted on average for 29.6% of total tax revenue and 9.9% of GDP across OECD countries (OECD, 2024). Among them, value-added taxes (VAT) and excises make up the bulk of receipts, contributing 20.8% and 8.2% of total tax revenue, respectively. These taxes are often praised for their economic efficiency, ease of collection, and neutrality in production decisions. However, their distributional implications remain a central issue in public economics and policymaking.

The embedded degree of progressivity of consumption taxes critically depends on the structure of the tax system. While these taxes may appear proportional when measured as a share of total consumption, their actual burden across households is shaped by differences in tax design and household consumption patterns. In particular, excise duties on specific goods and the common use of differentiated VAT rates across different classes of goods, combined with heterogeneity in consumption patterns across the income distribution, generate variation in effective tax burdens and consequent redistributive effects. Likewise, similar redistributive effects can be generated by informality and heterogeneity in tax pass-through. Yet, the literature has so far focused exclusively on labor income taxation as the primary channel for redistribution, overlooking the general equilibrium and welfare consequences of implicit progressivity in consumption tax schemes.

The aim of this paper is make a first step to address this gap by examining the macroeconomic and distributional consequences of the coexistence of progressivity in both labor income and consumption taxes within a heterogeneous-agent model with incomplete markets. The analysis proceeds in two steps. First, the model is calibrated to a realistic baseline. Then, a "progressivity rebalancing" exercise is conducted by simulating a set of revenue-neutral counterfactual tax regimes. In each simulation, the degree of labor income tax progressivity is exogenously perturbed, while the progressivity of consumption taxation is endogenously adjusted to maintain fiscal balance. A similar exercise is then performed to pin down effects generated by consumption tax design alone. This approach allows for a systematic exploration of how shifts in the composition of progressivity influence aggregate efficiency, inequality, and the degree of insurance provided by the fiscal system.

The analysis yields three main insights. First, the redistributive power of a tax instrument is closely linked to the structural characteristics of its base. Labor income taxation exhibits greater redistributive capacity than consumption taxation, both because it accounts for a larger share of total government revenue and because it operates on a substantially more unequal distribution. These features allow labor taxes to generate more pronounced transfers between the tails of the income distribution. Second, changes in tax progressivity affect welfare not only through static redistribution but also through the insurance channel. At the top of the distribution, income effects dominate behavioral responses, while at the bottom, a reduction in redistribution weakens implicit social insurance, increases exposure to consumption risk, and leads to adjustments in savings and consumption behavior. Risk-sharing thus emerges as a critical determinant of welfare outcomes, particularly for low-income households. Finally, progressive consumption taxation can serve as a powerful second-best instrument for redistributing progressivity within the tax system. By shifting the tax burden toward higher-consuming households — who are typically also higher earners — it creates fiscal space to reduce labor income tax distortions, particularly at the top of the distribution. This rebalancing enables the tax system to incentivize labor supply and capital accumulation, generating gains in efficiency and productivity. Crucially, it does so without sacrificing aggregate welfare or exacerbating inequality. Rather than relying on explicit transfers, progressive consumption taxation embeds redistribution and insurance directly into the structure of the tax code, allowing for efficiency-enhancing reforms in labor taxation to coexist with equity and risk-sharing objectives.

While this paper does not take a normative stance on optimality, the insights it offers aim to serve as a reference point for future research on optimal fiscal policy involving progressivity in both labor income and consumption taxation. The remainder of the paper is organized as follows. Section 2 elaborates on the determinants of progressivity in consumption taxes and provides useful theoretical references on which this work is grounded. Sections 3 and 4 lay out the quantitative model developed for the analysis, together with its calibration. Section 5 describe the results of the simulation exercises and identifies several directions for future work.

2 Progressivity in Consumption Taxes: Theory and Practice

Determinants of Progressivity in Consumption Taxes

In spite of widely used flat-rate designs, the incidence of consumption taxes along both the income and the consumption distribution is generally heterogeneous, with consequent distributional effects. The distribution of effective consumption tax paid across the household population reflects heterogeneity along two dimensions: an extensive margin, driven by differences in households' total consumption expenditure, and an intensive margin, determined by the composition of the consumption basket, which affects the overall consumption tax rate. As the first dimension is generally prevailing and saving rates increase with income, consumption tax designs are often found to be regressive with respect to the income distribution (Blasco et al., 2023; Thomas, 2022). However, the degree of incidence varies also along the consumption distribution due to heterogeneous consumption patterns across households.

Three main mechanisms stand behind this so-called bundle effect. Firstly, most existing consumption tax systems are characterized by non-uniform designs that include differentiated VAT rates, excise duties, and other fiscal instruments intended to tailor tax burdens across goods and services — often justified on grounds of equity. In practice, tax codes typically apply reduced VAT rates or exemptions to essential goods, such as food and medicines, while imposing higher rates on non-essential or "luxury" items, including leisure, entertainment, and high-end consumer goods. Table A1 provides an illustrative example, reporting the VAT rates applied to various goods and services in France as of April 2025. Since goods and services subject to reduced VAT rates typically constitute a larger share of the consumption baskets of low-income households, these households often face a lower average effective consumption tax rate compared to higher-income households. Several studies have pinned down how differential VAT rates affect the effective tax rates along the consumption distribution, finding proportionality or small progressivity for several advanced economies (Blasco et al., 2023; Institute for Fiscal Studies (IFS), 2011; O'Donoghue et al., 2004; Thomas, 2022).

The second factor affecting the incidence of consumption tax distribution, particularly in the case of developing economies, is informality. As low income households are more likely to purchase informally, even flat consumption taxes happen to be less regressive than expected. Bachas et al. (2024) show that the large share of informality at the bottom of the households distribution makes consumption taxes de facto progressive in a sample of 32 low and middle income countries. Finally, another factor potentially generating heterogeneity of consumption tax incidence is asymmetric pass-through of consumption tax rates. This phenomenom has been widely documented in the literature (see, for instance, Benzarti et al. (2020)) and could generate progressive or regressive effects if the pass-through is correlated with consumption patterns.

The modelling framework used for this analysis can well capture heterogeneity in effective consumption tax rate with the use of a single homogeneous good, which proxies for the households consumption basket. Even though the main focus remains to explore the effects of this "implicit" progressivity in consumption taxes already in place in most advanced economies, the adopted modelling framework allows to equally explore hypothetical steady-state equilibria achieved through consumption tax reform aiming for "explicit" progressivity, i.e., setting differential rates according to households' income.

Consumption Taxation: A Theoretical Perspective

This article broadly relates to the large literature focusing on consumption taxation with a Ramsey approach. Several papers have contributed to study the effects of consumption taxation, but their conclusions vary widely, reflecting sensitivity of the outcomes to the specific modelling assumptions regarding market completeness, agent heterogeneity, and risk. A seminal contribution in this strand is provided by Coleman (2000), who finds positive welfare gains from switching to a regime where consumption is taxed while labor is subsidised in a complete-markets infinite-horizon economy. Similarly, Correia (2010) suggests that shifting from classic labor income and capital taxation to flat consumption taxation generates large efficiency gains and reductions in inequality in an economy with infinitely-lived heterogeneous agents and complete markets. On the other hand, in a OLG model with heterogeneous agents, Nishiyama and Smetters (2005) show that in presence of uninsurable idiosyncratic risks replacing a progressive labor income tax system with a linear consumption tax reduces efficiency by loss of intragenerational risk sharing. On the same wavelength, Cahn et al. (2024), from which we borrow in terms of calibration, study the effect of a fiscal rebalancing, i.e. a partial replacement of payroll taxes with consumption taxes in a budget-neutral way in a heterogeneous-agent model calibrated to the French economy. In the aggregate, the reform generates increases in the supply of labor and capital at the expenses of an increase in inequality, as wealthier agents capture the majority of the reform gains, and a reduction in total welfare. More recently, Macnamara et al. (2025) have revamped the results from Coleman (2000) and Correia (2010) by showing how raising most fiscal revenues from consumption and providing redistribution and insurance via a lump-sum transfer is optimal in a standard heterogeneous-agent model with idiosyncratic risks.

The underlying quantitative framework used in the analysis, instead, relates more closely to the recent growing literature on optimal fiscal progressivity in heterogeneous-agent models with idiosyncratic risks and incomplete markets. Particularly, Heathcote et al. (2017) analyse the optimal labor income tax progressivity in relation to both its risk sharing role and the implied labor supply distorsions. Building on this work, Ferriere et al. (2023), from which I borrow for the quantitative methods, include means-based transfers in such setting as means to enhance redistribution and efficiency.

Finally, this paper places itself in a yet underdeveloped literature on progressivity in consumption taxes. An early contribution in this field has been achieved by Li (2013), who quantitatively evaluates welfare effects generated by the introduction of a simple progressive consumption tax. Using consumption deduction thresholds combined with flat-rate taxation, she finds that switching from labor income taxes to consumption taxes generates an increase in aggregate capital and labor supply due to precautionary motives and a reduction in welfare inequality. The closest work to mine, instead, is Shaw (2020), who evaluates the impact of a transition to a fiscal regime based on a non-linear progressive consumption tax in a rich life-cycle model with heterogeneous agents and uninsurable productivity risks. Echoing my results, he finds that the reform produces significant gains in output, consumption and welfare, mostly coming from the reduction of labor distorsions. Finally, introducing progressivity by mimicking common differential VAT rates, Parodi (2023) characterises optimal consumption and labor income taxation in a structural life-cycle featuring different flat-consumption tax rates on different consumption goods.

3 Quantitative model

To study the joint coexistence of progressivity in labor income and consumption taxes, I consider an economy populated by a normalised continuum of infinitely lived households à la Aiyagari (1994) facing idiosyncratic risks and supplying labor endogenously, a representative firm producing a single homogeneous good and a government levying a progressive labor income tax, a progressive consumption tax and a proportional tax on capital returns. Time is discrete, with each period indexed by $t \in \{0, 1, ..., \infty\}$. Markets are all perfectly competitive.

3.1 Households

Households in the economy differ along two dimensions, wealth and productivity, denoted respectively by a and ρ. Thus, in each period households are characterised by the individual state vector (a_t, ρ_t). Individual productivity ρ_t is *i.i.d.* across households and follows a Markov process over a bounded support $\mathcal{P}$ according to the stationary transition matrix $\pi(\rho_t, \rho_{t+1})$. Markets are incomplete and households cannot hedge against idiosyncratic productivity shocks. Lifetime utility has CRRA functional form and households choose consumption c_t and hours of labor supplied h_t in each period t according to

$$\max_{\{c_t, h_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \mathbb{E}_t \left(\frac{c_t^{1-\sigma}}{1-\sigma} - B \frac{h_t^{1+\psi}}{1+\psi} \right),$$

where β represents the subjective discount factor, σ the coefficient of relative risk aversion, B the disutility of labor relative to consumption and ψ is the inverse of the Frisch elasticity of substitution for labor. The fraction of after-tax income that is not consumed can be invested in the riskless financial asset yielding endogenous real interest rate r_t.

The inter-temporal consumption-saving choice is bound by the intra-temporal resource constraint

$$c_t + a_{t+1} = y_t + (1 + (1 - \tau_k) r_t) a_t - T_c(c_t) - T_y(y_t)$$
$$\text{with} \quad y_t = \rho_t w_t h_t$$

where y_t represents household's labor income and is proportional to the wage rate w_t, a_t denotes current period's asset holdings. $T_c(c_t)$ and $T_y(y_t)$ denote respectively consumption and labor income taxes to be paid, while τ_k is the proportional tax rate on capital returns. Households borrowing is subject to a limit $\underline{a}$, restricting the compact set of possible asset holdings $\mathcal{A}$ to $[\underline{a}, \overline{a}]$, with $\overline{a}$ being a large positive number.

3.2 Representative Firm

The representative firm operates a constant-returns-to-scale production technology to produce the economy's single good. Output at time t is determined by the Cobb-Douglas production function:

$$Y_t = F(K_t, L_t) = A_t K_t^{\alpha} L_t^{1-\alpha},$$

where Y_t is output, A_t denotes total factor productivity, and K_t and L_t represent aggregate capital and effective labor inputs. Capital is remunerated with a constant share $\alpha \in (0,1)$ of output and depreciates at a rate $\delta > 0$.
Each period, the firm chooses K_t and L_t to maximize profits:

$$\max_{K_t, L_t} A_t F(K_t, L_t) - (r_t + \delta) K_t - w_t L_t$$

The optimality conditions for capital and labor are given by:

$$r_t = A_t F_K(K_t, L_t) - \delta, \quad w_t = A_t F_L(K_t, L_t),$$

where F_K and F_L are the marginal products of capital and labor.

3.3 Government

The government finances its exogenous expenditure G by levying progressive taxes on consumption and labor income, along with a proportional tax on returns to capital. As I abstract from public debt, the intra-temporal government budget constraint must be satisfied period by period, i.e.

$$G_t = \int_{\mathcal{A}\times\mathcal{P}} \Big(T_c(c_t) + T_y(y_t)\Big)\, d\Gamma_t(a,\rho) \;\; + \tau_k r_t \int_{\mathcal{A}\times\mathcal{P}} a_t \, d\Gamma_t(a,\rho).$$

Both progressive taxes follow a non-linear Feldstein (1969) functional form, defined as

$$T_z(z) = z - \lambda_z z^{1-\tau_z} \quad \text{for} \quad z \in (c, y),$$

where the scale parameter $\lambda_z \in (0,1)$ determines the average level of taxation and τ_z governs the degree of progressivity. When $\tau_z = 0$, the tax is proportional with flat rate $1 - \lambda_z$; negative values imply regressivity, while $\tau_z > 0$ introduces progressivity, with marginal tax rates increasing in the tax base z.

