Optimal Public Debt with Redistribution *

Matias Bayas-Erazo†

This version: February 15, 2026
First version: November 21, 2023

Abstract

Fiscal policy choices affect both the degree of progressivity of the tax system and the amount of public debt in circulation. What is the connection between the two? In this paper, I consider a benevolent optimizing government and explore how both progressivity and indebtedness depend on the government's preferences for redistribution. I study fully dynamic optimal policy in standard heterogeneous-agent incomplete-markets economies, computing both the optimal long-run policy and the transition towards it. Somewhat surprisingly, I find that differences in preferences for redistribution lead to a negative correlation between progressivity and indebtedness, as a planner that cares more for redistribution favors *lower* levels of public debt.

Keywords: public debt, redistribution, heterogeneous-agent models, optimal fiscal policy

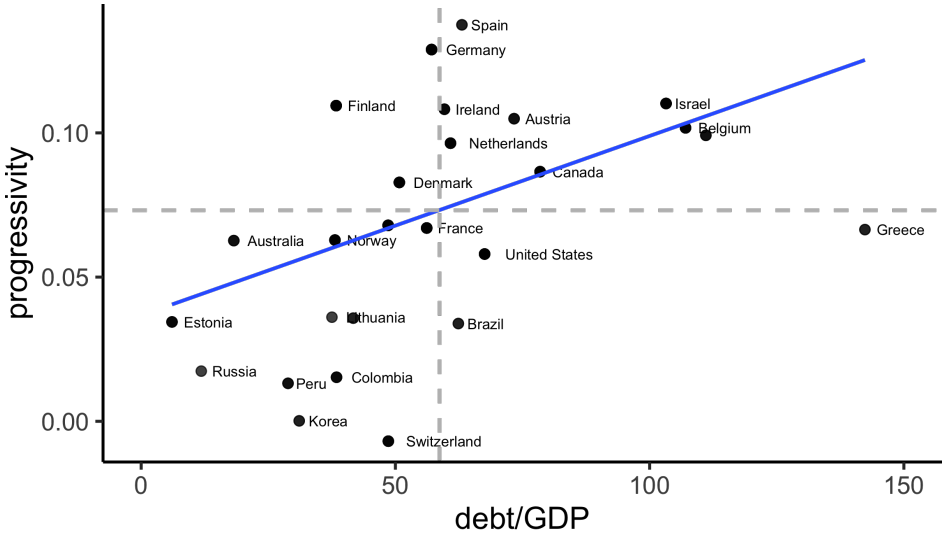
*This paper is a revised version of Chapter 1 of my PhD dissertation at Northwestern University. I am grateful to my advisors Guido Lorenzoni, George-Marios Angeletos, Alessandro Pavan, and Matthew Rognlie for their insights, guidance, and support throughout this project. I also thank Martin Beraja, Corina Boar, Axelle Ferriere, Joao Guerreiro, Fergal Hanks, Christian Hellwig, David Hémous, Diego Huerta, Florian Scheuer, Ludwig Straub, and Kwok Yan Chiu for comments that helped improve the paper. All remaining errors are my own.

†University of Zürich; matias.bayas-erazo@econ.uzh.ch

1 INTRODUCTION

There are ongoing debates on how public debt and progressive income taxes should be used. Several economists and policy makers have argued that, in an environment with low interest rates, governments can and should borrow more (see Blanchard, 2019). At the same time, progressive taxation is seen by many as a tool to address the recent increase in income and wealth inequality (Saez and Zucman, 2019; Heathcote et al., 2020). While there is extensive research on each of these fiscal instruments *in isolation*, there is less work exploring the connection between the two. Public debt and progressive taxation, however, are related both in theory and in practice.

Figure 1: Public debt and progressivity across countries, 1980-2015



Note: The data is taken from the IMF's *Historical Public Debt Database* and Qiu and Russo (2022).

Figure 1 plots the average progressivity of the tax system and the average debt-to-GDP ratio in the cross section. Countries with more progressive tax systems tend to have higher levels of public debt. But is this relationship natural from the perspective of optimal policy?

In theory, both public debt and progressive income taxes can play a role when it comes to insurance. When markets are incomplete, public debt provides liquidity and helps agents self-insure against idiosyncratic income risk (Woodford, 1990; Aiyagari and Mc-

Grattan, 1998). Meanwhile, the presence of income risk motivates the use of redistributive taxation (Mirrlees, 1974; Varian, 1980). A progressive income tax transfers resources from the lucky to the unlucky and acts as a form of social insurance. If countries choose a progressive tax system due to a high demand for insurance, then it makes sense for them to also favor higher levels of public debt. On the other hand, if progressive taxation lowers the aggregate demand for safe assets, it may increase the cost of borrowing, perhaps reducing the incentive to issue public debt. What, then, is the optimal *mix* of debt and progressivity, and how does it depend on social preferences for redistribution?

To address these questions, I compute the optimal mix of debt and redistributive taxation in standard heterogeneous-agent incomplete-markets models (Bewley, 1977; Huggett, 1993; Aiyagari, 1994). Households face uninsurable idiosyncratic income risk, supply labor, and self-insure by holding safe assets subject to borrowing constraints. A perfectly competitive firm hires labor and produces final output. The government controls the supply of safe assets and a progressive tax on labor income. In this context, I study fully dynamic optimal policy and solve for both the optimal long-run policy and the transition towards it.

In the long run, I find that planners with stronger preferences for redistribution implement more progressive tax systems, as one would expect. But, perhaps surprisingly, they also favor *lower* levels of public debt. In other words, differences in preferences for redistribution lead to a negative correlation between progressivity and indebtedness. This suggests something puzzling about the relationship between public debt and progressivity documented in Figure 1: if we order countries in terms of how much redistribution they like, those with more progressive tax systems should, in fact, favor lower levels of public debt.[1]

To explain what is behind this result, I derive a simple condition that characterizes the optimal long-run level of public debt in these economies. This condition shows that the optimal level of public debt depends on three key "sufficient statistics": i) the marginal

[1] Allowing for variation in risk across countries can explain the positive correlation of debt and progressivity in the cross-section. It is also likely that political economy considerations play an important role in explaining the relationship we observe in the data.

social value of public debt, ii) the sensitivity of interest rates with respect to changes in the level of public debt, and iii) the premium on public debt. This premium, the difference between the interest rate and the discount rate, arises due to incomplete markets. Together, these three objects determine the optimal level of public debt in the long run.

When I study the effects of progressive tax reforms, I show that increasing the progressivity of the tax system puts upward pressure on interest rates, thereby lowering the premium on public debt. This *interest rate channel of progressivity* follows from the fact that a progressive tax system reduces the need to self-insure against idiosyncratic income risk. This lowers the premium that private agents are willing to pay for holding safe assets and thus makes government borrowing more expensive. At the same time, a progressive tax system reduces the benefit of providing public liquidity. When the government is already offering insurance through the tax system, the value of self-insurance goes down, reducing the marginal social value of public debt. So both on the cost side and on the benefit side, more redistribution decreases the incentive to issue public debt. As a result, an inequality-averse planner, who naturally favors a redistributive tax system, finds it optimal to issue lower levels of public debt.

The key technical challenge is that the planner must keep track of a time-varying wealth distribution. To overcome this, I formulate the problem in sequence space, following the approach introduced by Auclert et al. (2023a). An additional complication is that the long-run optimum—what the literature calls the Ramsey steady state—may not exist under standard welfare criteria. Angeletos et al. (2022) study the Ramsey problem in a stylized economy that abstracts from heterogeneity and show that long-run satiation can be optimal. Auclert et al. (2023a) find a similar pattern in the full model: under utilitarian welfare and separable preferences, the planner issues debt until the interest rate equals the discount rate. At that point, however, asset holdings diverge, which is inconsistent with a well-defined steady state in this class of models.

Given this, I depart from utilitarian welfare criteria and consider planners that use generalized stochastic welfare weights (Davila and Schaab, 2022). These weights depend on endogenous outcomes, are used to aggregate individual utilities, and decouple society's concerns for fairness from agents' lifetime welfare. To the extent that infinitely-lived

agents represent altruistically linked generations, this approach can be interpreted as one that distinguishes between the welfare of each generation (Phelan, 2006; Farhi and Werning, 2007). In particular, it allows the planner to care about the inequality of future generations directly. I refer to such planners as generational planners. Although non-standard, this class of welfare functions nests the utilitarian criterion and allows me to speak about inequality-aversion while keeping the problem tractable.

I show that an interior Ramsey steady state exists for generational planners that are averse to wealth inequality. A planner that values redistribution from the asset-rich to the asset-poor no longer finds satiation optimal. This result highlights the role of inequality considerations for the optimal level of public debt, and helps overcome the difficulties that previous work has faced in this class of models.

Beyond the long-run optimum, I also study the optimal *transition* towards the Ramsey steady state. Computing the optimal transition is more involved: the planner must choose policy instruments at each date along a potentially long horizon. I address this challenge by anchoring policy paths to the Ramsey steady state and parameterizing deviations using a low-dimensional time basis. I then exploit the sequence-space structure to compute policy gradients analytically, combining insights from Auclert et al. (2023a) and Dyrda and Pedroni (2022). This makes the problem tractable on accessible hardware.

Along the optimal transition, the interest rate channel continues to play a central role. The planner starts by flattening the tax schedule: progressivity falls and average taxes decline. Because the tax system now provides less insurance, the private demand for safe assets rises, pushing interest rates down. The resulting increase in the premium on public debt creates fiscal space that helps finance the reform. This allows the planner to sustain lower taxes and higher consumption in the short run while gradually moving the debt-to-GDP ratio towards the long-run target.

To explore the robustness of these results, I consider a number of extensions. First, I consider steady-state welfare analysis à la Aiyagari and McGrattan (1998), computing the mix of debt and progressivity that maximizes steady-state welfare while ignoring transitional dynamics. Because this problem is computationally more tractable, I can use more standard social welfare functions. This verifies that the properties of the optimal mix are

not driven by the details behind the SWF: whenever social preferences are represented by a SWF that puts a relatively higher weight on the wellbeing of the asset poor, the optimal mix involves low debt and high progressivity. Quantitatively, I find that the steady-state problem overestimates the costs of debt issuance and overestimates the benefits of progressive tax systems. This leads the planner that ignores transitions to choose lower levels of debt and higher levels of progressivity.

Second, I analyze a version of the model with multiple safe assets and taxes on savings. This specification features a more general production technology and relaxes the assumption that the only supply of bonds outside the household sector comes from the government. The qualitative properties of the optimal mix remain unchanged but there are important quantitative differences. Across all kinds of planners, the optimal mix becomes less progressive and features lower levels of public debt. Notably, the negative correlation between progressivity and indebtedness becomes stronger. This is because a progressive tax system lowers equilibrium wages and increases interest rates, which scales down the stochastic component of consumer income (Davila et al., 2012). As consumers' effective exposure to risk decreases, the value of providing public liquidity falls more than in the baseline with fixed wages.

Third, I examine a simple open economy version in which some demand for safe assets comes from abroad. This setup allows me to test the importance of the interest rate channel in driving the main results. By varying the responsiveness of domestic interest rates to increases in net foreign debt, I can control the degree to which the interest rate channel is active (Schmitt-Grohé and Uribe, 2003). At one extreme, in a standard small open economy where the interest rate is completely exogenous, the interest rate channel is entirely muted. Between this extreme and the baseline closed economy, there are intermediate cases where the interest rate responds partially to changes in domestic conditions. This helps isolate the role of the interest rate channel: when the interest rate is less responsive to domestic conditions, the negative correlation between optimal progressivity and public debt becomes significantly weaker.

Fourth, I introduce more flexible labor income tax schedules. In the main body of the paper, I restrict attention to tax systems that exhibit a constant rate of progressivity

(CRP), as in Bénabou (2002) and Heathcote et al. (2017). This rules out the possibility of lumpsum transfers. Given their empirical relevance, I also consider labor income tax systems with negative intercepts. This variation does not affect the observation that planners that care about redistribution favor lower levels of debt but it does have implications for the optimal shape of average and marginal taxes. In particular, the optimal tax schedule now features increasing average tax rates and decreasing marginal tax rates, as in Ferriere et al., 2022.

Finally, I carry out an inverse optimum exercise to back out implied preferences for redistribution (Bourguignon and Spadaro, 2012; Heathcote and Tsujiyama, 2021). This gives a back-of-the-envelope calculation for the type of planners that would rationalize the observed mixes of debt and progressivity as solutions to the optimal policy problem analyzed in the first part of the paper. The results suggest that implied preferences for redistribution in selected advanced economies are inconsistent with standard Utilitarian and Rawlsian criteria: the covariance between implied welfare weights and assets/labor income is *positive*. In simpler terms, in order to explain the observed mix of debt and progressivity, the SWF must put a relatively higher weight on the welfare of the rich. Interestingly, the ranking implied by the model puts the US and Denmark at opposite ends of the spectrum in terms of their preference for redistribution.

Related literature

This paper contributes to the literature on optimal fiscal policy with heterogeneous agents that begins with Aiyagari (1995). Assuming the existence of the Ramsey steady state, that paper characterizes some properties of the long run optimum, including the modified golden rule and positive capital income taxes. Due to the difficulties involved in tracking the wealth distribution, most studies deviate from the original Ramsey problem and follow the seminal work of Aiyagari and McGrattan (1998), who compute the level of debt that maximizes welfare in steady state. Like most of the literature, they restrict attention to linear taxes and thus ignore the interaction with progressive taxation that is the focus of this paper. One notable exception is Flodén (2001), who studies the optimal steady state mix of debt and transfers. He also finds that debt may provide insurance more effectively

whenever equity considerations are excluded from the welfare objective. However, he does not relate the properties of the optimal mix to the interest rate channel, and more importantly, continues to ignore transitions. Angeletos et al. (2022) solve the full Ramsey problem in a stylized incomplete-markets economy that bypasses the computational challenges of the problem considered in this paper. They clarify how the approach taken by Aiyagari and McGrattan (1998) and Flodén (2001) ends up overestimating the costs of the services provided by public debt and thus underestimates its long-run quantity.