The interaction between λ_z and τ_z also regulates the break-even level where net taxes are zero, which is given by $z_0 = \lambda_z^{1/\tau_z}$. Agents with tax base below this threshold, $z < z_0$, receive a positive net transfer, while those with $z > z_0$ are net contributors.

3.4 Equilibrium

Definition 1: Given a tax regime $(\lambda_y, \tau_y, \lambda_c, \tau_c, \tau_k)$, an initial distribution $\Gamma_0(a,\rho)$ over the state space $\{\mathcal{A} \times \mathcal{P}\}$ and the associated transition matrix $\Pi(a,\rho)$, the **competitive stationary equilibrium** is defined as the set of a value function $V(a,\rho)$ and policy functions $g_a(a,\rho)$, $g_c(a,\rho)$, $g_h(a,\rho)$ for the households, a vector of firm's allocations (K, L), a stationary distribution $\Gamma(a,\rho)$ and a vector of equilibrium prices (r_t, w_t) such that for all t

1. Household's utility maximisation problem is solved by the value function and the associated policy functions. Recursively,

$$V(a,\rho) = \max_{\{c,a',h\}} \left\{ \frac{c^{1-\sigma}}{1-\sigma} - B\frac{h^{1+\psi}}{1+\psi} + \beta \sum_{\rho'} \pi(\rho'|\rho) V(a',\rho') \right\}$$

subject to

$$
\begin{aligned}
c + T_c(c) &= y - T_y(y) + (1 + (1-\tau_k)r)a - a' \\
y &= \rho h w \\
a' &= g_a(a,\rho) \\
c &= g_c(a,\rho) \\
h &= g_h(a,\rho) \\
c &\geq 0 \\
h &\in [0,1] \\
a' &\geq \underline{a}
\end{aligned}
$$

2. Equilibrium prices satisfy firm's optimality conditions:

$$
\begin{aligned}
r &= AF_K(K,L) - \delta \\
w &= AF_L(K,L)
\end{aligned}
$$

3. Government budget constraint is satisfied

$$
G = \int_{\mathcal{A}\times\mathcal{P}} [T_c(c) + T_y(y)]\, d\Gamma(a,\rho) \;\; + \tau_k r \int_{\mathcal{A}\times\mathcal{P}} a\, d\Gamma(a,\rho)
$$

4. Markets for labor, capital and goods clear

$$
\begin{aligned}
K &= \int_{\mathcal{A}\times\mathcal{P}} g_a(a,\rho)\, d\Gamma(a,\rho) \\
L &= \int_{\mathcal{A}\times\mathcal{P}} \rho \cdot g_h(a,\rho)\, d\Gamma(a,\rho) \\
C &= \int_{\mathcal{A}\times\mathcal{P}} g_c(a,\rho)\, d\Gamma(a,\rho) \\
AF(K,L) &= C + \delta K + G
\end{aligned}
$$

5. The distribution over individual states is stationary, i.e. $\Gamma'(a,\rho) = \Gamma(a,\rho)$, and consistent with the household's policy rules and transition matrix for states $\Pi(a,\rho)$.

4 Calibration

The model is calibrated to the French economy. A period is assumed to be one year. Several structural parameters are taken from the literature, others tailored to match specific moments that are relevant for the current analysis.

Productivity process

To model the households' productivity process, I assume its log follows an AR(1) process

$$
log(\rho_t) = \theta_\rho log(\rho_{t-1}) + \sigma_\rho \varepsilon_t \quad \text{with} \quad \varepsilon_t N(0,1).
$$

where the persistency θ_ρ and the volatility σ_ρ parameters are set to 0.9183 and 0.2976, as estimated by Cahn et al. (2024) for the French economy. Following Kopecky and Suen (2010), I discretise the productivity process with a 7-state Markov process using the Rouwenhorst (1995) method.

The productivity process so calibrated allows to reproduce quite well the pre-tax income distribution for France, while not resulting as accurate in capturing concentration of wealth at the top of distribution [1]. Figure 1 and Table 1 compare model performance to net wealth distribution moments, pre-tax income shares and income ratios in the data.

[1] A future version of this paper will include further adjustments to the income process - such as including a "superstar" state in the style of Cahn et al. (2024) - to improve the fit for the wealth distribution.

Ratio	Model	Data
T10/B50	7.3	7.2
T10/B90	4.1	4.2

Table 1: Pre-Tax Income Distribution Ratios — Baseline Model vs Data

Notes: Data for France (average 2010–2018). Source: Bozio et al. (2020).

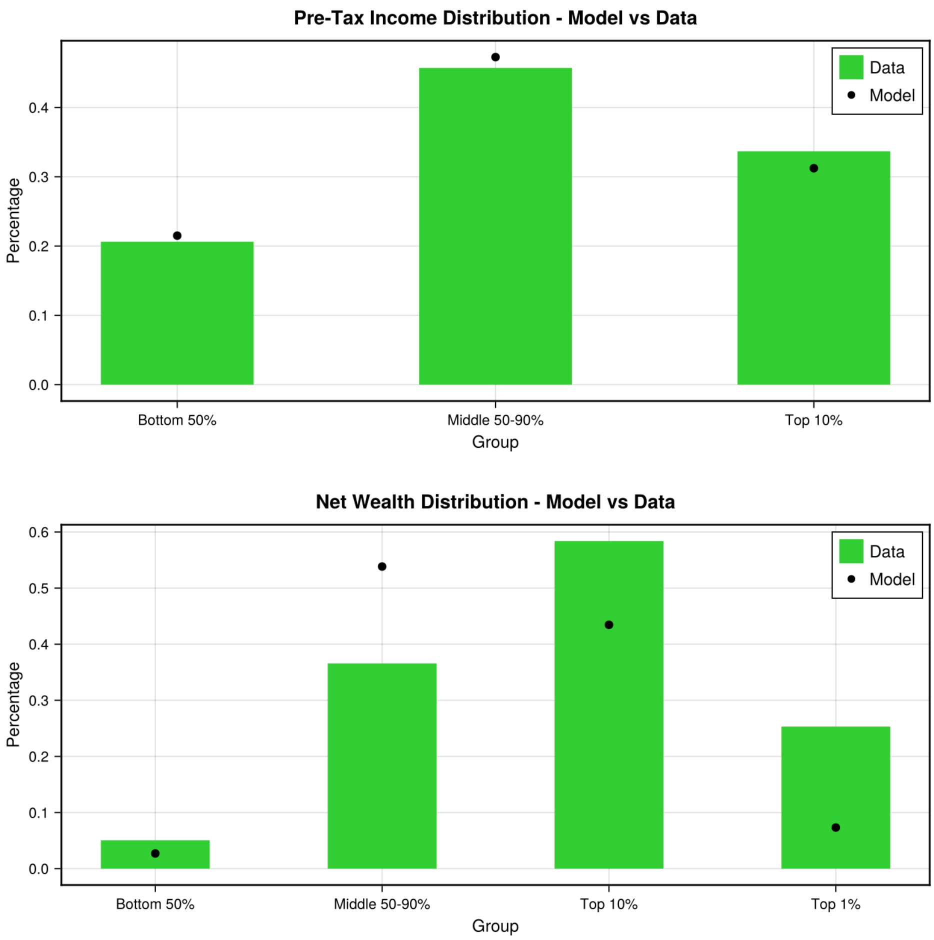

Figure 1: Pre-Tax Income and Net Wealth distributions - Baseline Model vs Data

Notes: Data for France (average 2018–2023). Source: World Inequality Database (2024).

Households and Firm

I set the relative risk aversion coefficient $\sigma = 1.5$ and the inverse of the Frisch labor supply elasticity $\psi = 2.5$, as standard in the literature. Firm's total factor productivity coefficient A is normalised to 1 and the capital share α is set to 1/3, consistently with the data. The borrowing limit for the household $\underline{a}$ is computed endogenously to avoid the insolvency trap, as described in Appendix C. Finally, the discount factor, the depreciation rate for capital and the disutility scale parameter for labor are adjusted to target i) capital-to-output ratio (2.9 for France [2]), ii) a reasonable interest rate consistent with discount factor (2.98) and iii) average hours worked (0.3), thus setting $\beta = 0.9675$, $\delta = 0.0840$ and $B = 85.00$.

[2] As in the AMECO Database (European Commission, 2024).

Tax regime

The tax regime in the model is defined by the set of five parameters $\lambda_y, \tau_y, \lambda_c, \tau_c, \tau_k$. To start with, we pin down the flat capital return tax rate $\tau_k = 0.352$ following Cahn et al. (2024). Then, we calibrate the labor income tax parameters λ_y and τ_y to target the average effective rate for the bottom 50% and top 10% of the distribution to their counterparties in the data for 2022, i.e., -1.5% and 31.0% [3]. This yields $\lambda_y = 0.7305$ and $\tau_y = 0.1875$. Finally, we adjust the two parameters defining the consumption tax regime λ_c and τ_c to target i) the Kakwani index for consumption taxes in France (0.0457), as estimated by Thomas (2022) and ii) the share of revenue from consumption taxes for France in 2024 (28%) (OECD, 2024). Thus, we set $\lambda_c = 0.8500$ and $\tau_c = 0.0343$. The resulting tax regime yields a G-to-Y ratio of 22.29%, matching closely the average level of 23.26% in the data for the period 1990-2023 (World Bank, 2024) [4]. Figure A2 reports average effective and marginal rates for labor income and consumption taxes in the baseline calibration by distribution decile.

Table 2 summarizes the calibration of the structural parameters of the baseline model.

Parameter	**Value**
β - Discount factor	0.9675
σ - Relative risk aversion	1.5
B - Disutility of labor	85.0
ψ - Inverse of Frisch Elasticity	2.5
$\rho_{AR(1)}$ - Productivity process persistency	0.9183
$\sigma_{AR(1)}$ - Productivity process volatility	0.2976
N. States - Rouwenhorst discretisation	7.0
α - Capital share	0.3333
δ - Depreciation rate	0.084
A - Total Factor Productivity	1.0
λ_y - Labor tax scale	0.7305
τ_y - Labor tax progressivity	0.1875
λ_c - Consumption tax scale	0.85
τ_c - Consumption tax progressivity	0.0343
τ_k - Capital return tax rate	0.352
$\underline{a}$ - Borrowing Limit	-1.1385

Table 2: Calibration Summary

5 Results

5.1 Baseline Model

The baseline model is calibrated as described in Section 4 and solved using Value Function Iteration. The complete numerical solution is described in Appendix C.1. Figure A1 reports descriptive figures for the household policy functions.

5.2 Simulating Alternative Tax Regimes

To assess the distributional and macroeconomic consequences of a progressivity rebalancing reform, I simulate counterfactual equilibria which entail different degrees of progressivity of consumption and wage taxes. In order to estimate these equilibria, I exogenously perturb the parameter regulating the degree of labor income tax progressivity, τ_y. For each perturbed value of τ_y, I endogenously solve for a new level of consumption tax progressivity, τ_c, that yields a tax-revenue-equivalent regime. Specifically, the goal is to operate a revenue-neutral progressivity reform, holding constant total government revenue G, as computed under the baseline tax system.

[3] Computed drawing from Direction Générale des Finances Publiques (DGFiP) (2024). These effective tax rates do not take into account net transfers outside the labor income taxation framework. A future version of the paper will include the transfers in the calibration.

[4] Government expenditure here excludes social security transfers. A future version of the paper will include them, reflecting this also in government expenditure.

The exercise spans 48 counterfactual regimes, defined by values of τ_y varying symmetrically in a range of $\pm 20\%$ around its baseline value. For each regime, I solve a two-level fixed-point problem to ensure that the equilibrium interest rate r clears the capital market, while the outer loop adjusts τ_c until the government's budget constraint is satisfied. The procedure is further described in Appendix C.2.

5.3 Simulation Results

New Tax Regimes

The set of simulated tax regimes is characterised by a fundamental pattern: progressivity in labor income taxes and progressivity in consumption taxes move in opposite directions. As more (less) curvature in tax curve implies more (less) redistribution and therefore lower (higher) tax revenue, the curvature parameter of the other progressive tax adjusts in the opposite way to compensate, ensuring budget neutrality. Table 3 summarises descriptive statistics for the baseline tax regime and the two extrema of the tax regime simulation exercise, i.e., the revenue-neutral tax regime with the lowest labor income tax progressivity and the highest consumption tax progressivity ("Augmented Consumption Progressivity" regime, ACP from now onwards) and the revenue-neutral tax regime with the highest labor income tax progressivity and the lowest consumption tax progressivity ("Augmented Labor Progressivity" regime, ALP from now onwards). In our calibration, the progressivity parameters in the two extrema correspond respectively to the two pairs ($\tau_y = 0.15, \tau_c = 0.066$) and ($\tau_y = 0.225, \tau_c = -0.001$).

Indicator	**Augmented Consumption Progressivity**	**Baseline**	**Augmented Labor Progressivity**
Government Expenditure	1.000	1.000	1.000
Labor Income Tax Progressivity - τ_y	0.15	0.1875	0.225
Consumption Tax Progressivity - τ_c	0.066	0.0343	-0.001
AETR (%) - Labor Income (Bottom 50%, Top 10%)	(5.0, 30.4)	(-1.6, 31.0)	(-8.6, 31.6)
AETR (%) - Consumption (Bottom 10%, Top 10%)	(2.9, 12.6)	(8.9, 13.7)	(15.2, 15.0)
Consumption Tax Effective Rates (%) (Min, Max)	(-2.0, 25.3)	(6.6, 20.6)	(14.8, 15.2)
Labor Taxes (% of Revenue)	63.46	57.13	50.34
Consumption Taxes (% of Revenue)	23.07	29.09	35.59
Capital Taxes (% of Revenue)	13.46	13.79	14.07

Table 3: Summary of Tax Regimes — Revenue-Neutral Extrema

Notes: *Augmented Consumption Progressivity* denotes the revenue-neutral tax regime with the highest (lowest) consumption (labor income) tax progressivity, defined by the pair $\tau_y = 0.150, \tau_c = 0.066$. *Augmented Labor Progressivity* denotes instead the revenue-neutral tax regime with the highest (lowest) labor income (consumption) tax progressivity, defined by the pair $\tau_y = 0.225, \tau_c = -0.001$. For readability, consumption tax rates are here computed following the standard standard VAT reporting, which expresses rates as a fraction of the pre-tax base instead of total consumption expenditures.

The ACP regime features a highly progressive consumption tax, with rates spanning from -2.0% to 25.3%, while the ALP regime is characterised by a slightly regressive consumption tax design, where the bottom 10% of the distribution pays an average effective VAT rate of 15.2%, 0.2% higher than the top 10%. Tax rates on labor income variate on average only to a limited extend for the top 10% of the distribution across the two regimes, whilst going from an AETR of 5.0% to -8.6% for the poorest half. Figure A2 reports detailed average effective and marginal tax rates for the ACP and the ALP tax regimes.

Rebalancing progressivity also affects the composition of tax revenue, as more progressivity implies smaller revenue for the government with the current household distribution. As a result, labor income taxes make almost two thirds of total revenue in the ACP scenario (63.46%) while only half in the ALP scenario (50.34%). As the changes in capital return tax revenue are only due to changes in the aggregate supply of capital and in the equilibrium interest rate, consumption taxes compensate almost all of the variation in revenue from labor income taxes by construction, achieving budget neutrality.

Overall, the ACP scenario results in a pronounced shift in the tax burden from the top 10% of taxpayers to the bottom 90%, effectively redistributing liability downward. Under ACP, the top decile contributes 37.4% of total tax revenue — roughly 1 percentage point less than in the baseline — while the bottom deciles offsets with an increased share. In contrast, the ALP regime exhibits the opposite pattern: a redistribution of the tax burden away from the bottom 90% toward the top 10%, though the concentration of change in the top decile is slightly less pronounced than under ACP.

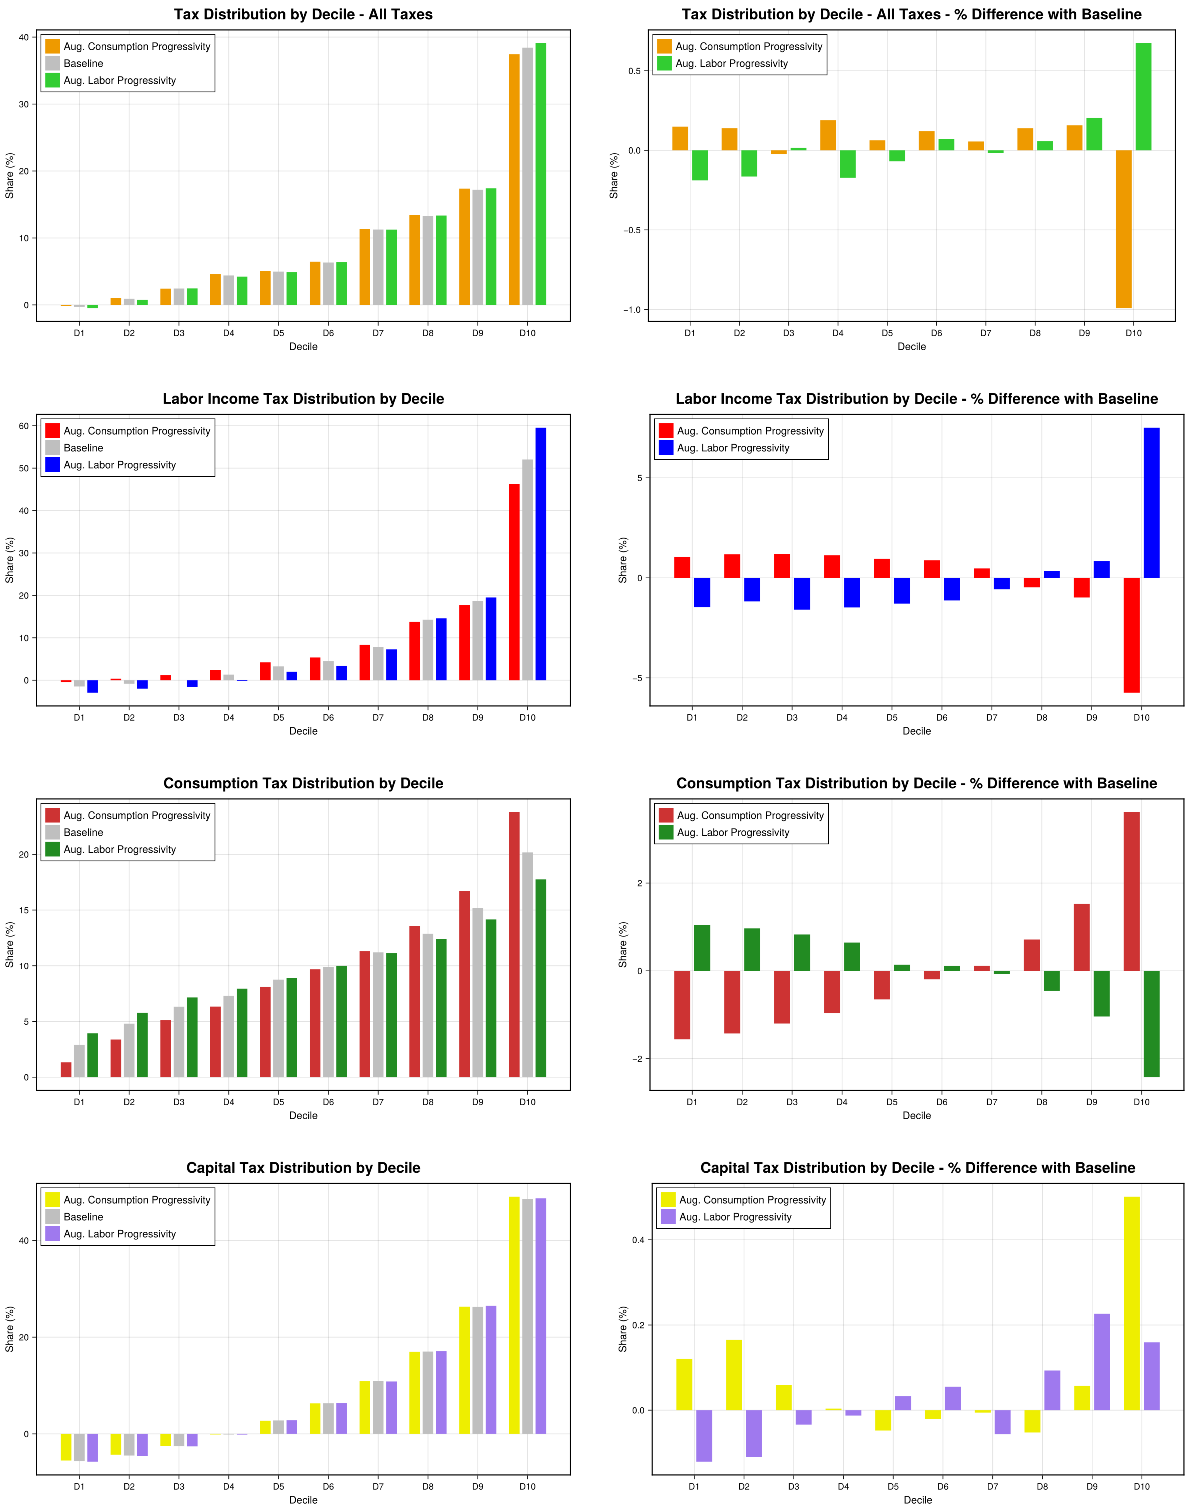

Figure 2: Decile distribution of tax liability: levels (left) and differences from baseline (right). Rows correspond to all taxes, wage taxes, consumption taxes and capital return taxes.

These changes are primarily driven by differences in labor income tax progressivity. In the ACP regime, reducing labor tax progressivity leads to a shift of more than 7 percentage points of total tax liability from the top 30% to the bottom 70% of the income distribution. While consumption taxes adjust in the opposite direction to maintain revenue neutrality, their impact on distribution is more limited. This reflects both their smaller share in total revenue — 29% for consumption taxes versus 57% for labor taxes in the baseline — and the more compressed distribution of consumption relative to labor income. Furtherly, changes in tax curvature produce far more concentrated effects in the labor income tax schedule than in the consumption tax schedule. This asymmetry stems from the fact that pre-tax labor income is significantly more unequal and exhibits a far more pronounced right skew than consumption.

These two factors — the share of government revenue generated by the tax and the degree of pre-existing inequality in the distribution of the tax base - emerge as the key determinants for the redistributive potential of the fiscal instrument. Taxes that finance a larger portion of public expenditure naturally exert a stronger influence on the post-tax distribution, while those applied to more unequal bases allow for greater redistribution when their progressivity is adjusted. In this context, labor income taxation stands out as the best suited tool for redistribution: it represents the principal source of government revenue and is levied on a highly concentrated and right-skewed income distribution. By contrast, consumption taxation contributes less to aggregate revenue and is applied to a base that is more evenly distributed across households, limiting its effectiveness in mitigating inequality. As a result, changes in labor income tax progressivity exert greater influence on both the distribution of resources and macroeconomic outcomes. These structural features explain why labor income taxation remains the primary lever for redistribution in the model, and why it plays a dominant role in shaping the equilibrium results discussed in the next section.

Lastly, changes in capital income tax revenue are minor and mirror endogenous shifts in the distribution of savings across households. Figure 2 displays tax burden levels and reallocations under the ACP and ALP regimes relative to the baseline by tax-paying population decile.

General Equilibrium: Main Economic Outcomes

Figure 3 and Figure 4 illustrate—both in aggregate and in distributional terms—how key general equilibrium outcomes evolve as the tax system is progressively rebalanced away from labor income taxation and toward consumption taxation, transitioning from the Augmented Labor Progressivity (ALP) regime on the right to the Augmented Consumption Progressivity (ACP) regime on the left.

The macroeconomic consequences of shifting progressivity toward consumption taxation (ACP scenario) are uniformly positive. Aggregate output rises steadily, with GDP increasing as high as 1.0% as labor income tax progressivity is reduced relative to the benchmark equilibrium. This expansion is driven by both supply-side and behavioral mechanisms. Lower marginal tax rates on labor income reduce disincentives at the top of the distribution — where labor supply is most elastic — by lowering the opportunity cost of work for high-productivity individuals. As a result, the decline in redistributional distorsions associated with labor income taxation amplifies existing behavioral patterns, with households that already worked the most increasing their labor effort further. Conversely, the bottom 30% of the distribution faces higher effective taxation under the ACP scenario, leading to a contraction in hours worked. Although aggregate hours increase only modestly (+0.2%), a compositional shift in labor supply toward higher-productivity individuals enhances effective labor input, thereby amplifying gains in output.

Under the ACP scenario, private consumption rises by roughly 0.8%, accounting for most of the increase in GDP. This expansion is almost entirely attributable to income effects, as consumption increases primarily among the top half of the distribution — i.e., among those who benefit most from the reduction in labor income tax progressivity. In contrast, the bottom half experiences declining consumption, as the reduction in redistribution through labor taxation is only partially offset by higher consumption tax progressivity.

Since government spending is held constant by construction, the remaining portion of the output gain in the ACP scenario must come from higher investment, fully reflected in the 2.0% rise in aggregate capital. This increase is driven by higher savings from both the wealthiest and the poorest quintiles. While the top-end behavior follows naturally from a rise in disposable income, the increase in savings among poorer households calls for a different explanation. Despite facing lower consumption tax rates under the ACP regime — which reduce the relative cost of consumption with respect to capital accumulation — these households increase their precautionary savings in response to diminished social insurance and greater income risk. Among others, this aligns with the findings of Nishiyama and Smetters (2005), who highlight the role of precautionary motives in environments where lower overall progressivity of the tax design implies reduced social insurance across households . The resulting accumulation of capital further supports gains in output and consumption through standard general equilibrium channels.

Finally, the interest rate declines by approximately 30 basis points, reflecting increased capital supply and the corresponding drop in the marginal product of capital. In parallel, the equilibrium wage rate rises by around 1.5%, supported by higher labor demand and productivity gains from capital deepening.