Relatedly, a series of papers in the quantitative Ramsey tradition point out that accounting for the transition path can lead to a very different optimal tax schedules in environments with heterogeneous agents. Bakış et al. (2015) focus on once-and-for-all changes in the tax system and find that accounting for transitions leads to a more progressive optimal tax system. Krueger and Ludwig (2016) look at the interaction between progressive taxation and education subsidies and also find that the optimal progressivity of the tax system depends on whether or not transitional dynamics are taken into account. Boar and Midrigan (2022) also study once-and-for-all reforms, evaluating welfare consequences along the transition. They find small welfare gains from enriching the set of instruments available to the planner. All of these papers abstract from public debt and are thus unable to relate the optimal level of debt to redistribution.

More recently, the literature has returned to the original Ramsey problem in Aiyagari (1995), developing different approaches to address the computational challenges. Acikgoz et al. (2023) and LeGrand and Ragot (2023) use a Lagrangian approach, inspired by Marcet and Marimon (2019). Dyrda and Pedroni (2022) directly search for the optimal sequence of policies after parameterizing them in the time domain. Auclert et al. (2023a) introduce a sequence-space approach for computing the Ramsey steady state and find that the standard heterogeneous agent model with utilitarian welfare criteria has no interior steady state.[2] LeGrand and Ragot (2023) opt for an inverse optimal taxation approach, estimating a social welfare function that makes the current tax system consistent with the planner's optimality conditions. Relative to these papers, I extend the sequence-space ap-

[2]To get around the non-existence of the RSS with separable preferences, Acikgoz et al. (2023) rely on GHH preferences, whereas Dyrda and Pedroni (2022) use a KPR utility function.

proach of Auclert et al. (2023a) to solve for both the Ramsey steady state and the optimal transition. To compute the transition, I parameterize policy paths in the time domain, as in Dyrda and Pedroni (2022), but use the sequence-space representation to efficiently evaluate equilibrium responses. By anchoring to the RSS and using gradient-based optimization, the approach runs on accessible hardware.

Given the focus on redistribution, this paper also speaks to a literature that is closer to the Mirrleesian approach to optimal taxation. The authors working with static models tend to emphasize the equity-efficiency tradeoff but Varian (1980) points out that redistributive taxation can be viewed as a form of social insurance. He shows that a government can effectively insure individuals against income risk, a theme that is explored further by the literature working with dynamic Mirrleesian models (see Golosov et al. (2006) and Farhi and Werning (2013), among others). Unlike this paper, the Mirrleesian approach tends to work in partial equilibrium and disregards the role of public debt. An important exception is Werning (2007), who studies nonlinear fiscal policy in a model with complete markets. He finds that the relationship between taxes and debt is indeterminate (i.e. Ricardian equivalence holds).[3] Chang and Park (2021) look at nonlinear tax policy in the standard incomplete markets model but assume away the role of public debt by forcing the government to balance the budget every period. Similarly, Ferriere et al. (2022) focus on the optimal design of transfers and progressivity and find that the optimal log-linear tax wih a transfer generates welfare gains almost as large as the Mirrleesian allocation.

2 Model

Consider a standard incomplete markets economy with a continuum of households who face uninsurable idiosyncratic income risk (Huggett, 1993; Aiyagari, 1994). The only asset is a one-period risk-free government bond that pays an interest rate r and can be freely traded up to some borrowing limit $\phi > 0$. Individual productivity θ evolves according to

[3]Relatedly, Bhandari et al. (2017b) show that Ricardian equivalence holds in the standard incomplete markets model when the government can also control the tightness of borrowing constraints. Bhandari et al. (2017a) consider fiscal policy and debt management jointly but restrict attention to proportional taxes.

some Markov process and determines the wage per unit of labor supplied by the agents. There is no aggregate uncertainty.

Given a sequence of interest rates $\{r_t\}$ and nonlinear labor income tax schedules $\{T_t(\cdot)\}$, individuals face an income fluctuation problem with endogenous labor supply. The value function of an agent entering the period with assets a and productivity θ in period t is

$$V_t(a,\theta) = \max_{\ell,c,a'} u(c) - v(\ell) + \beta \mathbb{E}_{\theta'|\theta}\left[V_{t+1}(a',\theta')\right] \quad \text{s.t} \quad \begin{cases} c + a' = (1+r_t)a + \theta\ell - T_t(\theta\ell) \\ a' \geq -\phi, \end{cases}$$

where c, ℓ, and a' are consumption, labor supply, and next period's asset holdings and $\beta \in (0,1)$ is the agent's discount factor. Pre-tax labor income is given by $y = \theta\ell$. For future reference, let $\boldsymbol{c}_t(x)$, $\boldsymbol{a}_t(x)$, and $\boldsymbol{y}_t(x)$ denote the policy functions for consumption, asset holdings, and labor income for an agent in state $x = (a,\theta)$. Also, denote by $D_t(\theta, A)$ the measure of households with productivity θ that have assets in set A at the beginning of period t.

For the baseline, I restrict attention to tax schedules that exhibit a constant rate of progressivity (CRP), as in Bénabou (2002) and Heathcote et al. (2017):

$$T_t(y) = y - \tau_t\, y^{1-p_t}, \tag{1}$$

for some $\{p_t, \tau_t\}$ with $p_t < 1$ and $\tau_t \in \mathbb{R}$, for all t. The parameter p_t indexes the progressivity of the tax schedule in period t, whereas τ_t governs the average level of taxes. Taxes are linear if $p_t = 0$, progressive if $p_t > 0$, and regressive if $p_t < 0$. With this functional form, after-tax income $\boldsymbol{z}_t(x) = \tau_t\, \boldsymbol{y}_t(x)^{1-p_t}$. Notice that this rules out the possibility of lumpsum transfers, an important feature of tax and transfer systems in practice. I address this limitation in Section 5, where I consider alternative labor income tax schedules. There, I show that the results are unchanged with a simple tax structure that still captures a form of progressive taxation: linear taxes with a lump-sum intercept. The same is true if I consider a CRP$_+$ tax system, a three parameter version of (1) that allows for a lump-sum intercept.

To close the model, I assume that the government supplies the risk-free bond subject

to a standard budget constraint

$$G + r_{t-1}B_{t-1} = B_t - B_{t-1} + \int T_t(\boldsymbol{y}_t(x))dD_t(x), \quad \text{for all } t. \tag{2}$$

In words, the government's exogenous spending needs $G \geq 0$ and the interest payments on the debt must be financed by aggregate tax revenues and net debt issuance in period t. On the production side, I assume that the effective labor ($\theta\ell$) of different productivity types is perfectly substitutable in production. This means that each unit of effective labor produces one unit of goods and that the real wage is equal to one. Goods market clearing requires

$$\int \boldsymbol{y}_t(x)dD_t(x) = \int \boldsymbol{c}_t(x)dD_t(x) + G, \quad \text{for all } t. \tag{3}$$

Finally, the asset market clearing condition is

$$\int \boldsymbol{a}_t(x)dD_t(x) = B_t, \quad \text{for all } t. \tag{4}$$

Note that the only supply of safe assets outside the household sector comes from the government. In particular, there is no capital. I start with this assumption in order to isolate the role of public debt but I relax it in Section 5, when I consider a more general production technology and allow firms to issue claims to physical capital.

Definition 1 An equilibrium is a sequence of interest rates $\{r_t\}$ and tax schedules $\{T_t\}$, a sequence of policy functions $\{\boldsymbol{c}_t(x), \boldsymbol{a}_t(x), \boldsymbol{y}_t(x)\}$ and distributions $\{D_t\}$ such that, given an initial distribution D_0:

(i) $\{\boldsymbol{c}_t(x), \boldsymbol{a}_t(x), \boldsymbol{y}_t(x)\}$ are optimal given $\{r_t, T_t\}$,

(ii) D_t is consistent with the policy functions and the Markov process for productivity,

(iii) the government's budget constraint (2) is satisfied,

(iv) the asset market clearing condition (4) and the goods market condition (3) hold. □

2.1 CALIBRATION

I calibrate the model to the US economy, following McKay et al. (2016) whenever possible. A period is one quarter and there is no borrowing. The discount factor β is chosen such that the ratio of aggregate liquid assets to GDP in the model is consistent with US data, given a real interest rate of 2%.[4] Following standard practice in the literature, I assume that θ follows an AR(1) process in logs with persistence ρ and an innovation variance σ_ϵ^2. These parameters are chosen to match Floden and Lindé (2001)'s estimates for the persistence of the US wage process and the standard deviation of earnings growth in Guvenen et al. (2014). I discretize the AR(1) process for productivity using the Rouwenhorst method on eight idiosyncratic states. The elasticity of intertemporal substitution is equal to one and the Frisch elasticity of labor supply is $1/2$, consistent with Chetty et al. (2011). For taxes, the value for progressivity p is taken from Heathcote et al. (2017). Government spending is 8.8% of annual GDP, the average ratio of government expenditures to output in the US over the period 1970 to 2013.[5] Given these choices, the level of taxes τ is pinned down by the government's budget constraint. Table 1 summarizes the parameters that come out from this procedure.

Table 1: Parameter values

Parameter	*Description*	*Value*	*Parameter*	*Description*	*Value*
β	discounting	0.988	G/Y	spending-to-GDP	0.088
ρ	persistence of AR (1)	0.966	B/Y	debt-to-GDP	1.4
			p	progressivity of	0.181
σ_ϵ	variance of AR(1)	0.033		taxes	
EIS	curvature in u	1	τ	level of taxes	0.641
$Frisch$	curvature in v	$1/2$	ϕ	borrowing limit	0

[4] McKay et al. (2016) calculate liquid assets from aggregate household balance sheets reported in the flow of funds and take the average ratio over the period 1970 to 2013. They arrive to a value of $A/Y = 1.4$.

[5] See Table 1.1.5. Gross Domestic Product in the National Income and Product Accounts.

3 INTEREST RATE CHANNEL OF PROGRESSIVITY

This section studies how progressivity affects the real interest rate in the model laid out in Section 2. The analysis helps isolate the interest rate channel of progressivity, which plays a central role in shaping the optimal mix of debt and progressivity. Because progressive tax reforms raise equilibrium interest rates, they can have unintended distributional consequences that benefit the asset-rich.

Consider a small permanent change in the progressivity of the tax schedule dp around the baseline economy while holding the initial level of debt fixed. The goal is to understand how this reform affects the equilibrium interest rate, the level of taxes, and individual payoffs. For simplicity, I focus on stationary outcomes. Lemma 1 characterizes the response of the real interest rate, the level of taxes, and individual value functions to this change in the tax system:

Lemma 1 *Given a small permanent change in progressivity dp,*

1. *The general equilibrium responses $d\boldsymbol{r}$ and $d\boldsymbol{\tau}$ are given by the solution to the following system:*

$$\begin{bmatrix} \frac{\partial \mathcal{A}^{ss}}{\partial r} & \frac{\partial \mathcal{A}^{ss}}{\partial \tau} \\ \frac{\partial \mathcal{T}^{ss}}{\partial r} - B & \frac{\partial \mathcal{T}^{ss}}{\partial \tau} \end{bmatrix} \begin{bmatrix} d\boldsymbol{r} \\ d\boldsymbol{\tau} \end{bmatrix} = \begin{bmatrix} -\frac{\partial \mathcal{A}^{ss}}{\partial p} \\ -\frac{\partial \mathcal{T}^{ss}}{\partial p} \end{bmatrix}.$$

2. *The response of individual outcomes dV solves:*

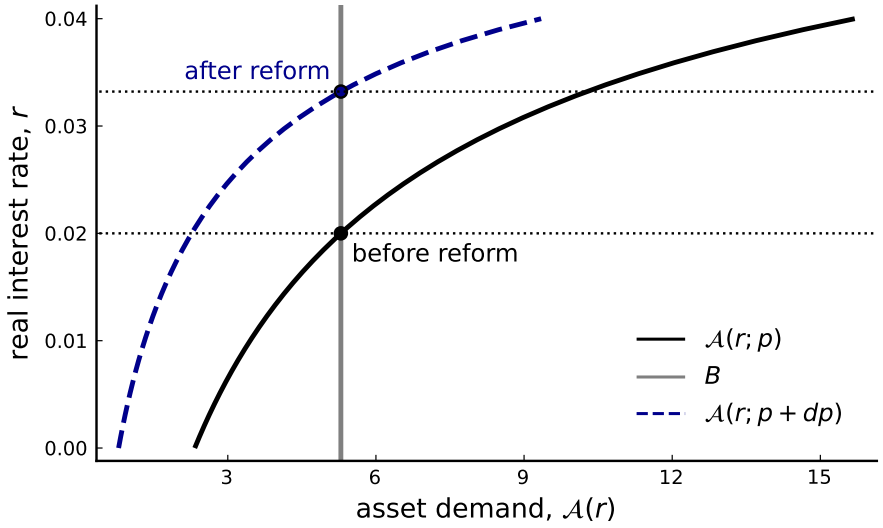
$$dV(x) = u'(\boldsymbol{c}(x)) \left(\boldsymbol{y}(x)^{1-p} d\boldsymbol{\tau} + a\, d\boldsymbol{r} - \boldsymbol{z}(x) \log \boldsymbol{y}(x) \right) + \beta \mathbb{E}_{\theta'|\theta} \left[dV(\boldsymbol{a}(x), \theta') \right]. \quad (5)$$

PROOF The first part follows from the implicit function theorem. The second part follows from the envelope theorem, taking into account general equilibrium effects. ■

The interest rate response to the reform depends on the partial equilibrium responses of aggregate asset demand $\mathcal{A}^{ss}(r, p, \tau) = \int a_i di$ and aggregate tax revenues $\mathcal{T}^{ss}(r, p, \tau) = \int T(y_i) di$ to changes in interest rates and taxes around the initial steady state. Because these objects are functions of the distribution of agents over idiosyncratic states, it is not possible to characterize them analytically. However, there is a simple intuition that helps understand the sign of the interest rate response.

For a given r, an increase in the progressivity of the tax system reduces aggregate asset demand because it lowers the volatility of after-tax income. There is more insurance happening via the tax system and thus less need to hold assets to insure against idiosyncratic income risk.[6] When the supply of debt is held fixed, the interest rate must rise in order to restore equilibrium in the asset market. This means that $d\boldsymbol{r} > 0$ in general. Figure 2 illustrates this mechanism by plotting the equilibrium in the asset market before and after this progressive tax reform. The point at which the solid black curve intersects with the grey line gives the equilibrium interest rate prior to the reform. The dashed blue line is the asset demand curve after the reform. It shifts to the left because the need for self-insurance is lower. Since the debt is held fixed, the intersection between the blue dashed line and the grey line gives the interest rate after the reform. The difference between the two dotted lines is what I call the interest rate channel of progressivity.

Figure 2: Equilibrium in the asset market before and after progressive tax reform



One of the main contributions of this paper is to isolate the implications of this mechanism for the optimal mix of debt and progressivity. A close relative of the interest rate channel appears in recent work by Mian et al. (2022), who show that inequality increases fiscal space in a two-agent model. Relatedly, Amol and Luttmer (2022) find that transfers

[6]There is an additional effect that works in the same direction: the increase in progressivity raises the level of taxes (i.e. $d\boldsymbol{\tau} < 0$). This lowers disposable income and thus shifts the demand for safe assets further to the left.

reduce the demand for safe assets and lowers the upper bound on deficits in a model with perpetual youth. In contemporaneous work, Kaplan et al. (2023) find that progressive tax systems reduce the maximum sustainable deficit of the government in a model that is similar to the one I use. Similarly, Chang and Park (2021) study how progressive tax reforms affect equilibrium interest rates in heterogeneous-agent incomplete-markets models and find results consistent with the mechanism I describe here. Taken together, their findings point to a potential conflict between progressive taxation and debt issuance. However, they do not study the optimal mix of these two instruments, which is the main focus of this paper.

The quantitative relevance of the interest rate channel depends on the sensitivity of aggregate safe asset demand to changes in the real interest rate. Intuitively, if aggregate safe asset demand is sufficiently elastic, then a small change in interest rates suffices to clear the market. While direct empirical estimates of how tax progressivity affects interest rates are not available, there is substantial evidence on how public debt affects interest rates.[7] Mian et al. (2022) survey this literature and find semi-elasticity estimates between 1.2% and 2.2%, whereas my calibrated model yields 0.6%. This suggests the interest rate channel would likely be stronger in models that better match these empirical elasticities.

In any case, the mechanism is quantitatively strong in standard calibrations of the heterogeneous-agent incomplete-markets model. Because of this, the indirect general equilibrium effects of progressive tax reforms, operating through the interest rate response, can outweigh the direct partial equilibrium effect. As a result, progressive tax reforms may have unintended effects.

UNINTENDED EFFECTS OF PROGRESSIVE TAX REFORMS

The second part of Lemma 1 shows that changes in the progressivity of the tax system affect individual life-time utilities through partial equilibrium (direct) and general equilibrium (indirect) channels. To see this more clearly, iterate (5) forward to write the response

[7] De Ferra et al. (2021) find that countries that are more unequal save less. This is not necessarily inconsistent with the mechanism I describe here because they rely on cross-sectional data and do not control for the degree of financial development across countries.

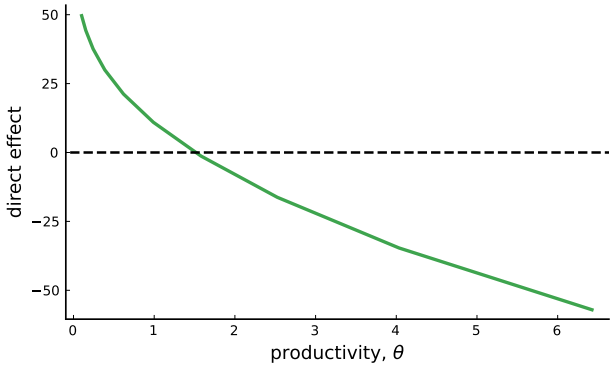
of individual outcomes as

$$dV(x) = \sum_{s=0}^{\infty} \beta^s \mathbb{E}\left[u'(c_s) \left(\underbrace{y_s^{1-p} d\boldsymbol{\tau} + a_s \, d\boldsymbol{r}}_{\text{indirect effect in } s} - \underbrace{z_s \log y_s}_{\text{direct effect in } s} \right) \mid x_0 = x \right].$$

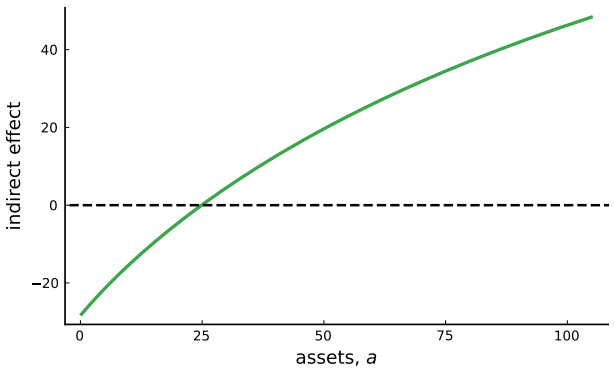
The second term inside the parenthesis is the expected direct effect s periods after the reform for an agent with initial state equal to x. This comes about because the change in p affects the slope of the tax schedule. By doing so, it lowers taxes paid at the bottom of the income distribution while increasing the taxes paid at the top. Figure 3a shows that this partial equilibrium effect tends to favor the agents that are income poor (i.e. those with low productivity) at the time of the reform.

The first term inside the parenthesis is the expected indirect effect s periods after the reform, which shows up due to the general equilibrium nature of the exercise– changes in progressivity trigger movements in the real interest rate and the level of taxes in order to ensure asset markets clear and the government's constraint holds. As Figure 3b shows, these general equilibrium effects tend to favor the agents that are asset rich at the time of the reform. Notice that it makes it possible for those at the top of the wealth distribution to *benefit* from progressive tax reforms. This is because the effect mainly operates through the interest rate response to the reform, which as I argued above, is positive.

Figure 3: Direct and indirect effects across the state space



(a) Direct effect along the productivity dimension

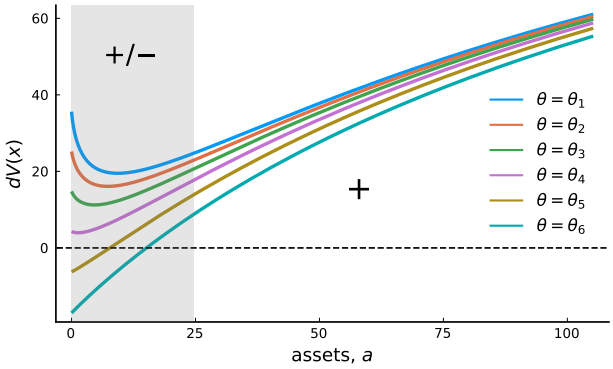


(b) Indirect effect along the asset dimension

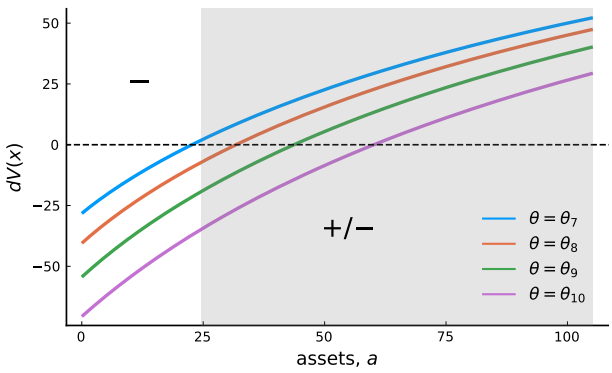
By combining the logic behind the direct and indirect effects, the state space can be divided into (a) regions where agents unambiguously win or lose from the reform, and

(b) regions where the outcome is ambiguous. First, those who are both income-poor and asset-rich will always favor progressive tax reforms: they benefit from lower taxes at the bottom of the income distribution and the higher interest rates induced by the reform. Second, those who are both income-rich and asset-poor will always be worse off: they pay higher taxes and do not have enough assets to benefit from the increase in interest rates. For the agents that are poor or rich along both dimensions, the ultimate effect is ambiguous and hinges on whether the direct or the indirect effect dominates, which is a quantitative question.

Figure 4: Individual responses across the state space



(a) Total effect at the bottom of the θ distribution



(b) Total effect at the top of the θ distribution

These different regions can be seen in Figure 4, where I plot the total change in the value function along the asset dimension, for different levels of productivity. The panel on the left shows the change in lifetime utilities for the agents at the bottom of the productivity distribution, whereas the panel on the right shows it for those at the top. The figures are constructed using the thresholds for assets and productivity where the indirect and direct effect switch signs. Productivity levels that are below the value of θ at which the direct effect becomes negative are included in the left panel, whereas the productivity levels that are above this threshold are in the right panel. The frontier between the grey region, where the total effect is ambiguous, and the white region, where there is no ambiguity, corresponds to the level of assets where the indirect effect switches sign.

The asset-rich who are also income-poor at the time of the reform always benefit from

progressive income taxes: the overall change in the value function is positive. Moreover, in this calibration, the interest rate response is so dominant that it allows the asset-rich to benefit from the reform, regardless of whether they are at the bottom or at the top of the productivity distribution. On the other hand, most of the agents with low wealth end up losing from the reform, except for those at the bottom of the productivity distribution. For the asset-poor, the indirect effect is mostly driven by the effect of the reform on the level of taxes, which is negative. Among these agents, only the low-productivity types are better off because they alone benefit from the direct effects of the reform.

To sum up, progressive tax reforms lead to higher interest rates and this indirect effect is quantitatively relevant. The interest rate channel can result in unintended effects of progressive tax reforms, where the asset-rich end up benefiting from the reform. As discussed above, indirect empirical evidence suggests that this channel would likely be stronger in more realistic models. Below, I show that this channel plays a key role when it comes to understanding the optimal mix of debt and progressivity and helps explain why inequality-averse planners favor relatively low levels of public debt.

4 OPTIMAL POLICY

This section introduces the Ramsey problem and computes the optimal plan. The fact that the Ramsey planner takes into account transitional dynamics complicates the analysis. I build on recent work by Auclert et al. (2023a) to arrive at a simple characterization for the steady state of the Ramsey plan. After verifying that an interior Ramsey steady state exists, I turn to the optimal long-run mix of debt and progressivity and the transition towards it.

4.1 RAMSEY PROBLEM

The Ramsey planner chooses *sequences* of interest rates and CRP tax codes to maximize the present discounted value of aggregate utility subject to implementability conditions. As mentioned earlier, this problem takes into account the transition path and allows both instruments to vary over time. I begin by characterizing the limiting steady state of the

Ramsey plan, which can be computed independently of the transition.[8] This reduces the dimensionality of the problem and provides a target for the transition analysis that follows.

Formally, the dynamic Ramsey problem is

$$\max_{\{r_t, B_t, p_t, \tau_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) \text{ s.t } \begin{cases} \mathcal{A}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = B_t, \\ G + (1 + r_{t-1}) B_{t-1} = B_t + \mathcal{T}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right), \end{cases} \tag{6}$$

Here, $\mathcal{U}_t(\cdot)$ is a sequence-space function that maps sequences of interest rates and CRP tax codes into "aggregate utility" at time t. Similarly, $\mathcal{A}_t(\cdot)$ and $\mathcal{T}_t(\cdot)$ map sequences of interest rates and taxes into aggregate asset holdings and aggregate tax revenues in period t. These functions aggregate optimal individual decisions using the distribution of agents in the economy. This ensures that the fiscal policy chosen by the planner can be implemented in equilibrium, provided that the government's budget constraint holds and the asset market clears. I assume that the Ramsey planner aggregates individual utilities as follows:

$$\mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = \int_i \omega_t(\theta_{it}, a_{it}) U(c_{it}, l_{it})\, di, \tag{7}$$

with the weights $\omega_t(\theta, a) \propto \exp\left(-\alpha_\theta \theta - \alpha_a a\right)$.[9] The parameters $\alpha_a \in \mathbb{R}$ and $\alpha_\theta \in \mathbb{R}$ govern the planner's aversion to inequality along the asset and productivity dimension, respectively. This SWF nests the utilitarian criterion when $\alpha_a = \alpha_\theta = 0$. When $\alpha_a > 0$, the planner dislikes inequality in asset holdings in the sense that it discounts the instantaneous utility of the asset rich, relative to a utilitarian benchmark. Similarly, when $\alpha_\theta > 0$, the planner's aversion to income inequality leads to a lower weight on the instantaneous utility of the income rich.

This social welfare function departs from standard welfare criteria for two reasons. First, it takes as input individual instantaneous utilities, as opposed to individual lifetime utilities. Second, it depends on weights that are endogenous. Formally, the planner uses

[8] See Acikgoz et al. (2023) for details.

[9] The weights depend on t because they are normalized by $\bar{\omega}_t = \int_i \exp(-\alpha_\theta \theta_{it} - \alpha_a a_{it})\, di$, the average weight in each period. This normalization ensures that the SWF does not penalize increases in asset holdings *per se*.

generalized dynamic stochastic weights to conduct welfare assessments, in the sense of Davila and Schaab (2022). These weights allow one to formalize new welfare criteria that society may find appealing and decouple society's concerns for fairness from individual utilities, as in Saez and Stantcheva (2016). Here, allowing the weights to depend on assets and productivity captures the view that society considers it fair to redistribute across these dimensions.

My preferred interpretation is that the social planner distinguishes between the welfare of each generation. To the extent that a model with infinitely lived agents represents altruistically linked generations of finitely lived agents, this SWF captures the preferences of a planner that cares about inequality of future generations *directly* (Phelan, 2006; Farhi and Werning, 2007).[10] When $\alpha_a > 0$, the planner lowers the weight on the utility of generations that have relatively higher assets because it dislikes the fact that future generations become more unequal than the initial generation. When $\omega(\theta, a) = 1$ for all a and all θ, the SWF weights the utilities of all generations equally. Following this interpretation, I refer to this type of planners as generational planners.