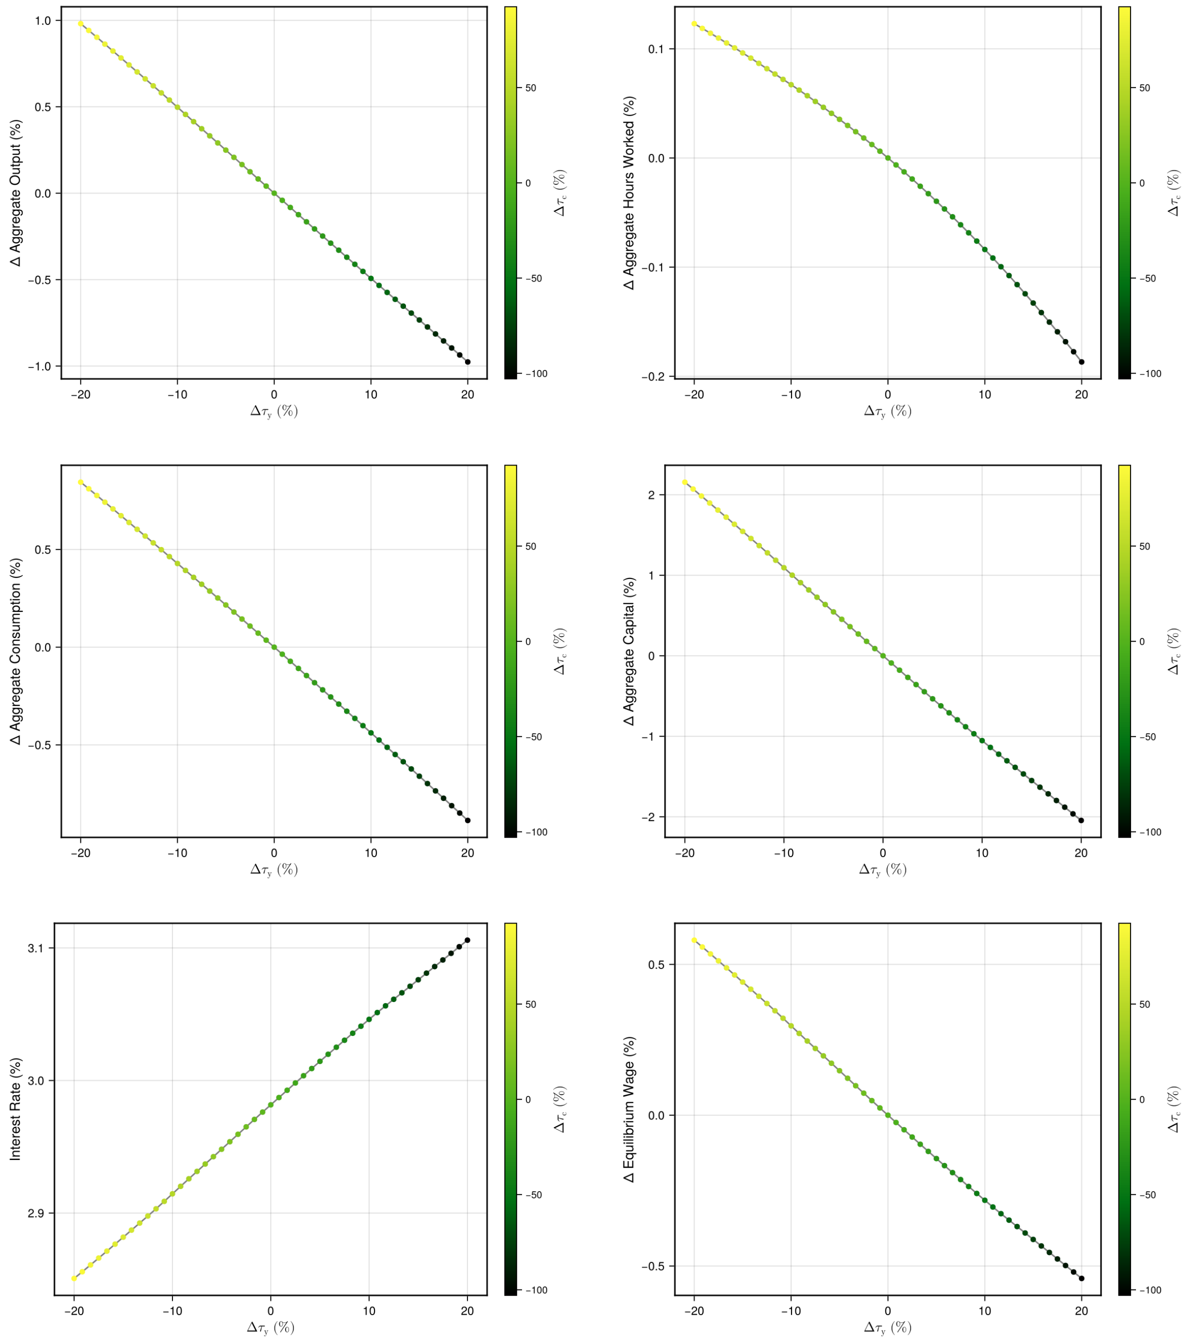

Figure 3: Changes in aggregate outcomes across counterfactual tax regimes: output, hours, consumption, capital, interest rate, and wage.

Notes: All panels report percentage variation across the simulated tax regimes with respect to the baseline level, except for interest rate, which is reported in absolute level for readability.

In sum, a transition towards the ACP regime results in a more efficient allocation of resources, reflected in broad-based improvements in output, capital accumulation, and aggregate consumption. These gains stem primarily from the reduction in labor income tax distortions, particularly at the top of the distribution, which enhances labor supply incentives and supports productivity growth. However, these efficiency improvements are asymmetrically distributed across the population. High-income households disproportionately benefit due to their greater responsiveness to marginal tax reductions, while lower-income households experience losses driven by two reinforcing mechanisms. First, the decline in labor tax progressivity reduces the redistributive capacity of the tax system, shifting the relative burden downward and compressing disposable income at the lower end of the distribution. Second, the erosion of implicit social insurance — previously embedded in the tax-transfer structure — exposes vulnerable households to greater consumption risk. In response, these agents increase precautionary savings and reduce current consumption, despite facing relatively lower consumption tax rates.

This tension between distortionary costs, redistributive effectiveness, and insurance provision lies at the core of the equity-efficiency trade-offs examined in the next section.

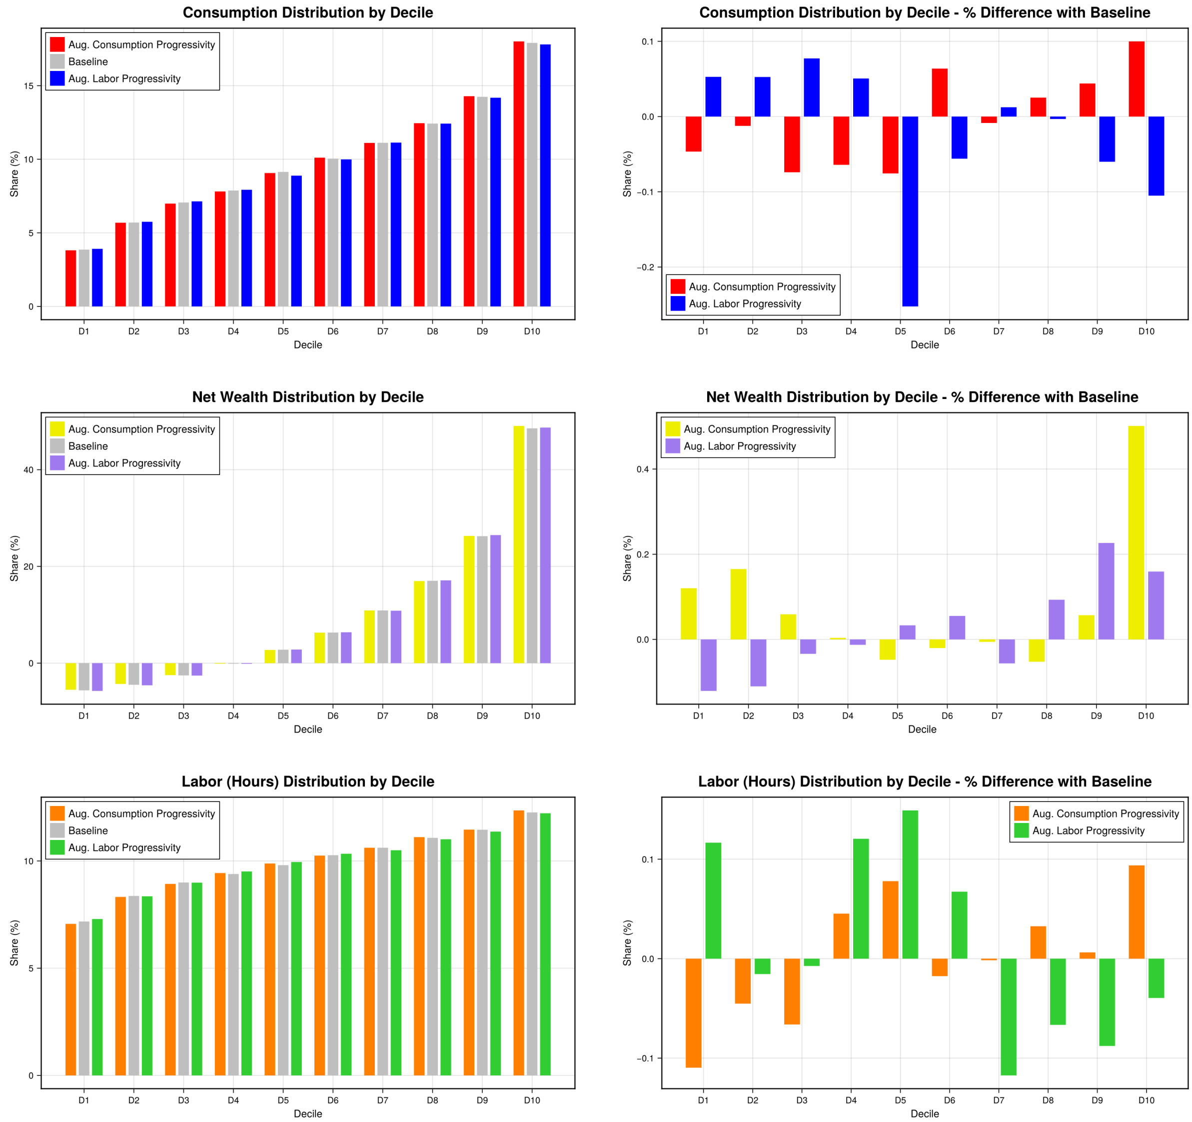

Figure 4: Decile distribution of key household outcomes: levels (left) and differences from baseline (right).

Welfare and Inequality

Table 4 summarises the effects of progressivity rebalancing on inequality and welfare across the baseline, ACP, and ALP regimes. Overall, inequalities rise as labor income tax progressivity decreases, in line with the distributional effects described in Section 5.3.

When decreasing wage tax progressivity, inequalities in pre-tax income increase slightly as a result of the compositional shifts in the labor supply distribution. The rise in inequalities between the ACP regime and the benchmark regime grows larger when analyzing the after-tax income distribution: as the tax design falls short of redistributive power, the Gini coefficient for after-tax income rises from 0.35 to 0.367 (+4.9%).

This increase in inequality is also passed-through to the consumption distribution, as the lack of redistribution is only partially mitigated by the more progressive consumption tax schedule. Moving from the benchmark tax regime to the ACP scenario, the Gini coefficient for consumption increases from 0.228 to 0.231 (+1.3%), reflecting the clear concentration of consumption gains in the top half of the distribution documented in Figure 4. On the other hand, concentration of wealth seems to move slightly in the other direction, as heightened precautionary motives for the poorest fraction of the population might very mildly compress the wealth distribution.

Similar patterns are observed also when considering other inequality measures for both income and consumption and they are more evident in the tails of the distribution - where changes in the tax design curvature adjustments have a stronger impact. As a result, the ratio of average after-tax incomes between the top 10% and the bottom half of the distributions grows from 5.048 to 5.479 (+8.5%). Figure A4 reports full estimations of inequality measures across the 49 tax regimes simulated.

Indicator	**Augmented Consumption Progressivity**	**Baseline**	**Augmented Labor Progressivity**
Gini Coefficient for Pre-Tax Income	0.424	0.422	0.42
Gini Coefficient for After-Tax Income	0.367	0.35	0.334
Gini Coefficient for Consumption	0.231	0.228	0.225
Gini Coefficient for Wealth	0.857	0.859	0.861
Pre-Tax Income - T10/B50 Average Ratio	7.34	7.263	7.183
After-Tax Income - T10/B50 Average Ratio	5.479	5.048	4.652
Consumption - T10/B50 Average Ratio	2.692	2.659	2.624
CEV (%)	0.123	0.0	-0.095

Table 4: Inequality and Welfare across Tax Regimes - Revenue-Neutral Extrema

Notes: *Augmented Consumption Progressivity* denotes the revenue-neutral tax regime with the highest (lowest) consumption (labor income) tax progressivity, defined by the pair $\tau_y = 0.150, \tau_c = 0.066$. *Augmented Labor Progressivity* denotes instead the revenue-neutral tax regime with the highest (lowest) labor income (consumption) tax progressivity, defined by the pair $\tau_y = 0.225, \tau_c = -0.001$.

Welfare outcomes, proxied by consumption-equivalent variation (CEV), highlight the asymmetric distributional consequences of rebalancing tax progressivity across fiscal instruments. Shifting progressivity away from labor income taxation and toward consumption, as in the ACP scenario, yields a modest but unambiguous aggregate welfare improvement of 0.123% relative to the baseline. Conversely, increasing labor income tax progressivity while flattening the consumption tax schedule, as in the ALP scenario, results in a mild welfare loss of approximately 0.10%.

Figure 5 decomposes these aggregate effects by income decile, revealing the granular distribution of welfare gains and losses, which broadly aligns with the distributional dynamics discussed earlier: the bulk of the variation occurs in the upper half of the distribution in both bad and good times. At the bottom, welfare changes are relatively muted and somewhat noisy, reflecting the interplay of competing forces. While the erosion of redistribution in the ACP regime raises the relative tax burden for low-income households, this is partially offset by progressive consumption tax rates and residual insurance effects. In contrast, the upper deciles exhibit more consistent and pronounced shifts. Here, welfare changes are primarily driven by the redistribution channel, but their magnitude reflects the higher behavioral responsiveness of top-income households to marginal tax changes—particularly in labor supply and saving decisions. This drives the outsized contribution of top deciles to the aggregate welfare shifts under both scenarios.

Overall, these results provide several insights on the mechanisms underlying general equilibrium and distributional dynamics triggered by differently-progressive fiscal regimes. In the ACP scenario, reductions in labor tax progressivity reduce disincentives to work and save for high-productivity agents, leading to higher output and capital accumulation. However, the accompanying erosion of redistribution and social insurance harms lower-income households, who face both a higher effective tax burden and

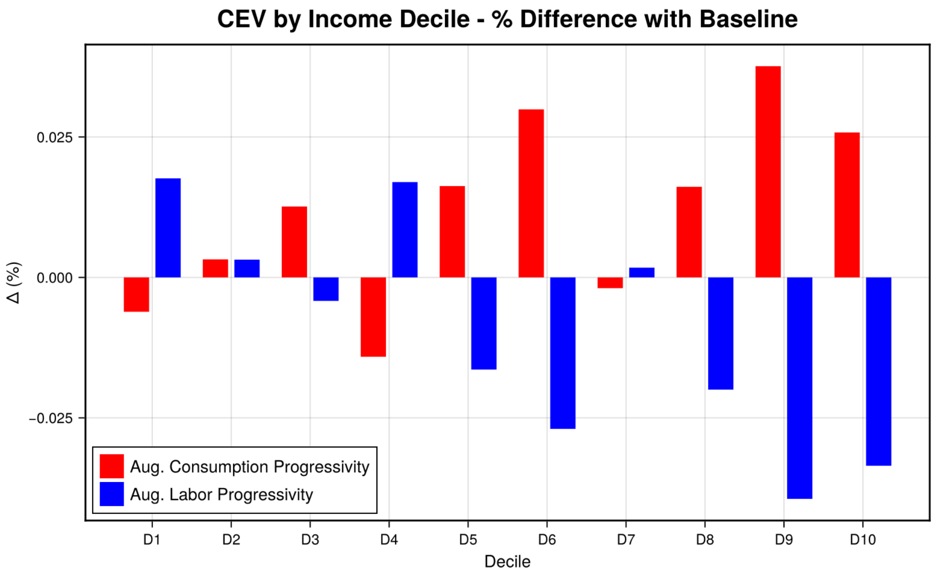

Figure 5: CEV Distribution with respect to Baseline - Revenue-Neutral Extrema

increased income risk. This explains the concentration of welfare losses at the bottom. Conversely, in the ALP regime, the reintroduction of progressivity into labor taxation enhances redistribution and insurance, benefitting poorer households. Yet this comes at the cost of efficiency losses, which are primarily borne by high-income agents through reduced returns to labor and capital.

Taken together, these findings reaffirm the central equity-efficiency trade-off that underlies tax policy design. Flattening labor taxes improves allocative efficiency and stimulates economic activity, but does so by shifting the tax burden downward and weakening both the redistributive function and the economy's capacity to insure against idiosyncratic shocks. In contrast, restoring progressivity through labor taxation enhances equity and consumption smoothing for vulnerable households, but introduces distortions that dampen aggregate performance.

More broadly, the results illustrate the dual role of redistribution and insurance in shaping welfare outcomes. Redistribution determines the allocation of net resources across households with differing lifetime income, while insurance governs intertemporal consumption stability for agents exposed to uncertainty. The welfare losses observed in the lower part of the distribution under the ACP regime stem not only from a decline in relative transfers but also from diminished protection against income shocks—prompting increases in precautionary savings and reductions in current consumption. Similarly, the gains experienced at the top are amplified by both lower average tax burdens and reduced exposure to redistribution.

Ultimately, these findings suggest that the desirability of rebalancing tax progressivity depends not only on its aggregate efficiency effects but also on its implications for equity and risk-sharing. The interplay between these dimensions is particularly salient in economies with incomplete markets and heterogeneous agents, where the social value of insurance and redistribution can be substantial even when they come at the cost of modest efficiency losses.

5.4 Isolating the effect of consumption taxes

To validate and complement the results presented in Section 5.3, I propose two new alternative tax designs to the ACP reform which are both characterised by i) the same low-progressivity labor income tax calibrated in the ACP scenario and ii) revenue-neutrality with respect to the ACP - and therefore to the baseline - regime. The former is built fixing the labor income tax and the same curvature parameter τ_c as in the baseline, adjusting this time the level parameter λ_c to target budget neutrality (labelled as "LC-EQ" scenario). The latter, instead, is constructed by adjusting both the curvature parameter τ_c and the scale parameter λ_c for consumption taxation to match i) budget neutrality and ii) the same share of revenue from consumption observed in the baseline scenario (labelled therefore as "CRev-EQ"). Analyzing the general equilibrium and distributional outcomes of these two new counterfactuals vis-à-vis the ACP equilibrium will help disentangling the effect of consumption tax rebalancing alone, as the three scenarios feature the same labor income tax design and the same overall tax revenue.

Variable	Baseline	ACP	LC-EQ	CRev-EQ
Government Expenditure	1.0	1.0	1.0	1.0
λ_y	0.731	0.731	0.731	0.731
τ_y	0.188	0.150	0.150	0.150
λ_c	0.850	0.850	0.879	0.537
τ_c	0.034	0.066	0.034	0.451
τ_k	0.352	0.352	0.352	0.352
AETR (%) - Labor Income (B50%, T10%)	(-1.6, 31.0)	(5.0, 30.4)	(5.1, 30.5)	(3.5, 29.9)
AETR (%) - Consumption (B10%, T10%)	(8.9, 13.7)	(2.9, 12.6)	(5.7, 10.8)	(-26.3, 30.1)
Consumption Tax Eff. Rates (%) (Min, Max)	(6.6, 20.6)	(-2.0, 25.3)	(3.2, 17.9)	(-56.0, 74.6)
Labor Taxes (% of Revenue)	57.13	63.47	64.12	57.15
Consumption Taxes (% of Revenue)	29.09	23.07	22.41	29.09
Capital Taxes (% of Revenue)	13.79	13.47	13.47	13.76

Table 5: Fiscal Parameters Across Alternative Tax Regimes - Robustness Checks

Table 5 presents the main details of the new tax systems, compared against the baseline and the ACP regimes. Keeping the same curvature as in the baseline, the LC-EQ consumption tax generates average effective rates that are comparable in width range to those of the baseline but lower for absolute level by roughly 3 percentage points, to compensate for the increase in revenue from labor income taxation. The calibration results are broadly comparable to those of the ACP regime in terms of revenue shares, with a limited redistribution of revenue from consumption to labor income taxes. Under the LC-EQ regime, the bottom 47% of the household distribution pays a higher effective consumption tax rate than in the ACP regime. On the other side, the CRev-EQ fitting generates a much more progressive consumption tax schedule, with average effective rates ranging between -56.0% for the poorest household to 74.6% for the richest. Under this regime, the bottom 38% of the households enjoys a (much) lower effective tax rate on their consumption basket. Figure A5 reports the Average Effective Rate and Average Marginal Rate curves under the consumption tax regimes considered.

Indicator	ACP	LC-EQ	CRev-EQ
Δ Aggregate Output (%)	0.98	1.67	-6.0
Δ Aggregate Consumption (%)	0.84	1.73	-7.57
Δ Aggregate Capital (%)	2.16	3.07	-8.07
Δ Aggregate Labor (Hours) (%)	0.12	0.81	-7.02
Δ Equilibrium Wage (%)	0.58	0.68	-1.11
Interest Rate (%)	2.85	2.83	3.24
Gini Coefficient for Pre-Tax Income	0.424	0.424	0.434
Gini Coefficient for After-Tax Income	0.367	0.367	0.376
Gini Coefficient for Consumption	0.231	0.235	0.196
Gini Coefficient for Wealth	0.857	0.853	0.928
Pre-Tax Income - T10/B50 Average Ratio	7.34	7.32	7.701
After-Tax Income - T10/B50 Average Ratio	5.479	5.467	5.713
Consumption - T10/B50 Average Ratio	2.692	2.745	2.328
CEV (%)	0.123	-0.166	1.832

Table 6: Main Economic and Distributional Indicators Across Alternative Regimes Robustness Checks

Notes: Variations, denoted by Δ, are computed with respect to the baseline steady state.

Table 6 summarises the main economic and distributional indicators for the new counterfactual scenarios, while Figure 6 reports the decomposition of relevant main changes by population decile. The first insight from the table is a negative correlation between progressivity in consumption taxes and main economic outcomes when holding the labor income tax schedule constant: moving from the ACP to the LC-EQ scenario, which entangles a consumption tax that is less redistributive and less confiscatory on average, the economy benefits from sizeable gains in consumption (+0.88%), capital (+0.89%), labor hours supplied (+0.69%) and eventually output (+0.68%). On the contrary, a much more progressive and confiscatory consumption tax design depresses output by 6 percentage points holding the labor tax schedule constant, with consumption, assets and labor supply falling even more. In line with earlier results, these aggregate changes are mainly driven by the upper half of the population and underpin the essential redistribution-efficiency trade-off.

Switching from the ACP regime to the LC-EQ one, households in the top decile of the distribution substitute part of their savings with consumption, as it is now cheaper. The opposite takes place on the other end, as less redistribution implies not only higher consumption tax rates for the poorest households but also a lower degree of risk sharing in the economy with a consequent higher demand for precautionary savings. The loss of insurance against consumption risk is reflected also in the welfare losses experienced by the bottom 40% of the distribution under the LC-EQ regime, which more than offset the gains in the upper tail, driving an aggregate decrease in well-being. As for inequalities, the sum of the underlined behavioral responses compresses the distribution of wealth on the one hand and widens the consumption distribution on the other.

In the CRev-EQ scenario, a greatly more redistributive consumption tax schedule not only heavily redistributes consumption from the upper to the lower half of the distribution, but also reduces the need for precautionary savings for the bottom 70% of the population. On the other hand, as households in the top deciles of the distribution face confiscatory rates on consumption, they increase their appetite for asset holdings - which become much more concentrated - and leisure, as a result of a two-folded substitution effect. Overall, the high welfare gains enjoyed by the bottom 60% of the distribution due to an increase in risk-sharing - with a 1% peak for the poorest decile - more than offset the welfare losses faced by the top three deciles, resulting in a positive aggregate welfare gain equivalent to a 1.83% lifetime increase in consumption.

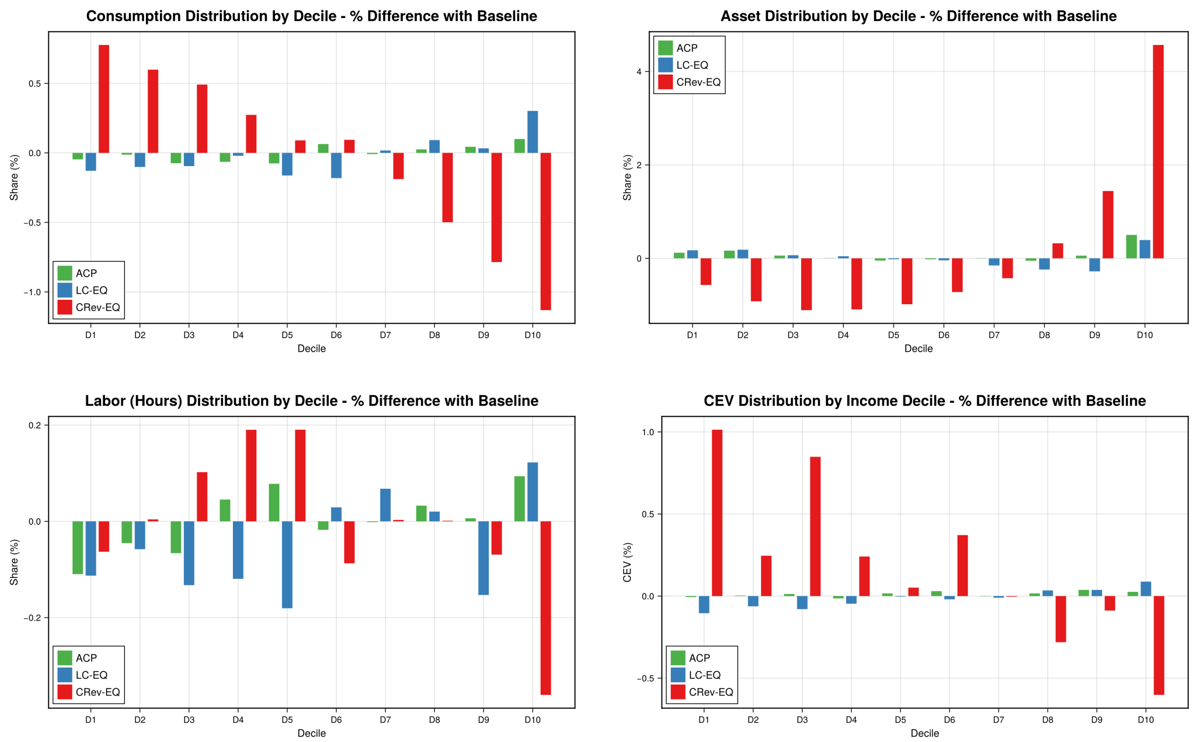

Figure 6: Changes in consumption, asset, labor hours supply and CEV distributions

Notes: Income distribution is used as a reference for the distribution of CEV gains/losses.