Departing from a utilitarian welfare criterion helps resolve the complications that arise when searching for the steady state of the Ramsey plan. The literature continues to debate the existence of an interior Ramsey steady state in the standard heterogeneous agent model with separable preferences. Chien and Wen (2022) claim that an interior Ramsey steady state does not exist when the elasticity of intertemporal substitution is less than one. In an environment with deterministic income fluctuations, LeGrand and Ragot (2023) prove that the steady state exists for separable CRRA utility functions provided that the planner is not utilitarian. In a more standard calibration of the model, the results in Auclert et al. (2023a) suggest that a version of the Friedman rule may be optimal when the Ramsey planner uses a utilitarian welfare criterion. In other words, a utilitarian planner finds it optimal to satiate the demand for public debt, issuing debt to the point where the interest rate is equal to the discount rate. However, this is not consistent with an inte-

[10] I do not want to take a stand on the philosophical debate of whether the welfare of future generations should enter the planner's objective. I am simply arguing that *if* society cares about the inequality of future generations, then the SWF would resemble (7).

rior steady state in this class of models: aggregate asset demand asymptotes to infinity as $\beta(1+r) \to 1$.

Below, I show that this is not the case when the planner cares to redistribute from the asset-rich to the asset-poor. In this case, an interior Ramsey steady state exists.

4.2 Sufficient statistics for the optimal long-run choice of debt

To solve the Ramsey problem, I rely on a set of first order conditions that I obtain by perturbing the optimal plan. Here, I derive a simple "sufficient-statistic" representation for the optimal long-run choice of debt. In Appendix A, I derive the full set of first-order conditions for the Ramsey problem and elaborate on the computational approach to compute the steady state.

Fix an arbitrary sequence of CRP tax codes $\{p_t, \tau_t\}$. For any $u = 0, 1, \ldots,$ the following must be true:

$$\sum_{t=0}^{\infty}\sum_{s=0}^{\infty}\beta^{t-u}\frac{\partial \mathcal{U}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \sum_{t=0}^{\infty}\sum_{s=0}^{\infty}\beta^{t-u}\lambda_t\frac{\partial \mathcal{T}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \lambda_u - \beta\lambda_{u+1}(1+\boldsymbol{r}_u) - \sum_{t=0}^{\infty}\beta^{t-u}\lambda_t\frac{\partial \boldsymbol{r}_t}{\partial B_u}B_{t-1} = 0. \quad (8)$$

Here, λ_t denotes the Lagrange multiplier on the government's budget constraint in period t and $\boldsymbol{r}_t(\cdot)$ is a sequence-space function that maps sequences of CRP tax codes $\{\tau_t, p_t\}$ and public debt $\{B_t\}$ into the interest rate that clears the asset-market at time t. If (8) does not hold, then a small perturbation dB_u would increase social welfare, contradicting the optimality of the Ramsey plan.

To understand the intuition behind each term let me first ignore the summations. The first term says that increasing the amount of debt at time u has a direct effect on welfare via interest rates. The strength of this effect is mediated by the sensitivity of interest rates to changes in the level of public debt. The second term arises due to income effects on labour supply. The third term captures the effects on the government's budget: λ_u is the shadow value of public funds in period u, and issuing debt relaxes the budget at time u. However, this debt has to be repaid next period at the given interest rate, which explains the term $-\beta\lambda_{u+1}(1+\boldsymbol{r}_u)$. Finally, because interest rates are endogenous in this model, the change in debt issuance is also going to change interest rate payments on the existing debt. Now, the double summations are there due to the forward/backward nature of

the system. Individuals are forward looking and changes in the distribution propagate slowly. Therefore, welfare, tax revenues, and interest rates depend on the entire sequence of fiscal policy chosen by the government.

Despite this, Proposition 1 shows that this condition simplifies around the steady state:

Proposition 1 *Suppose that* $\lim_{u\to\infty} \lambda_u \to \lambda^{RSS} > 0$. *Then, the optimal long-run level of debt* B^{RSS}, *if it exists, solves*

$$\left[\frac{\mathcal{E}_r^{\mathcal{U}}}{\lambda^{RSS}} + \mathcal{E}_r^{\mathcal{T}}\right] \mathcal{E}_B^{\boldsymbol{r}} + \{1 - \beta(1+\boldsymbol{r})\} - \mathcal{E}_B^{\boldsymbol{r}}\, B^{RSS} = 0, \tag{9}$$

where $\mathcal{E}_X^F \equiv \lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial F_t}{\partial X_u}$.

PROOF See Appendix A.4. ∎

From this result, it is clear that the optimal long-run level of public debt can be written in terms of three key "sufficient statistics". First, the marginal social value of public debt. Ignoring income effects on labor supply this is simply the discounted long-run response of aggregate utility to changes in the interest rate, in dollar terms: $\frac{\mathcal{E}_r^{\mathcal{U}}}{\lambda^{RSS}}$. Second, the so-called premium on public debt, which is the difference between the interest rate and the discount rate: $1 - \beta(1+\boldsymbol{r})$. This arises due to incomplete markets and captures the rents that the government extracts from the private sector due to its ability to issue risk-free debt. Finally, the sensitivity of interest rates to changes in public debt: $\mathcal{E}_B^{\boldsymbol{r}}$. This object captures the extent to which the government can affect the *path* of interest rates by issuing more or less debt. In models with complete markets, $\mathcal{E}_B^{\boldsymbol{r}} = 0$ and the condition reduces to $1 = \beta(1+r)$. Since this is already guaranteed by the consumer's Euler equation in steady state, the optimal level of public debt is indeterminate. In models with incomplete markets, $\mathcal{E}_B^{\boldsymbol{r}} \neq 0$ and the optimal level of public debt is determined by the relative strength of the three terms.

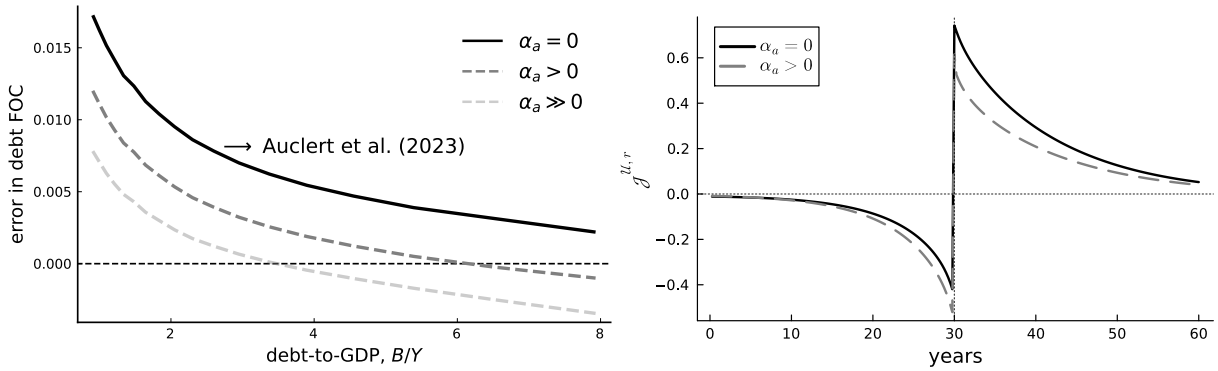
Equation (9) is quite general and it holds beyond the specific structure of the model considered in this paper. It applies to a broad class of incomplete-markets economies, including those in the spirit of Holmström and Tirole (1998) where debt acts as collateral, and those with overlapping generations where debt acts as a vehicle for life-cycle savings (Diamond, 1965). The mapping from fiscal policy choices to sufficient statistics changes but the relevant objects are always the ones uncovered in Proposition 1.

4.3 EXISTENCE OF INTERIOR STEADY STATE WITH INEQUALITY-AVERSE PLANNERS

I now verify that an interior Ramsey steady state exists whenever the Ramsey planner uses weights that decrease with asset holdings (i.e. whenever $\alpha_a > 0$). To explain why this is the case, I compare the asymptotic response of aggregate utility to changes in the interest rate with and without aversion to inequality.

Figure 5a plots the error in the first order condition for the optimal long-run choice of debt as the debt-to-GDP ratio varies, holding fixed the progressivity of the tax system. The solid line confirms the results in Auclert et al. (2023a): there is no interior RSS in the standard heterogeneous-agent model with separable CRRA preferences and utilitarian SWF. The dashed lines, on the other hand, show that a unique, interior RSS exists whenever the planner puts a relatively higher weight on the utility of generations with low asset holdings ($\alpha_a > 0$). In Appendix A, I also verify that existence and uniqueness go through when the planner also optimizes over the progressivity of the tax system.

Figure 5: Existence of interior RSS with inequality-averse generational planners



(a) The search for B^{RSS} and inequality aversion (b) $\mathcal{E}^{\mathcal{U}_r}$ with and without inequality aversion

To understand what drives this result, I look at the Jacobians of aggregate utility to changes in the interest rate: $\mathcal{J}^{\mathcal{U},r}_{t,s} = \frac{\partial \mathcal{U}_t}{\partial r_s}$.[11] From Proposition 1, the β-discounted sum of a far-out column of $\mathcal{J}^{\mathcal{U},r}$, $\mathcal{E}^{\mathcal{U}}_r = \lim_{s\to\infty} \sum_{t=0}^{\infty} \beta^{t-s} \frac{\partial \mathcal{U}_t}{\partial r_s}$, governs the marginal social value of issuing public debt around the RSS. This is the only component in (9) that depends on the planner's aversion to inequality. Therefore, understanding how this varies with the

[11] See Auclert et al. (2021) for more on these sequence-space Jacobians.

planner's aversion to inequality is key.

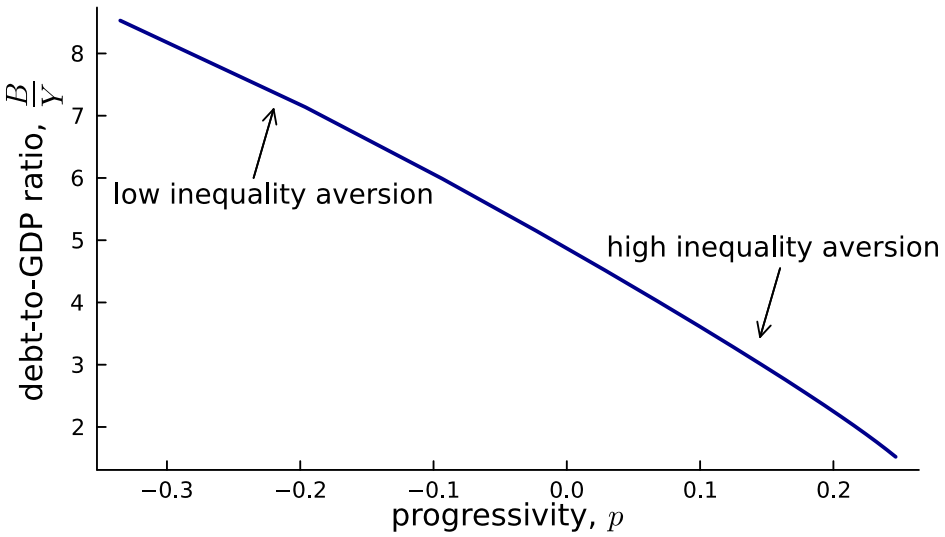
Figure 5b plots the long-run discounted response of aggregate utility for a utilitarian planner (solid line) and an inequality-averse planner (dashed line). A generational planner that underweights the utility of the asset-rich assigns a lower marginal benefit to increases in public debt compared to a utilitarian planner. This difference arises for two reasons. First, the effects due to the anticipatory response, which lowers consumption and thus lowers aggregate utility prior to the shock, are stronger at the bottom of the wealth distribution. This is because the marginal utility of consumption is higher for those who are asset-poor. Second, the benefits once the shock hits are discounted by the generational planner because a higher r, once realized, mostly benefits those who are asset rich. Together, these two effects imply that the marginal benefit of increasing public debt is lower for an inequality-averse planner than for a utilitarian planner. This explains why introducing aversion to inequality allows me to find an interior Ramsey steady state.

4.4 Optimal long-run mix of debt and progressivity

Now that we know that an interior steady state exists with inequality-averse planners, I turn to the optimal long-run mix of debt and progressivity. I solve the problem for different degrees of inequality aversion and plot the optimal long-run mix in the space of debt-to-GDP and progressivity in Figure 6. What varies along the frontier is α_a, the planner's aversion to wealth inequality.

In the bottom right corner of the figure we have the planners that are highly averse to wealth inequality. They naturally prefer a progressive tax system as these redistribute towards the agents they care about the most. Given the positive correlation between asset holdings and labor income, a progressive labor income tax helps the asset-poor. Now, because of the interest rate channel discussed in Section 3, this means they face relatively higher interest rates in the long run: redistributive taxation reduces the need to self insure against idiosyncratic income risk. In order to induce agents to hold a given supply of safe assets, the government must eventually offer agents a higher interest rate. Hence, it is optimal to issue lower levels of public debt in the long run. To put it in terms of the sufficient statistics isolated in Proposition 1, the premium on public debt is lower when

Figure 6: Optimal long-run mix of debt and progressivity



the planner favors a redistributive tax system.

In addition, the value of issuing public debt falls when the need for self-insurance is lower. When the government provides more insurance through the tax system, $\mathcal{E}_r^{\mathcal{U}}$, which captures the response of aggregate utility to changes in interest rates, drops. This is one of the key determinants of the marginal social value of public debt, as shown in Proposition 1.[12] The combination of these two effects explains the negative correlation between progressivity and indebtedness.

If differences in preferences for redistribution explain most of the variation in fiscal policy across countries, this result suggests something puzzling about the relationship between debt and progressivity documented in Figure 1. It also goes against the view that left-leaning governments, who typically favor redistributive policies, are more likely to be fiscally irresponsible. Of course, the real world is more complex, and the mix of debt and progressivity might not be driven solely by normative considerations. Political economy considerations likely play a key role and policy makers may not be aware of how the government's stance on progressive taxation interacts with its capacity to incur debt. Another possibility is that most of the variation in fiscal policy across countries

[12]The overall effect on the marginal social value of public debt might be ambiguous, though, because income effects cause $\mathcal{E}_r^{\mathcal{T}}$ to increase with higher progressivity.

is not driven by differences in preferences for redistribution. In fact, it is possible to rationalize the relationship documented in Figure 1 from the perspective of optimal policy if differences in idiosyncratic income risk drive most of the variation (see Appendix A for a figure).