Overall, risk sharing emerges as a key driver of welfare in the economy: reducing labor income tax progressivity results in reduced redistribution which is accompanied by a weakening of the economy's insurance function. As public transfers or implicit insurance mechanisms shrink, households — especially those facing higher income volatility — experience increased consumption risk and respond by raising precautionary savings, further depressing consumption. This downward adjustment is driven both by

income effects, as higher rates reduce now households resources, and substitution effects, as increased volatility raises marginal utility of future consumption.
However, this loss in redistribution and social insurance can be partially offset by the social planner through an increase in the progressivity of consumption taxes. When doing so, as in the ACP scenario, welfare losses for the bottom half of the distribution are offset and total welfare grows as a result of the gains for the richest. In contrast, revenue-neutral reforms that do not adjust consumption tax progressivity - as in the case of the LC-EQ regime - fail to compensate for the loss of insurance, leading to a contraction in aggregate welfare.

In this sense, progressivity in consumption taxes emerges as a potentially powerful fiscal instrument. By tilting the tax burden toward higher-consuming households, it creates space for reducing distortions in labor income taxation — particularly for the top half of the distribution — without sacrificing redistribution or aggregate welfare. This rebalancing allows the tax system to incentivize labor supply and capital accumulation, generating productivity gains and welfare improvements at the top, while preserving or even enhancing the welfare of lower-income households. Importantly, it achieves this not through explicit transfers, but by embedding implicit redistribution and insurance within the tax schedule itself. In doing so, progressive consumption taxation acts as a second-best policy tool: it indirectly compensates for lost redistribution from flatter labor taxes, supports risk-sharing by reducing tax pressure in periods of low consumption and maintains aggregate efficiency by minimizing disincentives to work or save. As such, it can be seen as a flexible and politically feasible complement to traditional progressive income taxation in constrained policy environments.

5.5 Further Steps

In addition to an overall improvement in the robustness of the results presented, a number of extensions could enrich the analysis and support a more comprehensive approach to optimal fiscal policy involving progressive consumption taxation.

First, a natural step is to explicitly incorporate social security contributions, which represent a substantial portion of labor taxation in most advanced economies. Their interaction with labor income taxes can amplify distortions in work and saving decisions, and their exclusion may understate the total burden on earnings—particularly at the lower end of the income distribution. Including them would allow for a more accurate assessment of the marginal cost of public funds and the equity-efficiency trade-offs inherent in the broader tax-transfer system.
Second, while this paper focuses on steady-state comparisons, a more complete evaluation of tax reforms should also account for transitional dynamics. Policy changes can entail significant short- and medium-run costs, especially for households that are far from their long-run optimal allocations. Capturing these transition effects is crucial for evaluating the political and welfare viability of reform paths.
Third, embedding the analysis in a life-cycle framework would offer deeper insight into the intertemporal and intergenerational implications of progressive consumption taxation. Among other reasons, the intergenerational transfer of resources from a progressivity rebalancing in such a setting could entail interesting insights. As taxing consumption shifts the burden toward older cohorts and retirees and away from young, saving workers, reforms may entail redistributive consequences across cohorts — an aspect of particular policy interest in the context of ageing societies. Incorporating such dimensions would allow for a richer understanding of how tax progressivity should be allocated not only across income levels, but also across time and demographic groups.

6 Conclusions

This paper investigates how shifting the composition of tax progressivity between labor income and consumption taxation affects macroeconomic efficiency, inequality, and welfare in a heterogeneous-agent economy with incomplete markets. Using a quantitative model calibrated to the French economy, I simulate a series of revenue-neutral fiscal regimes in which progressivity in labor income taxation is exogenously adjusted and compensated by offsetting changes in the progressivity of consumption taxation. This setup enables a structured analysis of "progressivity rebalancing" reforms and consequent implications for aggregate outcomes and household-level welfare.

The findings reveal three key insights. First, the redistributive capacity of a tax is closely tied to the characteristics of its base. Labor income taxes are found to be significantly more powerful as redistributive instruments than consumption taxes, due both to the larger revenue share they represent and the higher inequality of the labor income distribution. Second, welfare effects of changing tax progressivity operate through both redistribution and risk-sharing channels. While high-income households mainly respond through income effects, lower-income households are especially sensitive to the implicit insurance embedded in the tax-transfer system. When redistribution is reduced, these households face higher income volatility and increase precautionary savings, leading to lower consumption and welfare. Third, progressive consumption taxation can function as a flexible second-best policy tool. By shifting some of the redistributive burden away from labor income taxes—particularly at the top of the distribution—it enables efficiency gains without incurring proportional costs in terms of equity or welfare. Importantly, it does so without requiring explicit transfers, embedding insurance and redistribution directly into the tax system.

From a policy perspective, these results have clear relevance. In settings where increasing labor income tax progressivity is politically or institutionally constrained—or where reducing it is desirable for efficiency—progressive consumption taxation offers a flexible alternative. By shifting part of the redistributive burden away from labor income, it enables efficiency gains while maintaining equity and insurance through the tax system itself. Differential VAT rates, for instance, are already a prominent feature of European tax systems. Enhancing their redistributive orientation—either explicitly through progressive consumption taxes or implicitly through smarter rate structures—can help governments preserve redistribution and insurance while minimizing economic distortions. Moreover, in an era of rising inequality and growing demands on public insurance mechanisms, understanding how tax design affects both efficiency and consumption risk is essential.

In spite of not seeking to provide definitive prescriptions for optimal taxation, this paper offers a structured framework and first set of insights that future work can build upon. In particular, future research should consider the inclusion of social security contributions, transitional dynamics, and intergenerational effects in life-cycle settings to further assess the long-run desirability and incidence of progressive consumption tax reforms from different angles. As fiscal systems increasingly face constraints on both equity and efficiency grounds, rethinking the composition of progressivity across tax instruments will remain a critical area of inquiry.

References

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

Bachas, P., Gadenne, L., & Jensen, A. (2024). Informality, Consumption Taxes, and Redistribution. *The Review of Economic Studies*, *91*(5), 2604–2634. https://doi.org/10.1093/restud/rdad095

Benzarti, Y., Carloni, D., Harju, J., & Kosonen, T. (2020). What Goes Up May Not Come Down: Asymmetric Incidence of Value-Added Taxes. *Journal of Political Economy*, *128*(12), 4438–4474. https://doi.org/10.1086/710558

Blasco, J., Guillaud, E., & Zemmour, M. (2023). The inequality impact of consumption taxes: An international comparison. *Journal of Public Economics*, *222*, 104897. https://doi.org/10.1016/j.jpubeco.2023.104897

Bozio, A., Garbinti, B., Goupille-Lebret, J., Guillot, M., & Piketty, T. (2020). *Predistribution vs. Redistribution: Evidence from France and the US* (Working Paper No. 2020-24). Center for Research in Economics and Statistics.

Cahn, C., Fève, P., & Matheron, J. (2024). The Long-Run Effects of Fiscal Rebalancing in a Heterogeneous-Agent Model. https://doi.org/10.2139/ssrn.4886600

Coleman, W. J. (2000). Welfare and optimum dynamic taxation of consumption and income. *Journal of Public Economics*, *76*(1), 1–39. https://doi.org/10.1016/S0047-2727(99)00043-2

Correia, I. (2010). Consumption Taxes and Redistribution. *American Economic Review*, *100*(4), 1673–1694. https://doi.org/10.1257/aer.100.4.1673

Direction Générale des Finances Publiques (DGFiP). (2024, April). *L'impôt sur les revenus perçus en 2022* (tech. rep.). Paris. https://www.impots.gouv.fr/sites/default/files/media/9_statistiques/0_etudes_et_stats/0_publications/dgfip_statistiques/2024/num22_04/dgfip_stat_22_limpot-sur-le-revenu-en-2022.pdf

European Commission. (2024). Ameco database [Accessed: 2025-05-16].

Feldstein, M. S. (1969). The Effects of Taxation on Risk Taking. *Journal of Political Economy*, *77*(5), 755–764. Retrieved January 9, 2025, from https://www.jstor.org/stable/1829965

Ferriere, A., Grübener, P., Navarro, G., & Vardishvili, O. (2023). On the Optimal Design of Transfers and Income Tax Progressivity. *Journal of Political Economy Macroeconomics*, *1*(2), 276–333. https://doi.org/10.1086/725034

French Government. (2025). Taxe sur la valeur ajoutée (TVA). https://entreprendre.service-public.fr/vosdroits/N13445

Heathcote, J., Storesletten, K., & Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework*. *The Quarterly Journal of Economics*, *132*(4), 1693–1754. https://doi.org/10.1093/qje/qjx018

Institute for Fiscal Studies (IFS). (2011). Quantitative analysis of vat rate structures. In Institute for Fiscal Studies et al. (Eds.), *A retrospective evaluation of elements of the eu vat system* (pp. 361–443). Report prepared for the European Commission, TAXUD/2010/DE/328.

Kopecky, K. A., & Suen, R. M. H. (2010). Finite state Markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, *13*(3), 701–714. https://doi.org/10.1016/j.red.2010.02.002

Li. (2013). Welfare Effects of Consumption Taxes [Job Market Paper].

Macnamara, P., Pidkuyko, M., & Rossi, R. (2025). Taxing Consumption in Unequal Economies.

Nishiyama, S., & Smetters, K. (2005). Consumption Taxes and Economic Efficiency with Idiosyncratic Wage Shocks. *Journal of Political Economy*, *113*(5), 1088–1115. https://doi.org/10.1086/432137

O'Donoghue, C., Mantovani, D., & Baldini, M. (2004). Modelling the redistributive impact of indirect taxes in Europe: An application of EUROMOD [Number: EM7/01 Publisher: EUROMOD at the Institute for Social and Economic Research]. *EUROMOD Working Papers*. Retrieved January 31, 2025, from https://ideas.repec.org//p/ese/emodwp/em7-01.html

OECD. (2024). Consumption Tax Trends 2024. Retrieved May 11, 2025, from https://www.oecd.org/en/publications/consumption-tax-trends-2024_dcd4dd36-en.html

Parodi, F. (2023). Taxation of Consumption and Labor Income: A Quantitative Approach. *American Economic Journal: Macroeconomics*, *15*(4), 177–216. https://doi.org/10.1257/mac.20200289

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In T. F. Cooley (Ed.), *Frontiers of business cycle research* (pp. 294–330). Princeton University Press.

Shaw, K. (2020). Are Consumption Taxes Really Better Than Labour Income Taxes? Theoretical and Quantitative Implications of the Choice of Tax Base.

Thomas, A. (2022). Reassessing the regressivity of the VAT. *Fiscal Studies*, *43*(1), 23–38. https://doi.org/10.1111/1475-5890.12290

World Bank. (2024). World development indicators: Government expenditure [Accessed: 2025-05-16].

World Inequality Database. (2024). World inequality database [Accessed: 2024-05-16].

A Main Appendix

A.1 Differential VAT rates

Good/Service	Category	VAT Rate
Therapeuthic medical services	Health & Beauty	0%
Public Education Services	Education	0%
Newspapers & Periodicals	Arts & Press	2.1%
Reimbursed Medicines	Health & Beauty	2.1%
Basic Food Products	Food & Beverages	5.5%
Standard Home Electricity	Energy	5.5%
Books	Arts & Press	5.5%
Renewable Energy Heating Systems	Energy	5.5%
Works of Art	Arts & Press	5.5%-20%
Public Transport	Services	10%
Cinema	Sport & Culture	10%
Forestry Services	Agriculture	10%
Restaurant Meals (Non-Alcoholic)	Hospitality	10%
Hotel Accommodations	Hospitality	10%
Fossil Fuel Heating Systems	Energy	20%
Alcoholic Beverages	Food & Beverages	20%
Sport Facilities' Access	Sport & Culture	20%
Cosmetics	Health & Beauty	20%
Private Education Services	Education	20%
Luxury Goods (e.g., Jewelry)	Retail	20%

Table A1: Selected VAT Rates in Metropolitan France by Product/Service as of April 2025 (French Government, 2025).

A.2 Baseline Model

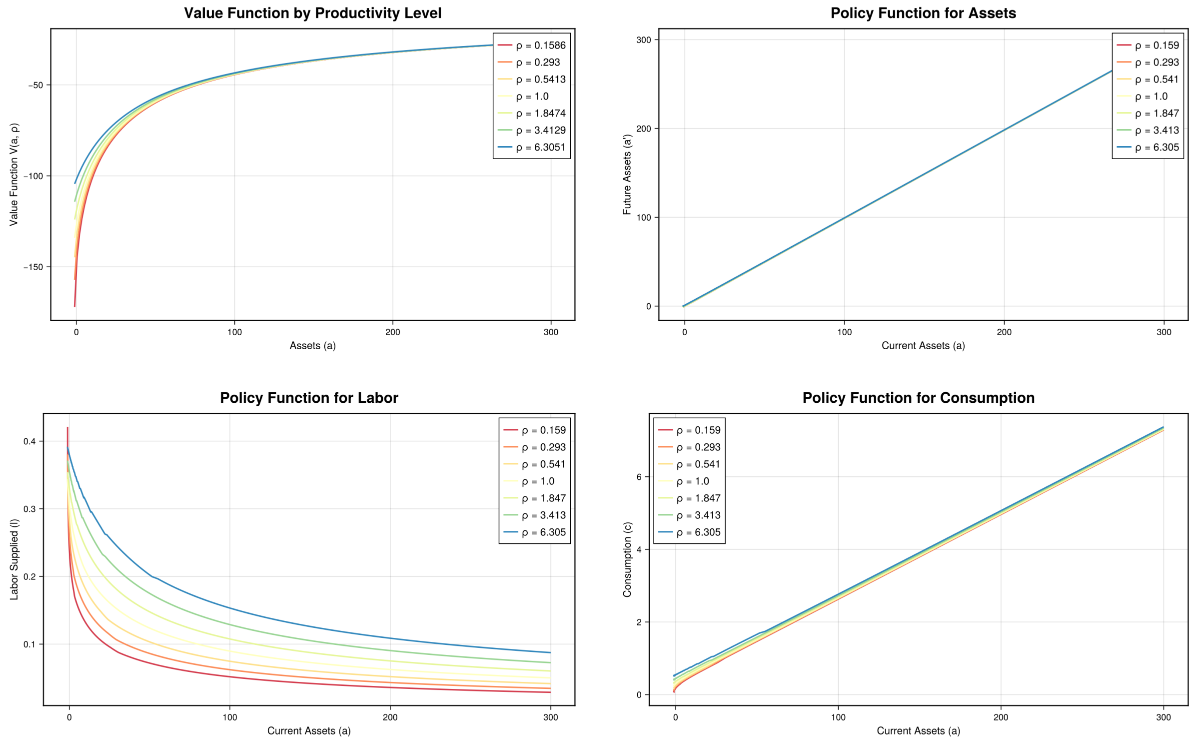

Figure A1: Value Function and Policy Functions for Assets, Labor and Consumption by Productivity Level - Baseline Calibration

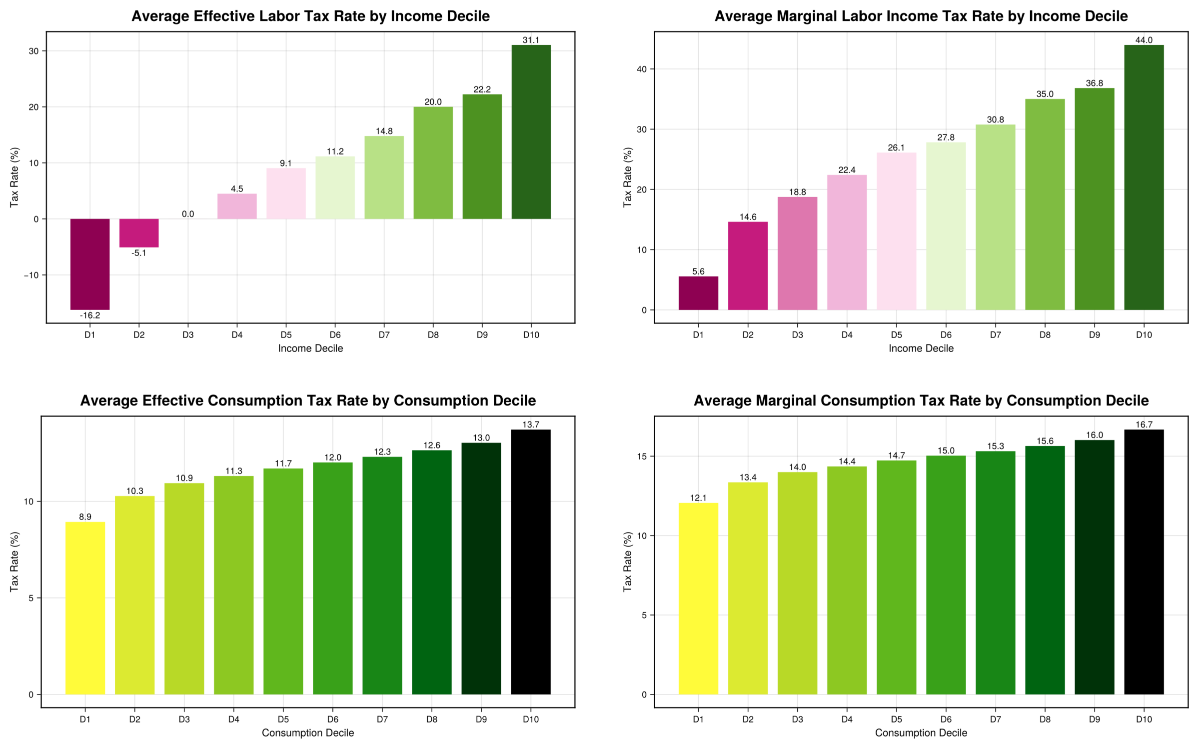

Figure A2: Decile Profiles of Effective and Marginal Tax Rates by Tax Type - Taseline Calibration

Notes: For readability, consumption tax rates are here computed following the standard standard VAT reporting, which expresses rates as a fraction of the pre-tax base instead of total consumption expenditures.

A.3 Simulation Results

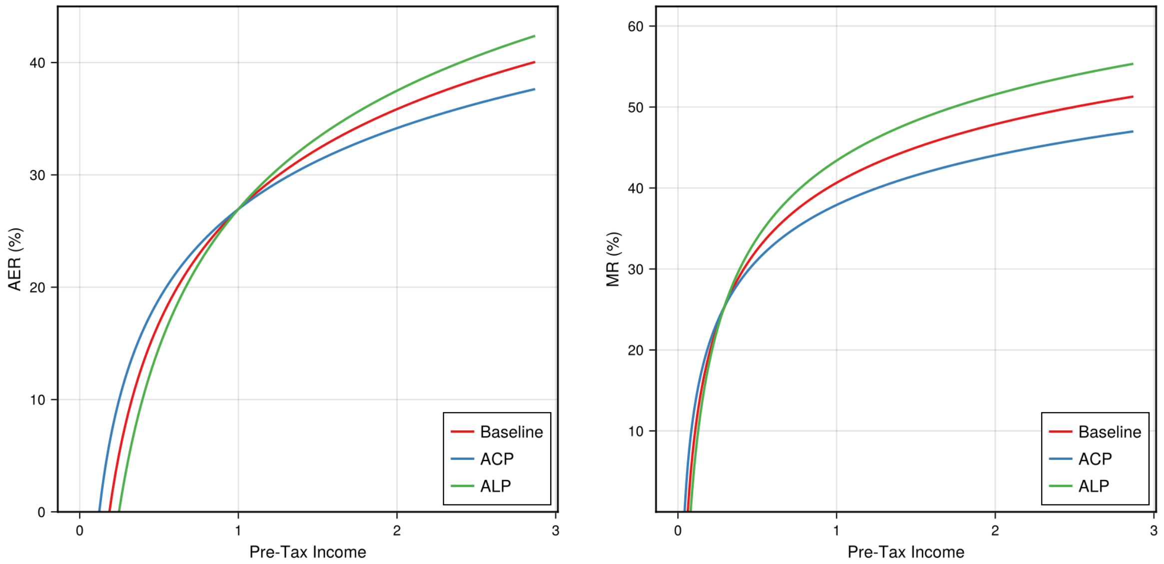

Figure A3: Average Effective Rates (AER) and Marginal Rates across Labor Income Tax Regimes

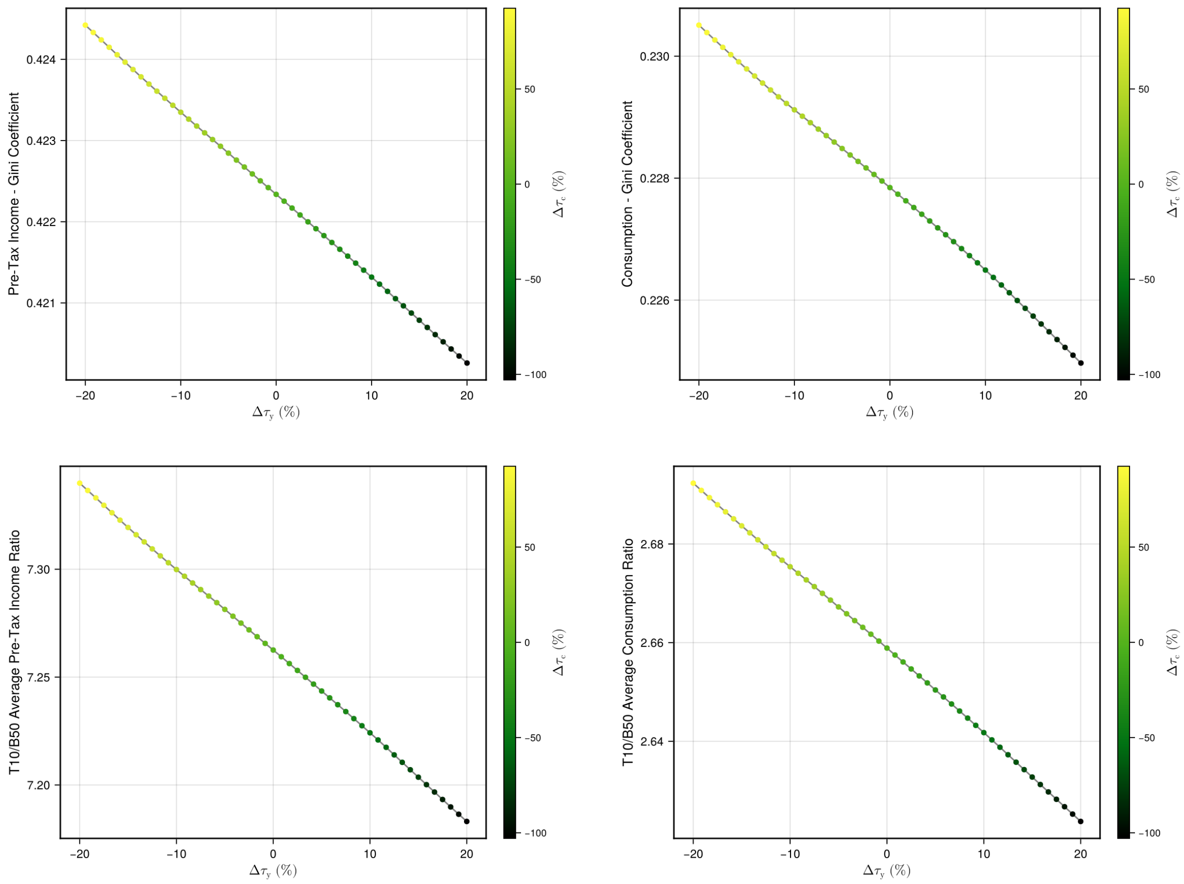

Figure A4: Inequality Measures for Income (left) and Consumption (right) across Counterfactual Regimes: Gini Coefficient and Top-10%-Bottom-50% Average Level Ratio.

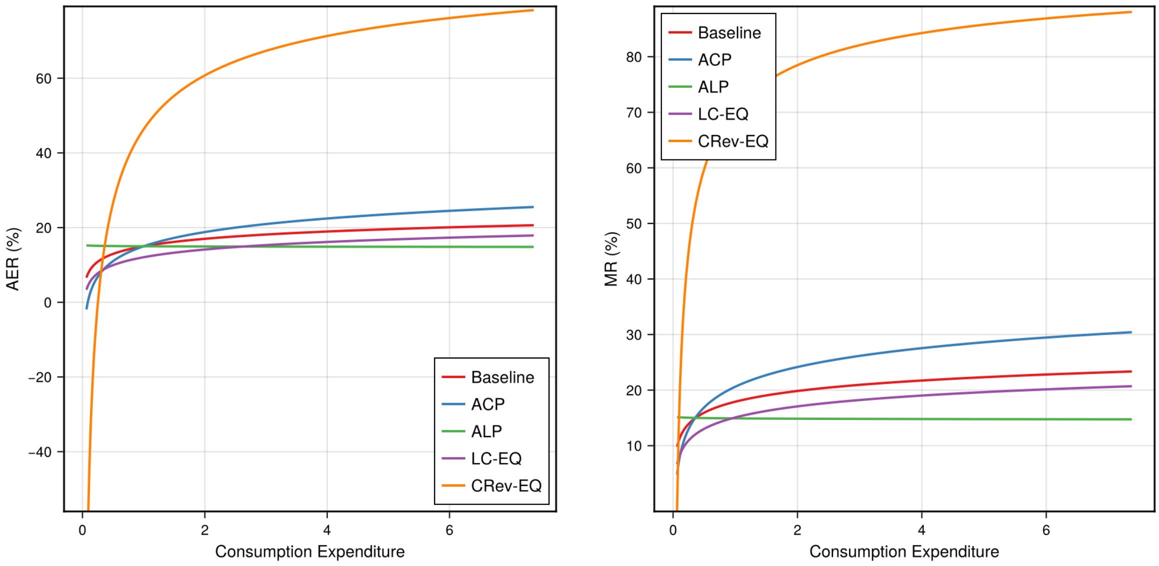

Figure A5: Average Effective Rates (AER) and Marginal Rates across Consumption Tax Regimes

Notes: Computed crossing points for AER curves: i) ACP/LC-EQ curves: 0.342; ii) ACP/CRev-EQ curves: 0.303. Through the use of the stationary distribution of households and the policy matrix for consumption, we can derive the share of households below the crossing points: i) ACP/LC-EQ: 0.473; ii) ACP/CRev-EQ: 0.383.