Next, I show that the interest rate channel is also active along the optimal transition, and that the properties of the optimal mix identified above are robust. In particular, they hold regardless of whether transitional dynamics are taken into account and do not depend on the details behind the SWF used in this section. They also hold in more general versions of the model that allow for (a) multiple safe assets, (b) more flexible labor income tax schedules, and (c) taxes on savings.

4.5 Optimal transition to the Ramsey steady state

With the optimal long-run policy in hand, I now turn to the optimal transition: given an initial distribution, what is the optimal time path of fiscal instruments that takes the economy to the Ramsey steady state?

The Ramsey problem (6) optimizes over sequences of policy variables. The computational challenge is that the planner chooses policy instruments at each date along a potentially long horizon. As described in the introduction, I address this by anchoring policy paths to the Ramsey steady state and parameterizing deviations using a low-dimensional Chebyshev time basis. I then combine the sequence-space formulation of Auclert et al. (2023a) with the direct parameterization method of Dyrda and Pedroni (2022). The implementability constraints are handled via an augmented Lagrangian:

$$\mathcal{L}_{\rho_{AL}}(x, \lambda) = -W(x) + \lambda^{\top} \boldsymbol{g}(x) + \frac{\rho_{AL}}{2} \|\boldsymbol{g}(x)\|^2, \tag{10}$$

where $x = \{r_t, p_t, \tau_t\}_{t=0}^{T}$ stacks the policy variables over the transition horizon, $W(x) = \sum_{t=0}^{T} \beta^t \mathcal{U}_t$ is the discounted sum of aggregate utility, $\boldsymbol{g}(x)$ collects the goods market clearing residuals at each date[13], λ is the vector of Lagrange multipliers, and $\rho_{AL} > 0$ is a

[13] By Walras' Law, the original problem (6) can be reformulated with goods market clearing as the only constraint, after dropping the government budget constraint and the associated choice variable B_t. See Appendix A.3 for details.

penalty parameter. I solve this problem using policy gradient techniques adapted from machine learning, exploiting the sequence-space structure to compute the gradients analytically. The gradients of (10) involve sequence space Jacobians around *arbitrary paths*.[14] To reduce the computational burden, I parallelize the construction of the Jacobian matrix and use the recursive structure of the household problem to speed up the computation of each column. The details are in Appendix D.

Figure 7 shows the optimal transition paths for an inequality-averse planner ($\alpha_a = 0.025$) starting from the baseline calibration in Section 2. The planner starts by flattening the tax schedule: progressivity falls on impact and the average tax burden declines. This reduction in distortions triggers an immediate boom in output and consumption.

The interest rate channel, which drives the long-run results in Section 3, continues to operate along the transition. Because the tax system now provides less insurance, the private demand for safe assets rises, pushing interest rates down. The resulting increase in the premium on public debt creates fiscal space: the government can service its debt more cheaply and sustain lower taxes while gradually moving the debt-to-GDP ratio towards the long-run target. Robustness checks in Appendix D confirm that these patterns hold regardless of the specific parameters or basis functions used to approximate the path.

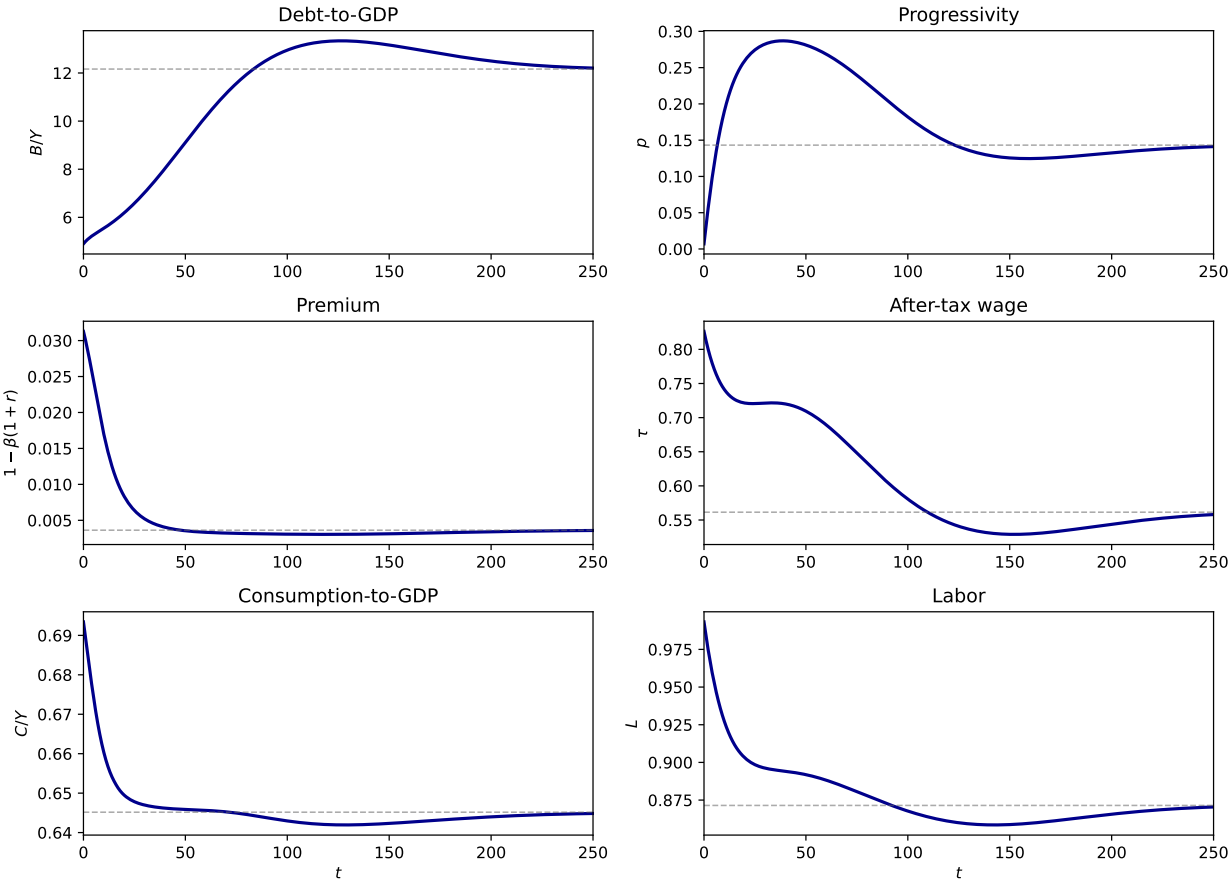
5 Extensions

5.1 Optimal steady state analysis

Due to the difficulties involved in tracking the wealth distribution, the methods for solving the dynamic optimal policy problem in this paper have been developed fairly recently (Acikgoz et al., 2023; Auclert et al., 2023a; Dyrda and Pedroni, 2022; Ragot and Le Grand, 2023). To make progress, Aiyagari and McGrattan (1998) focused on stationary outcomes, computing the level of debt that maximizes welfare in steady state, ignoring transition dynamics. For the purposes of this paper, this remains a useful exercise because it is computationally tractable and permits the study of richer models. Moreover, it can be thought of as the limit of the dynamic problem when the social discount factor approaches one.

[14] This prevents the use of the efficient algorithm in Auclert et al. (2021).

Figure 7: Optimal transition paths



Note: The planner has inequality aversion $\alpha_a = 0.025$. The transition horizon is $T_{\text{free}} = 425$ quarters using a Chebyshev time basis of degree 5 per instrument and an anchoring function $e_t = (1 - t/T_{\text{free}})^{\kappa}$ with $\kappa = 10$ that anchors the path to the Ramsey steady state. The transition appends a constant tail of length $T_{\text{tail}} = 375$ at terminal policies to approximate the infinite-horizon continuation. Dashed lines indicate Ramsey steady-state values.

Given a social welfare function $\mathcal{W}$, the optimal steady state problem (OSS) is to choose a time-invariant CRP tax-code $\{\tau, p\}$ and the steady state level of public debt B in order to maximize

$$\max_{\{r,B,\tau,p\}} \mathcal{W}(r,\tau,p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(r,\tau,p) = B, \\ G + rB = \mathcal{T}(r,\tau,p) \end{cases} . \tag{11}$$

Here, $\mathcal{A}(r,\tau,p)$ and $\mathcal{T}(r,\tau,p)$ denote aggregate asset holdings and aggregate tax revenues in steady state. The two constraints ensure that the government budget is balanced and that the asset market clears.

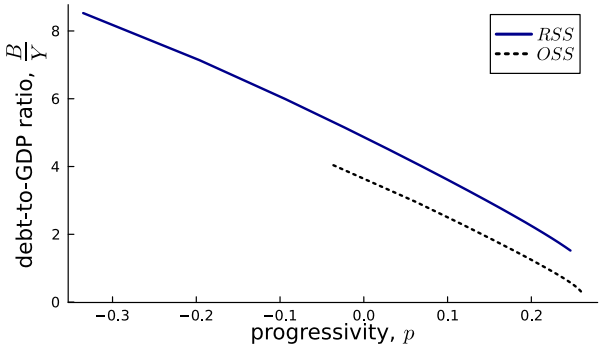
I first use an analogous version of the social welfare function to the one used in the fully dynamic problem. That is, I assume the planner uses weights $\omega(a,\theta) \propto e^{-\alpha_a a - \alpha_\theta \theta}$ to aggregate instantaneous utilities. As in Section 4, I focus on how the optimal mix changes as the planner's aversion to wealth inequality α_a varies. The optimality conditions and the details behind the computations are in Appendix B.

Figure 8a traces out the optimal mix of debt and progressivity for the OSS problem (dotted black line) and compares it to the frontier of the fully dynamic problem (solid blue line). The qualitative properties of the optimal mix are unchanged: planners that have stronger preferences for redistribution implement more progressive tax systems and relatively lower levels of debt. Since the OSS problem can be thought of as the limit of the dynamic problem when the planner's discount factor approaches one, this shows that the main results of the paper are robust to the choice of social discount factor. The reasons for the negative correlation between progressivity and indebtedness are the same as in the dynamic problem, as a similar sufficient statistics formula applies (see Appendix B.4).

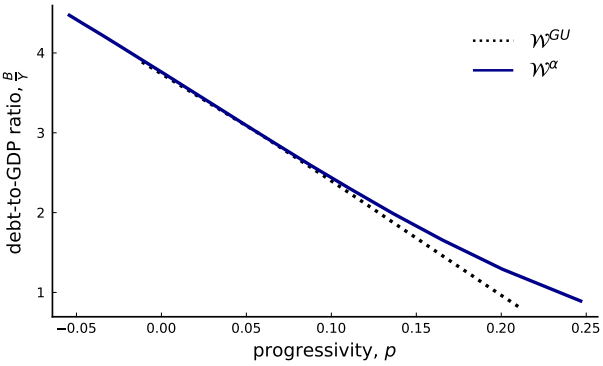
Transitional dynamics only matter for the quantitative properties of the optimal mix. The optimal set in the OSS shifts down and to the right. In other words, the optimal long-run mix in the OSS problem features relatively lower levels of public debt and higher progressivity. To see this more clearly, the bottom two panels of Figure 8 plot how each instrument varies with the planner's aversion to wealth inequality. For given α_a, the planner that ignores transitions favors relatively lower levels of public debt (bottom left panel). This is consistent with the results in Angeletos et al. (2022), who clarify how this concept of long run optimality overestimates the costs of debt issuance and thus under-

estimates the optimal long-run level of public debt. The intuition for this result is that the planner that ignores transitions does not take into account the benefit of issuing debt along the transition, which helps reduce reliance on distortionary taxation. The bottom right panel shows that it also overestimates optimal long-run progressivity. This is also related to the fact that the planner does not take into account the distortionary effects of taxation along the transition.

Figure 8: The optimal long-run mix ignoring transitions and across SWFs



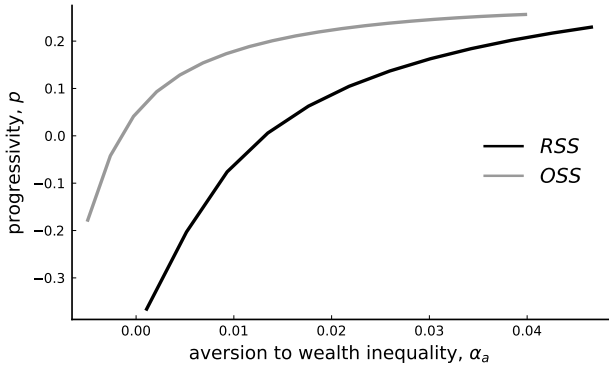
(a) The optimal long-run mix ignoring transitions



(b) The optimal long-run mix across SWFs



(c) Optimal B/Y with and without transitions



(d) Optimal p with and without transitions

To explore the robustness of the results to my choice of SWF, I also solve (11) with alternative SWFs. The purpose of this is to convince the reader that the main results are not driven by my choice of SWF. First, I consider a generalized utilitarian criterion with weights that may depend on the productivity and asset holdings. Formally,

$$\mathcal{W}^{GU}(r,\tau,p) = \int \omega(\theta_i,a_i)V(\theta_i,a_i)\,di, \tag{12}$$

with the weights parameterized as before. Relative to a generational planner, this kind of social planner only cares about inequality of the initial generation: the weight on the utility of future generations does not depend on the asset holdings or productivity of that generation. In this sense, the social preferences behind equation (7) may reflect a stronger notion of aversion to inequality than the ones behind equation (12). Second, following Bénabou (2002) and Boar and Midrigan (2022), I also consider a social welfare function that separates society's aversion to inequality from household's aversion to intertemporal fluctuations. Letting $\bar{c}_i$ denote the consumption certainty-equivalent of agent i[15], social welfare for a "Bénabou planner" is given by

$$\mathcal{W}^{\alpha}(r,\tau,p) = \left(\int \bar{c}_i^{1-\frac{1}{\alpha}}\, di\right)^{\frac{1}{1-\frac{1}{\alpha}}}. \tag{13}$$

The parameter $\alpha \in (0,\infty)$ governs society's aversion to inequality. When $\alpha \to \infty$, the objective captures pure economic efficiency and puts no value on equity of consumption per se – redistribution has value only to the extent that it relaxes borrowing constraints or reduces idiosyncratic risk. As Bénabou (2002) puts it, efficiency concerns are thus separated from pure equity concerns. If α equals the agents' elasticity of intertemporal substitution, one recovers the standard utilitarian criterion. As $\alpha \to 0$, the objective reduces to that of a Rawlsian planner who only cares about the welfare of the poorest agents.

The results for these SWFs are displayed in Figure 8b. This verifies that the negative correlation between progressivity and indebtedness is not driven by the details behind the SWF. Whenever social preferences are such that the planner has a taste for redistribution, the optimal long-run mix features lower debt and higher progressivity.

5.2 Multiple safe assets and taxes on savings

In the baseline model, the only supply of bonds outside the household sector comes from the government. This section considers an extension with a more general production technology $F(K,L)$ and allows firms to issue claims to capital. This introduces an alternative asset that households can use to smooth their consumption. The details behind

[15] Formally, $\bar{c}_i$ is the constant level of consumption a household would need to receive, without working, in order to achieve the equilibrium lifetime utility $V(\theta_i, a_i)$.

the calibration of this model and a complete definition of the equilibrium are relegated to Appendix C.1.

Given the absence of aggregate risk, capital and government bonds are perfect substitutes. This means that in equilibrium the interest rate on both assets must be the same. Firm optimality implies that before-tax factor prices w_t and r_t must satisfy

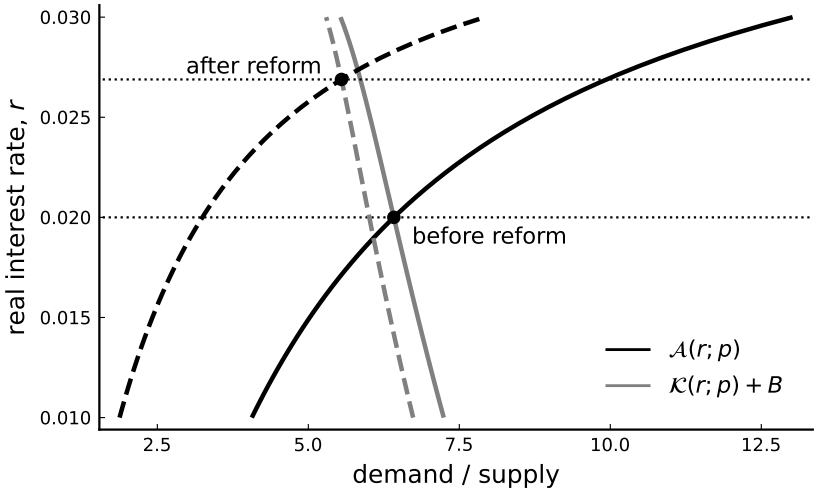
$$w_t = F_L(K_{t-1}, L_t) \quad \text{and} \quad r_t = F_K(K_{t-1}, L_t) - \delta.$$

The key difference with respect to the baseline model is that the asset market clearing condition becomes

$$B_t + K_t = \int \boldsymbol{a}_t(x) dD_t(x).$$

This implies that the supply of safe assets is not perfectly inelastic. The interest-rate elasticity of the supply of safe assets will of course depend on the curvature of the production function. Here, I work with a standard Cobb-Douglas production function $F(K, L) = K^{\alpha} L^{1-\alpha}$, which generates a fairly inelastic supply of capital. This has implications for the interest rate channel discussed in Section 3. Specifically, the response of the equilibrium interest rate to changes in the progressivity of the tax system becomes less pronounced. There is an additional margin of adjustment and thus the change in interest rates that is required to clear the asset market is smaller (see Figure 9).

Figure 9: The interest rate channel in the model with capital



The dynamic optimal policy problem for this economy is similar to the one in Section

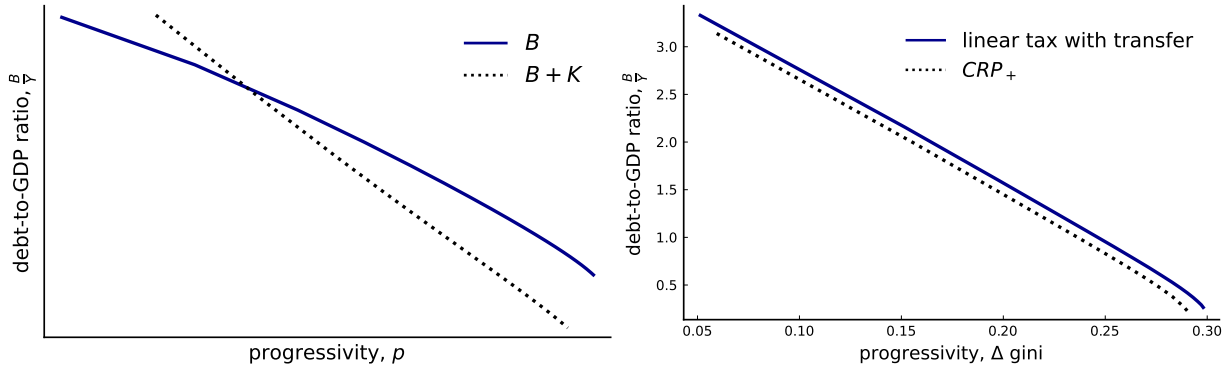
4 except that I allow the planner to tax the return on savings at rate τ_k.[16] In Appendix C.2, I derive the optimality conditions for this problem and show that the planner chooses τ_k in order to implement the modified golden rule ($F_K = \delta + \beta^{-1} - 1$). This result is already in Aiyagari (1995) but I verify that it goes through when the planner has a taste for redistribution. As Acikgoz et al. (2023) discuss, the fact that distributional concerns do not interfere with the efficient level of investment is reminiscent of the production efficiency result in Diamond and Mirrlees (1971). There is no need to implement an inefficiently high level of capital to help agents self-insure against risk because of the availability of public debt.

Once the capital-labor ratio has been determined, one then essentially solves a reparameterized version of the baseline model to compute the optimal long-run mix. This helps understand why the qualitative properties of the optimal long-run mix, displayed in Figure 10a, are similar to the baseline model. Interestingly, the negative correlation between progressivity and indebtedness becomes *stronger*: the dotted line is steeper than the blue line. This is despite the fact that the interest rate channel is less powerful in the model with capital. The reason for this is due to an effect that operates through wages, which is absent in the baseline model with exogenous wages. An increase in the progressivity of the tax system increases r but lowers w. This scales down the stochastic component of consumer income. Thus, the value of providing public liquidity falls more than in the baseline with fixed wages. Accounting for this effect explains why the optimal level of debt falls more than in the baseline model. Relatedly, lowering B increases w, which is good for redistribution. This means that the planner that cares about redistribution should lower debt by more relative to the baseline model.

The quantitative differences are due to the fact that the additional supply of capital reduces reliance on government debt to fulfill the demand for safe assets, thereby lowering the optimal level of debt compared to the baseline model. In addition, debt crowds out capital and this additional cost of debt issuance calls for lower levels of debt in the long run.

[16] A linear tax on savings makes no difference in the baseline model.

Figure 10: Optimal long-run mix with capital and transfers



(a) Optimal long-run mix with capital

(b) Optimal long-run mix with transfers

5.3 Open economy considerations

I now extend the baseline model to allow for foreign demand for safe assets. This provides a natural way to isolate the role played by the interest rate channel. Following Schmitt-Grohé and Uribe (2003), I introduce debt-elastic interest rates. This means the domestic interest rate is now given by

$$r = r^{\text{ROW}} + \psi_2(e^{D_f} - 1), \tag{14}$$

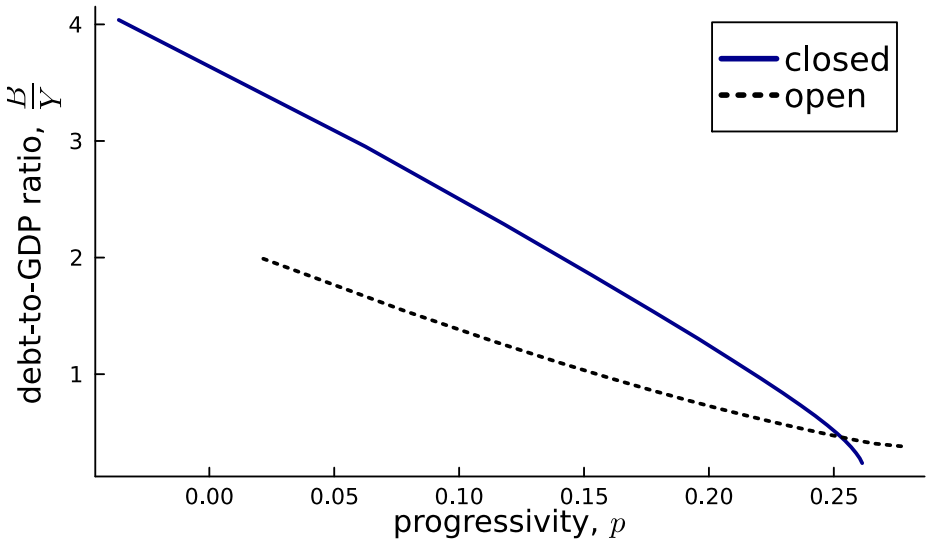
where r^{ROW} is the world interest rate, D_f is net foreign debt, and $\psi_2 \geq 0$ is a parameter that govern the sensitivity and the level of domestic interest rates, respectively. The parameter ψ_2 controls the strength of the interest rate channel: when $\psi_2 \to 0$, the economy approaches the small open economy limit with $r = r^{\text{ROW}}$, effectively shutting down the interest rate channel; as ψ_2 increases, the model behaves more like the closed economy baseline.

The asset market clearing condition becomes $A + D_f = B$, where $A = \int a_i di$ is domestic private asset demand and B is government debt. The government's budget constraint remains unchanged.

Figure 11 shows the optimal mix of debt and progressivity across different values of inequality aversion for both the closed economy (solid line) and a partially open economy (dotted line). Two patterns emerge. First, the negative relationship between progressivity

and debt becomes less pronounced. Second, the frontier shifts to the right, indicating that the optimal progressivity of the tax system is higher in the open economy.

Figure 11: Optimal long-run mix: closed vs. open economy



Note: The figure shows the optimal mix of debt and progressivity for different values of inequality aversion. The solid line represents the closed economy benchmark, while the dotted line represents an open economy with debt-elastic interest rates ($\psi_2 = 0.00005, r^{\text{ROW}} = 0$).

Both of these results support the idea that the interest rate channel is a key driver of the properties of the optimal mix. With foreign demand for safe assets, the link between domestic progressivity and equilibrium interest rates weakens. This allows the planner to implement a more progressive tax system without incurring the higher interest rates that would otherwise arise in the closed economy case. This is consistent with the idea that the interest rate channel is a key driver of the negative correlation between progressivity and indebtedness identified in this paper.

5.4 Alternative labor income tax schedules

I also study the properties of the optimal mix with alternative labor income tax schedules. First, I consider a simpler tax system that still captures a form of progressive taxation: linear taxes with a lump-sum intercept. Specifically, I assume that taxes are now given

by $T(y) = (1-\tau)y - T_0$, with $T_0 \geq 0$. For simplicity, I focus on the problem where the planner ignores transitions.

The results for this tax system are displayed in Figure 10b (solid blue-line). The insight that planners that care more about redistribution should favor lower levels of debt goes through. The reason for this is related to the interest rate channel that I have emphasized throughout the paper. An inequality-averse planner naturally favors the use of lump-sum transfers because this lowers average tax rates at the bottom, favoring the poor. But transfers, by acting as a form of social insurance against idiosyncratic risk, reduce the need for self-insurance. This lowers the aggregate demand for safe assets and puts upward pressure on interest rates, thereby making government borrowing more expensive.

Figure 10b also presents the results for a three-parameter version of (1):

$$T(y) = y - \tau y^{1-p} - T_0,$$

where, once again, $T_0 \geq 0$. This CRP$_+$ tax system introduces a negative intercept into the tax system analyzed so far and gives the planner more flexibility. In particular, it allows the planner to disentangle average tax rates from marginal tax rates. A number of recent papers have argued that this improves the empirical fit to the overall tax and transfer system in the United States.[17]

As above, inequality-averse planners rely on lump-sum transfers to separate average tax rates from marginal tax rates and redistribute towards the poor. They continue to favor lower levels of debt, but unlike the baseline, the optimal tax sytem now features *decreasing* marginal tax rates ($p < 0$). The intuition for this result is that this allows the planner to achieve redistribution while preserving efficiency (Ferriere et al., 2022). This is related to the fact that p is no longer a sufficient statistic for the progressivity of the tax system. A more global measure of progressivity, such as the change in the Gini coefficient as one moves from before-tax income distributions to after-tax income distributions, is now needed. Once this is taken into account, the negative correlation between progressivity and indebtedness remains (notice the x-axis in Figure 10b).

[17] The two-parameter tax function tends to overestimate taxes paid at the top and underestimate transfers at the bottom. See Boar and Midrigan (2022) and Ferriere et al. (2022) for more details.

To sum up, incorporating lump-sum transfers into the analysis does not change the observation that planners that care about redistribution favor lower levels of debt provided that one uses an appropriate measure of progressivity. It does, however, have implications for the optimal shape of average and marginal taxes. In particular, the optimal mix features increasing average tax rates but decreasing marginal tax rates.