B Analytical Model

Household's recursive problem

The problem of the household can be written as

$$\max_{\{c_t,h_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \mathbb{E}_t[u(c_t,h_t)]$$

with

$$u(c_t,h_t) = \frac{c_t^{1-\sigma}}{1-\sigma} - B\frac{h_t^{1+\psi}}{1+\psi}$$

subject to the budget constraint

$$c_t + a_{t+1} = y_t - T_y(y_t) + (1+(1-\tau_k)r_t)a_t - T_c(c_t)$$

$$\text{with} \quad y_t = \rho_t w_t h_t,$$

$$T_z(z_t) = (1 - \frac{\lambda_z}{z^{\tau_z}})z = z_t - \lambda_z z^{1-\tau_z} \quad \text{for} \quad z \in \{y,c\},$$

the feasibility constraints

$$c_t \geq 0, \quad h_t \in [0,1] \quad ,$$

and the credit constraint

$$a_t \geq -b.$$

Recursively, this can be otherwise stated as

$$V(a,\rho) = \max_{c,h,a'} [u(c,h) + \beta \mathbb{E}_{\rho'} V(a',\rho')] \text{ with } u(c,h) = \frac{c^{1-\sigma}}{1-\sigma} - B\frac{h^{1+\psi}}{1+\psi} \tag{B.1}$$

subject to:

$$c + T_c(C) = y - T_y(y) + (1+(1-\tau_k)r)a - a'$$

$$a' \geq -b$$

$$c \geq 0, \quad h \in [0,1].$$

Solving analytically the recursive problem with respect to h and a', the following FOCs can be derived

$$\frac{\partial c}{\partial h} c^{-\sigma} - Bh^{\psi} = 0 \tag{B.2}$$

$$\frac{\partial c}{\partial a'} c^{-\sigma} + \beta \mathbb{E}_{\rho'} \left[\frac{\partial V(a',\rho')}{\partial a'} \middle| \rho \right] = 0. \tag{B.3}$$

As c appears twice in the budget constraint and the taxation function is not analytically invertible, we exploit implicit differentiation to find

$$\frac{\partial c}{\partial h} = \frac{\lambda_y(1-\tau_y)(\rho w)^{1-\tau_y}}{2-\lambda_c(1-\tau_c)c^{-\tau_c}} h^{-\tau_y} \tag{B.4}$$

$$\frac{\partial c}{\partial a} = \frac{1+(1-\tau_k)r}{2-\lambda_c(1-\tau_c)c^{-\tau_c}} \tag{B.5}$$

$$\frac{\partial c}{\partial a'} = -\frac{1}{2-\lambda_c(1-\tau_c)c^{-\tau_c}}. \tag{B.6}$$

Substituting (B.4) in the FOC for labor (B.2), I derive the optimality condition for labor h with respect to consumption c which I will later exploit in the numerical solution

$$h_{opt} = \left[\frac{\lambda_y(1-\tau_y)(\rho w)^{1-\tau_y}}{B(2-\lambda_c(1-\tau_c)c^{-\tau_c})} c^{-\sigma} \right]^{\frac{1}{\psi+\tau_y}} . \tag{B.7}$$

Starting from the FOC for assets (B.3), exploiting equations (B.6) and (B.5), together with the Envelope Theorem, we derive the Euler Equation for Consumption

$$\frac{\partial u(c,\ell)}{\partial c} = \beta \mathbb{E}_{\rho'} \left[\frac{\frac{\partial (c+T_c(c))}{\partial c}}{\frac{\partial (c'+T_c'(c'))}{\partial c'}} (1+\tilde{r}') \frac{\partial u(c',\ell')}{\partial c'} \right] \tag{B.8}$$

i.e.,

$$c^{-\sigma} = (2-\lambda_c(1-\tau_c)c^{-\tau_c})\beta \mathbb{E}_{\rho'} \left[\frac{1+\tilde{r}'}{2-\lambda_c'(1-\tau_c')c'^{-\tau_c}} c'^{-\sigma} \right] \tag{B.9}$$

C Numerical Solution

To yield the results in the paper, three steps are performed: i) resolution of the baseline model, ii) equivalent regime search and iii) extended equivalent regime search (robustness check). The following section describe each of these processes.

C.1 Baseline Model

Solving the baseline model requires determining the steady-state real interest rate r that clears the capital market. I proceed as follows.

1. **Grids**. Generate grids for productivity ρ and assets a. Productivity levels are pinned down using the Rouwenhorst (1995) discretisation method and the transition matrix $\pi(\rho,\rho')$ is found via fixed point iteration. The grid for assets is constructed using a polynomial transformation, which increases the spacing between grid points as asset levels rise. Specifically, the asset grid is defined as:

   $$a_i = \underline{a} + (\overline{a} - \underline{a}) \cdot \left(\frac{i}{N_a - 1} \right)^{\gamma}, \quad i = 0,1,\ldots,N_a - 1,$$

   where $\underline{a}$ and $\overline{a}$ denote the minimum and maximum asset levels, $N_a = 400$ is the number of grid points, and $\gamma > 1$ controls the curvature of the grid (with $\gamma = 4$ in the baseline). This concentrates points near the borrowing constraint and spreads them out in the upper tail.

   $\underline{a}$ is endogenously determined to avoid the insolvency trap, i.e., the possibility that households cannot repay interests on their debt. Therefore, I conservatively assume that even in the case of zero taxes (and therefore zero subsidies), households with the lowest productivity $\rho_{\min}$ must be able to repay the proceedings on their debt for a large interest rate $\overline{r}$ and a large supply of labor $\overline{h}$. Having derived $\rho_{\min} = 0.1586$ and set $\overline{r} = 0.10$ and $\overline{h} = 0.8$, I derive the borrowing constraint from the household budget constraint for the baseline calibration as follows

   $$\underline{a} = \frac{\rho_{min}\overline{h}w(\overline{r})}{\overline{r}} = -1.139$$

   recalling also that the Cobb-Douglas functional form implies that the equilibrium wage w can be expressed as a function of r through the firm's problem FOCs, i.e.

   $$w = w(r) = (1-\alpha)A\left(\frac{\alpha A}{r+\delta}\right)^{\frac{\alpha}{1-\alpha}} .$$

2. **Optimal labor policy.** To reduce the dimensionality of the household's problem, I exploit the intratemporal first-order condition for labor supply (see equation (B.7)), which allows me to express optimal labor as a function of consumption, $h_{opt}(c)$. Substituting this into the budget constraint yields a mapping from states and asset choices to implied consumption, $c^*(a,\rho,a')$, and thus to

implied labor, $h^*(a, \rho, a')$. Both objects are precomputed using a hybrid root-finding method (secant with fallback on bisection) and stored on a three-dimensional grid over (a, ρ, a'), and are used as interpolated inputs during value function iteration. We denote these interpolated functions as:

$$\widehat{c}(a, \rho, a') \quad \text{and} \quad \widehat{h}(a, \rho, a').$$

3. **Interpolated Value Function Iteration.** Given an initial guess $V(a, \rho)$, solve the household problem by value function iteration. At each grid point, the agent chooses next period's asset holdings a' to maximize:

$$\widehat{V}(a, \rho) = \max_{a'} \left\{ \frac{\widehat{c}(a, \rho, a')^{1-\sigma}}{1-\sigma} - B\frac{\widehat{h}(a, \rho, a')^{1+\varphi}}{1+\varphi} + \beta \sum_{\rho'} \pi(\rho' \mid \rho) \widehat{V}(a', \rho') \right\},$$

subject to $a' \geq -b$.

The functions $\widehat{c}(\cdot)$ and $\widehat{h}(\cdot)$ are interpolated over the precomputed (a, ρ, a') grid using linear interpolation. The continuation value $\widehat{V}(a', \rho')$ is also interpolated to evaluate expected future utility for off-grid values of a'.

At each step, this procedure yields the household's optimal policy functions for asset accumulation (a, ρ), labor supply $g_h(a, \rho)$, and consumption $g_c(a, \rho)$, which are jointly consistent with the value-maximizing decision.

The iteration continues until the value function converges:

$$\left\| \widehat{V} - V \right\| < \varepsilon^V, \quad \text{with} \quad \varepsilon^V = 10^{-6}.$$

4. **Stationary distribution.** Given the asset policy function $g_a(a, \rho)$ derived in step 2, I compute the stationary distribution $\Gamma(a, \rho)$ implied by fixed point iteration. For a given initial guess $\mu(a, \rho)$, the distribution is updated iteratively using:

$$\widehat{\Gamma}(a', \rho') = \sum_{i=1}^{N_a} \sum_{j=1}^{N_\rho} \mathbb{L}\left\{a' = a^*(a_i, \rho_j)\right\} \pi(\rho' \mid \rho_j)\, \Gamma(a_i, \rho_j),$$

where $\mathbb{L}$ denotes a linear interpolation operator that maps mass from off-grid choices $a^*(a_i, \rho_j)$ onto the asset grid. Specifically, if $a' \in [a_{i-1}, a_i]$, the interpolation weights are:

$$\mathbb{L}(a_i, a') = \mathbb{I}(a' \in [a_{i-1}, a_i]) \cdot \frac{a' - a_{i-1}}{a_i - a_{i-1}}.$$

This procedure is iterated until the distribution converges, i.e., until:

$$\left\| \widehat{\Gamma} - \Gamma \right\| < \varepsilon^\Gamma, \quad \text{with} \quad \varepsilon^\Gamma = 10^{-10}.$$

5. **General equilibrium.** Given the stationary distribution $\Gamma(a, \rho)$ and the policy functions $g_a(a, \rho)$, $g_h(a, \rho)$, and $g_c(a, \rho)$, I compute the implied aggregate quantities and check for market clearing. Specifically, the markets for capital, labor, and goods clear when:

$$\begin{aligned}
K &= \int_{\mathcal{A}\times\mathcal{P}} g_a(a, \rho)\, d\Gamma(a, \rho), \\
L &= \int_{\mathcal{A}\times\mathcal{P}} \rho \cdot g_h(a, \rho)\, d\Gamma(a, \rho), \\
C &= \int_{\mathcal{A}\times\mathcal{P}} g_c(a, \rho)\, d\Gamma(a, \rho), \\
AF(K, L) &= C + \delta K + G.
\end{aligned}$$

To ensure capital market clearing, I compute the difference between capital supply K and capital demand $\widehat{K}$, where the latter is defined by the firm's optimality condition:

$$\widehat{K} = L \left(\frac{r + \delta}{\alpha A} \right)^{-1/(1-\alpha)}.$$

I iterate on the equilibrium interest rate r until the capital market excess demand error $|K - \widehat{K}|$ falls below the tolerance threshold $\varepsilon^K = 10^{-6}$. The fixed point in r is computed using a Newton-Raphson method.

C.2 Simulating Alternative Tax Regimes

The simulation of tax-revenue-equivalent alternative equilibria begins with the construction of a uniform grid for perturbed values of τ_y on a $\pm 20\%$ range around its baseline value. I set $N_{\tau_y} = 48$ to optimise computational time, as the procedure is parallelised on a 16-core machine.

For each value of τ_y, I run a nested two-level Newton solver where
i) the *inner loop* solves for the equilibrium interest rate r such that capital market clears, using the same accuracy criterion $\varepsilon^K = 10^{-6}$ as in the benchmark; ii) the *outer loop* solves for the value of τ_c that satisfies:

$$G(\tau_y, \tau_c) = G_{\text{target}}.$$

This is done via Newton-Raphson iteration on the scalar function $G(\tau_c)$, with a convergence tolerance $\varepsilon^G = 10^{-6}$.

For each candidate (τ_y, τ_c) pair, I solve the full model following the same procedure described for the baseline model in Section C.1.

The same procedure with optimisation on λ_c instead of τ_c is then used to compute the LC-EQ equilibrium. Finally, the Crev-EQ equilibrium is computed using an analogous three-dimensional Newton-Rhapson method: I solve the equilibrium and find λ_c and τ_c in order to ensure i) capital market clearing, ii) revenue neutrality, and, additionally, iii) the same level of consumption tax revenue as in the baseline scenario.