6 INVERTING THE OPTIMUM

This section uses the normative model to rank social preferences for redistribution in the US and a collection of advanced economies.[18] I start by asking how a Utilitarian planner evaluates the observed mix of debt and progressivity in the US. Then, I show that parsimonious deviations from a utilitarian SWF struggle to explain the empirical mix. Finally, I ask what kind of social preferences for redistribution are consistent with the data for the US and other advanced economies. This exercise is inspired by the *inverse optimal taxation problem* (Bourguignon and Spadaro, 2012; Heathcote and Tsujiyama, 2021). The takeway is that implied social preferences for redistribution appear inconsistent with both Utilitarian and Rawlsian criteria. Moreover, the ranking implied by the model puts the US and Denmark at opposite ends of the spectrum in terms of their preference for redistribution.

6.1 UTILITARIAN PLANNERS & US FISCAL POLICY

For this subsection, I focus on the optimal policy problem that ignores transitional dynamics. This is because the optimal plan does not converge to an interior steady state when the planner is Utilitarian.

From the perspective of a Utilitarian planner, the US issues too little debt and is too progressive. The average value of public debt in the US between 1995 and 2007 was around 61.5% (Dyrda and Pedroni, 2022), well below 312%, the optimal debt-to-GDP ratio in the optimal steady state problem with a utilitarian SWF analyzed in Section 5.[19] The

[18] I do this given the features of the model and data availability.

[19] The number used in the calibration for the baseline model in Section 2 is 140 %, which is somewhat higher because in that model the notion of debt, being the only safe asset in the economy, is broader. This

progressivity of the US tax system, estimated by Heathcote et al. (2017) using PSID data from 2000 to 2006, is 0.181 and is also far from what is favored by a Utilitarian planner who ignores transitional dynamics (0.048).

When public debt is not part of the analysis, studies typically find that the optimal tax system under a utilitarian criterion should be more progressive than what we observe in the data (Heathcote and Tsujiyama, 2021). The difference here stems from the general equilibrium effects of progressive tax reforms. In particular, the interest rate channel highlighted in Section 3 pushes the planner to favor less progressivity relative to the status quo even under the utilitarian criterion. This highlights the importance of determining the optimal progressivity of the tax system together with the level of public debt.

These large differences are not driven by the absence of capital in the baseline model. In the model with multiple safe assets, the overall conclusion is the same. First, if the planner is not allowed to tax capital, the optimal level of debt is closer to what we see in the data, but the gap between optimal progressivity and estimated progressivity *increases*. In the model with capital taxes analyzed in Section 5.2, the situation reverts back to the case without capital. Allowing for more flexible forms of labor income taxation (i.e. lumpsum transfers as in Section 5.4) does not affect the conclusion.

6.2 Backing out inequality-aversion

It is then reasonable to ask if parsimonious deviations from utilitarian SWF criterion can bring the values implied by the normative theory closer to what we see in the data. If we allow for a single parameter α that captures the planner's aversion to inequality, this turns out not to be the case. Indeed, if we move towards a generational planner that is averse to wealth inequality, the optimal debt-to-GDP ratio can be made arbitrarily close to the one in the US. However, the required degree of inequality aversion implies an optimal progressivity that is *higher* than the one we observe in the data.

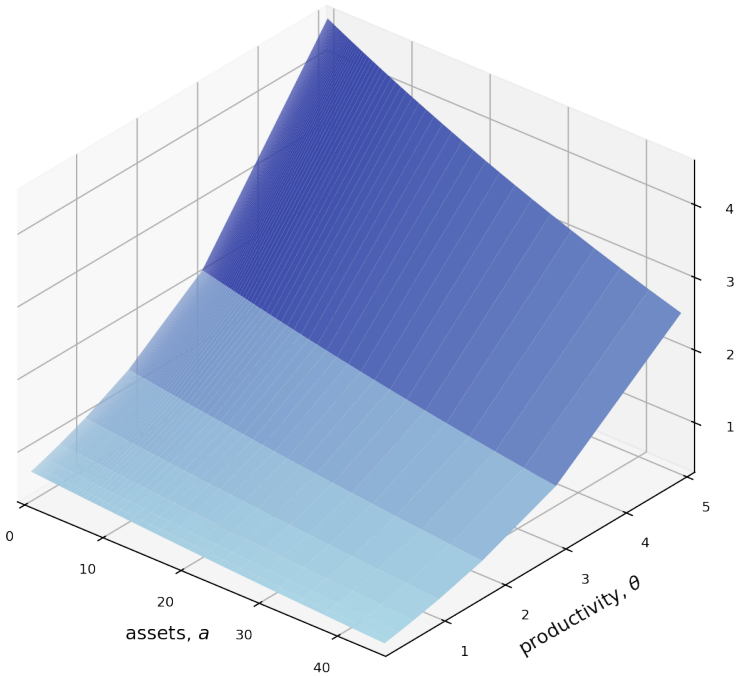
To determine what kind of social preferences can rationalize the observed mix, I turn to more flexible-forms of inequality aversion. I focus on generational planners, whose inequality-aversion along the asset and productivity dimension is indexed by α_a and α_θ,

is also well-below the optimal B.

respectively. The exercise consists in finding the α_a and α_θ such that the solution to the first order conditions of the fully dynamic optimal policy problem are consistent with $p^* = p^{US}$ and $B^*/Y^* = B^{US}/Y^{US}$, where p^{US} and B^{US}/Y^{US} are the observed values of progressivity and debt-to-GDP in the US economy. Figure 12 presents the results.

The weights are "standard" if we restrict attention to the asset dimension: the asset poor are relatively more important than the asset-rich for the US planner, which is consistent with Rawlsian welfare criteria. Along the productivity dimension, US social preferences for redistribution appear inconsistent with both Utilitarian and Rawlsian criteria: the weights increase. To understand what drives this result, recall that inequality-averse planners favor lower levels of debt. But if we choose an aversion to inequality that makes the optimal debt-to-GDP consistent with the data, from Figure 6, we know that the optimal progressivity would exceed the one observed in the data. In other words, the US is not progressive enough for the level of debt it has. In order to be able to match the US tax system, the weights must *increase* along the productivity dimension. This, in turn, implies that the covariance between welfare weights and both asset and labor income is *positive*.

Figure 12: Implied welfare weights for the US



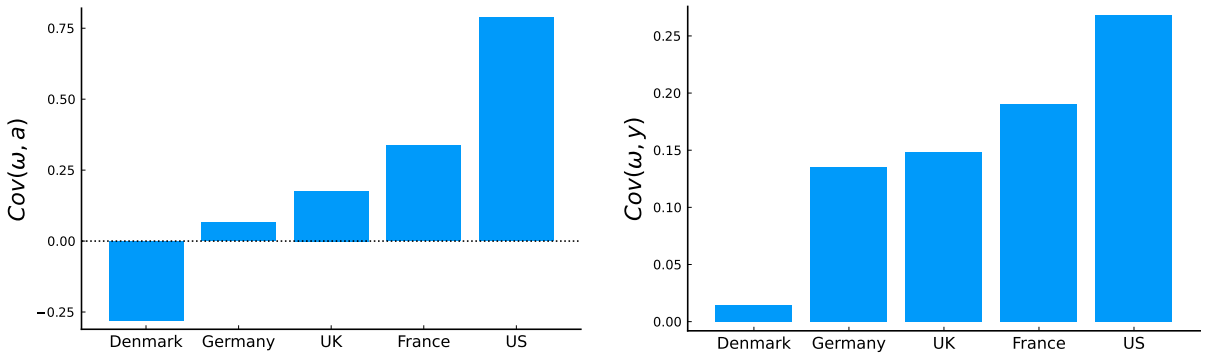
I also extend the exercise beyond the US, considering a collection of advanced economies that have consistent estimates for p and B/Y. I summarize the results by reporting the implied covariance between welfare weights and asset holdings, $Cov(\omega, a)$, as well as the covariance with labor income, $Cov(\omega, y)$. For the six advanced economies I consider, I find that welfare weights are inconsistent with Utilitarian or Rawlsian criteria: the implied correlation of welfare weights with both assets and labor income is *positive*. Interestingly, the ranking across countries in Figure 13 puts the US and Denmark at opposite ends of the spectrum. The estimated social preferences for the US are far from Utilitarian, with the weights covarying strongly with both assets and labor income. Denmark is closest to the Utilitarian benchmark, with welfare weights that are almost independent of assets and labor income.

In models of political economy with probabilistic voting a lá Persson and Tabellini (2002), the government chooses fiscal policy in order to maximize a weighted sum of agents' utilities. The weights capture the political power of different types of agents. In this sense, the weights can be interpreted as capturing the political influence of the rich and the poor. The fact that the US weights covary strongly with asset and labor income suggests that the US government is more responsive to the preferences of the rich. Denmark, on the other hand, is a country where the fiscal policy of the government responds to the preferences of all agents somewhat equally. Diving deeper into the political economy of these countries is beyond the scope of this paper, but I believe that this is an interesting direction for future research.

7 Conclusion

This paper studies the optimal mix of public debt and redistributive taxation in standard heterogeneous-agent incomplete-markets models. I find that inequality-averse planners should favor lower levels of public debt. This is due to an interest rate channel: progressive income taxes reduce the need to self-insure against idiosyncratic risk, lowering the demand for safe assets and increasing the fiscal cost of public debt. This result is robust to the presence of multiple safe assets, alternative social welfare functions, and restrictions

Figure 13: Inferred covariances of welfare weights and assets/income in advanced economies



(a) Implied $Cov(\omega, a)$

(b) $Cov(\omega, y)$

on the tax system.

Along the optimal transition, the interest rate channel continues to play a central role. The planner starts by flattening the tax schedule, reducing the insurance provided by the tax system. This pushes interest rates down and creates fiscal space, allowing the government to sustain lower taxes and higher consumption while gradually moving the debt-to-GDP ratio towards the long-run target. Ignoring transitional dynamics leads to underestimating the optimal level of public debt and overestimating the optimal degree of progressivity.

On the methodological side, I extend the sequence-space approach of Auclert et al. (2023a) to compute optimal transitions via policy gradient methods. Allowing for aversion to wealth inequality helps overcome the difficulties that arise when searching for a Ramsey steady state in this class of models.

Finally, the analysis abstracts from political economy considerations, which are likely to shape the planner's ability to commit to a long-run policy mix. This is an important direction for future work. Nevertheless, the results make clear that the level of public debt and the progressivity of the tax system should not be set independently: a government's stance on redistribution directly affects its capacity to borrow.

REFERENCES

Acikgoz, O., Hagedorn, M., Holter, H. A., and Wang, Y. (2023). The Optimum Quantity of Capital and Debt. *SSRN Electronic Journal*.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659.

Aiyagari, S. R. (1995). Optimal Capital Income Taxation with Incomplete Markets, Borrowing Constraints, and Constant Discounting. *Journal of Political Economy*, 103(6):1158–1175. Publisher: University of Chicago Press.

Aiyagari, S. R. and McGrattan, E. R. (1998). The optimum quantity of debt. *Journal of Monetary Economics*, 42(3):447–469.

Amol, A. and Luttmer, E. G. J. (2022). Permanent primary deficits, idiosyncratic long-run risk, and growth. Technical report, Federal Reserve Bank of Minneapolis.

Angeletos, G.-M., Collard, F., and Dellas, H. (2022). Public Debt as Private Liquidity: Optimal Policy. Technical Report w22794, National Bureau of Economic Research, Cambridge, MA.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Cai, M., Rognlie, M., and Straub, L. (2023a). Optimal Long-Run Fiscal Policy with Heterogeneous Agents.

Auclert, A., Rognlie, M., and Straub, L. (2023b). Determinacy and existence in the sequence space. *Manuscript*.

Auclert, A., Rognlie, M., and Straub, L. (2023c). The intertemporal keynesian cross. *Journal of Political Economy*, Forthcoming.

Bakış, O., Kaymak, B., and Poschke, M. (2015). Transitional dynamics and the optimal progressivity of income redistribution. *Review of Economic Dynamics*, 18(3):679–693.

Bewley, T. (1977). The permanent income hypothesis: A theoretical formulation. *Journal of Economic Theory*, 16(2):252–292.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2017a). Fiscal Policy and Debt Management with Incomplete Markets*. *The Quarterly Journal of Economics*, 132(2):617–663.

Bhandari, A., Evans, D., Golosov, M., and Sargent, T. J. (2017b). Public debt in economies with heterogeneous agents. *Journal of Monetary Economics*, 91:39–51.

Blanchard, O. (2019). Public Debt and Low Interest Rates. *American Economic Review*, 109(4):1197–1229.

Boar, C. and Midrigan, V. (2022). Efficient redistribution. *Journal of Monetary Economics*, 131:78–91.

Bourguignon, F. and Spadaro, A. (2012). Tax–benefit revealed social preferences. *The Journal of Economic Inequality*, 10(1):75–108.

Bénabou, R. (2002). Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency? *Econometrica*, 70(2):481–517.

Chang, Y. and Park, Y. (2021). Optimal Taxation with Private Insurance. *The Review of Economic Studies*, 88(6):2766–2798.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, 101(3):471–475.

Chien, Y. and Wen, Y. (2022). The Ramsey Steady-State Conundrum in Heterogeneous-Agent Economies. Technical report.

Davila, E. and Schaab, A. (2022). Welfare Assessments with Heterogeneous Individuals.

Davila, J., Hong, J. H., Krusell, P., and Ríos-Rull, J.-V. (2012). Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica*, 80(6):2431–2467.

De Ferra, S., Mitman, K. E., Romei, F., et al. (2021). *Why does capital flow from equal to unequal countries?* Centre for Economic Policy Research.

Diamond, P. (1965). National debt in a neoclassical growth model. *American Economic Review*, 55(5):1126–1150.

Diamond, P. A. and Mirrlees, J. A. (1971). Optimal Taxation and Public Production I: Production Efficiency. *The American Economic Review*, 61(1):8–27.

Dyrda, S. and Pedroni, M. (2022). Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Income Risk. *The Review of Economic Studies*, page rdac031.

Farhi, E. and Werning, I. (2007). Inequality and Social Discounting. *Journal of Political Economy*, 115(3):365–402.

Farhi, E. and Werning, I. (2013). Insurance and Taxation over the Life Cycle. *The Review of Economic Studies*, 80(2):596–635.

Ferriere, A., Grübener, P., Navarro, G., and Vardishvili, O. (2022). On the Optimal Design of Transfers and Income-Tax Progressivity.

Floden, M. and Lindé, J. (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic Dynamics*, 4(2):406–437.

Flodén, M. (2001). The effectiveness of government debt and transfers as insurance. *Journal of Monetary Economics*, 48(1):81–108.

Golosov, M., Tsyvinski, A., Werning, I., Diamond, P., and Judd, K. L. (2006). New Dynamic Public Finance: A User's Guide. *NBER Macroeconomics Annual*, 21:317–387. Publisher: The University of Chicago Press.

Guvenen, F., Ozkan, S., and Song, J. (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy*, 122(3):621–660. Publisher: The University of Chicago Press.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal Tax Progressivity: An Analytical Framework. *Quarterly Journal of Economics*, 132(4):1693–1754.

Heathcote, J., Storesletten, K., and Violante, G. L. (2020). Presidential Address 2019: How Should Tax Progressivity Respond to Rising Income Inequality? *Journal of the European Economic Association*, 18(6):2715–2754.

Heathcote, J. and Tsujiyama, H. (2021). Optimal Income Taxation: Mirrlees Meets Ramsey. *Journal of Political Economy*, pages 000–000.

Holmström, B. and Tirole, J. (1998). Private and Public Supply of Liquidity. *Journal of Political Economy*, 106(1):1–40. Publisher: The University of Chicago Press.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5):953–969.

Kaplan, G., Nikolakoudis, G., and Violante, G. L. (2023). Price Level and Inflation Dynamics in Heterogeneous Agent Economies.

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

LeGrand, F. and Ragot, X. (2023). Should we increase or decrease public debt? optimal fiscal policy with heterogeneous agents. Working Paper, SciencesPo.

Marcet, A. and Marimon, R. (2019). Recursive Contracts. *Econometrica*, 87(5):1589–1631. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA9902.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The Power of Forward Guidance Revisited. *American Economic Review*, 106(10):3133–3158.

Mian, A. R., Straub, L., and Sufi, A. (2022). A Goldilocks Theory of Fiscal Deficits.

Mirrlees, J. A. (1974). Notes on Welfare Economics, Information, and Uncertainty. In Balch, M., McFadden, D., and Wu, S. Y., editors, *Essays on Economic Behavior under Uncertainty*, page 0. North-Holland.

Persson, T. and Tabellini, G. (2002). *Political economics: explaining economic policy*. MIT press.

Phelan, C. (2006). Opportunity and Social Mobility. *Review of Economic Studies*.

Qiu, X. and Russo, N. (2022). Income Tax Progressivity: A Cross-Country Comparison.

Ragot, X. and Le Grand, F. (2023). Optimal Policies with Heterogeneous Agents: Truncation and Transitions. preprint, SSRN.

Saez, E. and Stantcheva, S. (2016). Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review*, 106(1):24–45.

Saez, E. and Zucman, G. (2019). *The triumph of injustice: How the rich dodge taxes and how to make them pay*. WW Norton & Company.

Schmitt-Grohé, S. and Uribe, M. (2003). Closing small open economy models. *Journal of international Economics*, 61(1):163–185.

Straub, L. and Werning, I. (2020). Positive Long-Run Capital Taxation: Chamley-Judd Revisited. *American Economic Review*, 110(1):86–119.

Trabandt, M. and Uhlig, H. (2011). The laffer curve revisited. *Journal of Monetary Economics*, 58(4):305–327.

Varian, H. R. (1980). Redistributive taxation as social insurance. *Journal of Public Economics*, 14(1):49–68.

Werning, I. (2007). Optimal Fiscal Policy with Redistribution. *The Quarterly Journal of Economics*, 122(3):925–967.

Woodford, M. (1990). Public Debt as Private Liquidity. *The American Economic Review*, 80(2):382–388. Publisher: American Economic Association.

A APPENDIX TO RAMSEY PROBLEM

This appendix provides additional details on the Ramsey problem and its solution. I start by deriving necessary conditions for the Ramsey steady state by following the sequence-space approach recently introduced by Auclert et al. (2023a). I then describe the algorithm to compute the optimal long-run mix of debt and progressivity. Finally, I discuss alternative formulations of the RSS problem and derive the sufficient-statistic representation of the optimal long-run value of public debt in Proposition 1.

A.1 OPTIMALITY CONDITIONS FOR RSS PROBLEM

Recall that the dynamic Ramsey problem is

$$\max_{\{r_t,B_t,p_t,\tau_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{r_s\},\{\tau_s\},\{p_s\}\right) \quad \text{s.t} \quad \begin{cases} \mathcal{A}_t\left(\{r_s\},\{\tau_s\},\{p_s\}\right) = B_t, \\ G + (1+r_{t-1})B_{t-1} = B_t + \mathcal{T}_t\left(\{r_s\},\{\tau_s\},\{p_s\}\right) \end{cases}, \forall t.$$

Given *sequences* for debt $\{B_t\}$ and progressivity $\{p_t\}$, one can take into account the period-by-period constraints by solving for *sequence space functions* $\boldsymbol{r}_t(\{p_s\},\{B_s\})$ and $\boldsymbol{\tau}_t(\{p_s\},\{B_s\})$. Then, the problem becomes

$$\max_{\{p_t,B_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{\boldsymbol{r}_s\left(\{p_u\},\{B_u\}\right)\},\{\boldsymbol{\tau}_s\left(\{p_u\},\{B_u\}\right)\},\{p_s\}\right).$$

Given an optimal plan $\{p_t^*, B_t^*\}$, a necessary condition for optimality is that any perturbation dp_u and dB_u shouldn't affect welfare. This yields the following pair of optimality conditions for any $u = 0, 1, \ldots$:

$$\sum_{t=0}^{\infty} \beta^{t-u} \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial r_s} \frac{\partial \boldsymbol{r}_s}{\partial p_u} + \sum_{t=0}^{\infty} \beta^{t-u} \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial \tau_s} \frac{\partial \boldsymbol{\tau}_s}{\partial p_u} + \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial p_u} = 0 \tag{15}$$

$$\sum_{t=0}^{\infty} \beta^{t-u} \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial r_s} \frac{\partial \boldsymbol{r}_s}{\partial B_u} + \sum_{t=0}^{\infty} \beta^{t-u} \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial \tau_s} \frac{\partial \boldsymbol{\tau}_s}{\partial B_u} = 0 \tag{16}$$

To simplify (15) and (16) we take $u \to \infty$. Now, for any sequence space function $F_t\left(\{X_s\}\right)$, define the discounted sum $\mathcal{E}_X^F$ of the long-run response,

$$\mathcal{E}_X^F \equiv \lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial F_t}{\partial X_u}.$$

Now, combining the *quasi-Toeplitz* property of Jacobians in stationary models[20] with the convolution theorem, we can write the composition of Jacobians as the product of discounted sums. That is,

$$\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial \mathcal{U}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial p_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial \mathcal{U}_t}{\partial r_u}\cdot\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial \boldsymbol{r}_t}{\partial p_u}=\mathcal{E}_r^{\mathcal{U}}\cdot\mathcal{E}_p^{\boldsymbol{r}} \tag{17}$$

To see why this is the case, observe that by defining $m=t-u$ and $n=s-u$,

$$\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial \mathcal{U}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial p_u}=\sum_{m=-u}^{\infty}\beta^{m}\sum_{n=-u}^{\infty}\frac{\partial \mathcal{U}_{m+u}}{\partial r_{n+u}}\frac{\partial \boldsymbol{r}_{n+u}}{\partial p_u}$$

The quasi-Toeplitz property of the Jacobians implies that as $u\to\infty$,

$$\sum_{m=-u}^{\infty}\beta^{m}\sum_{n=-u}^{\infty}\frac{\partial \mathcal{U}_{m+u}}{\partial r_{n+u}}\frac{\partial \boldsymbol{r}_{n+u}}{\partial p_u}\to\sum_{m=-\infty}^{\infty}\beta^{m}\sum_{n=-\infty}^{\infty}f_{m-n}g_n,$$

where $\boldsymbol{f}=\{f_t\}$ and $\boldsymbol{g}=\{g_t\}$ are the symbol vectors of $\frac{\partial \mathcal{U}_t}{\partial r_s}$ and $\frac{\partial \boldsymbol{r}_t}{\partial p_s}$, respectively. The term inside the summation is the convolution of the functions $\boldsymbol{f}$ and $\boldsymbol{g}$. Since quasi-Toeplitz Jacobians decay at least exponentially away from the diagonal, $\{f_t\}$ and $\{g_t\}$ are summable. The convolution theorem then implies that

$$\sum_{m=-\infty}^{\infty}\beta^{m}\sum_{n=-\infty}^{\infty}f_{m-n}g_n=\sum_{k=-\infty}^{\infty}\beta^{k}f_k\sum_{k=-\infty}^{\infty}\beta^{k}g_k$$

Now,

$$\sum_{k=-\infty}^{\infty}\beta^{k}f_k=\sum_{k=-\infty}^{\infty}\beta^{k}\lim_{u\to\infty}\frac{\partial \mathcal{U}_{k+u}}{\partial r_u}=\lim_{u\to\infty}\sum_{k=-u}^{\infty}\beta^{k}\frac{\partial \mathcal{U}_{k+u}}{\partial r_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial \mathcal{U}_t}{\partial r_u}=\mathcal{E}_r^{\mathcal{U}}\quad,$$

where the second equality follows from the dominated convergence theorem and the third equality by substituting $t=k+u$. Analogously,

$$\sum_{k=-\infty}^{\infty}\beta^{k}g_k=\sum_{k=-\infty}^{\infty}\beta^{k}\lim_{u\to\infty}\frac{\partial \boldsymbol{r}_{k+u}}{\partial p_u}=\lim_{u\to\infty}\sum_{k=-u}^{\infty}\beta^{k}\frac{\partial \boldsymbol{r}_{k+u}}{\partial p_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial \boldsymbol{r}_t}{\partial p_u}=\mathcal{E}_p^{\boldsymbol{r}}.$$

[20] An infinite matrix $\boldsymbol{J}=\{J_{t,s}\}\in\mathbb{R}^{\mathbb{N}\times\mathbb{N}}$ is a quasi-Toeplitz matrix with symbol vector $\boldsymbol{j}=\{j_t\}\in\mathbb{R}^{\mathbb{Z}}$ if i) for any $t,s\geq 0$ we have $\lim_{u\to\infty}J_{t+u,s+u}=j_{t-s}$ and ii) the entries of $J_{t,s}$ decay at least exponentially away from the diagonal. See Auclert et al. (2023a) and Auclert et al. (2023b) for details.

This establishes (17). Following similar steps, one can show that

$$\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial\mathcal{U}_t}{\partial\tau_s}\frac{\partial\boldsymbol{\tau}_s}{\partial p_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\mathcal{U}_t}{\partial\tau_u}\cdot\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\boldsymbol{\tau}_t}{\partial p_u}=\mathcal{E}_\tau^{\mathcal{U}}\cdot\mathcal{E}_p^{\boldsymbol{\tau}}$$
$$\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial\mathcal{U}_t}{\partial r_s}\frac{\partial\boldsymbol{r}_s}{\partial B_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\mathcal{U}_t}{\partial r_u}\cdot\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\boldsymbol{r}_t}{\partial B_u}=\mathcal{E}_r^{\mathcal{U}}\cdot\mathcal{E}_B^{\boldsymbol{r}}$$
$$\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial\mathcal{U}_t}{\partial\tau_s}\frac{\partial\boldsymbol{\tau}_s}{\partial B_u}=\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\mathcal{U}_t}{\partial\tau_u}\cdot\lim_{u\to\infty}\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\boldsymbol{\tau}_t}{\partial B_u}=\mathcal{E}_\tau^{\mathcal{U}}\cdot\mathcal{E}_B^{\boldsymbol{\tau}}.$$

This yields two necessary conditions for the Ramsey steady state

$$\mathcal{E}_r^{\mathcal{U}}\cdot\mathcal{E}_p^{\boldsymbol{r}}+\mathcal{E}_\tau^{\mathcal{U}}\cdot\mathcal{E}_p^{\boldsymbol{\tau}}+\mathcal{E}_p^{\mathcal{U}}=0 \tag{18}$$
$$\mathcal{E}_r^{\mathcal{U}}\cdot\mathcal{E}_B^{\boldsymbol{r}}+\mathcal{E}_\tau^{\mathcal{U}}\cdot\mathcal{E}_B^{\boldsymbol{\tau}}=0 \tag{19}$$

The discounted sum of the general equilibrium long-run responses $\mathcal{E}_p^{\boldsymbol{r}}$, $\mathcal{E}_p^{\boldsymbol{\tau}}$, $\mathcal{E}_B^{\boldsymbol{r}}$, and $\mathcal{E}_B^{\boldsymbol{\tau}}$ can then be computed from the system of equations derived below.

System for GE responses

I start by deriving a system of equations for $\mathcal{E}_p^{\boldsymbol{r}}$ and $\mathcal{E}_p^{\boldsymbol{\tau}}$. Perturbing the asset market-clearing condition at time t by dp_u,

$$\sum_{s=0}^{\infty}\frac{\partial\mathcal{A}_t}{\partial r_s}\frac{\partial\boldsymbol{r}_s}{\partial p_u}+\sum_{s=0}^{\infty}\frac{\partial\mathcal{A}_t}{\partial\tau_s}\frac{\partial\boldsymbol{\tau}_s}{\partial p_u}+\frac{\partial\mathcal{A}_t}{\partial p_u}=0,\ \forall t.$$

Multiply condition at time t by β^{t-u} and sum across t to get

$$\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial\mathcal{A}_t}{\partial r_s}\frac{\partial\boldsymbol{r}_s}{\partial p_u}+\sum_{t=0}^{\infty}\beta^{t-u}\sum_{s=0}^{\infty}\frac{\partial\mathcal{A}_t}{\partial\tau_s}\frac{\partial\boldsymbol{\tau}_s}{\partial p_u}+\sum_{t=0}^{\infty}\beta^{t-u}\frac{\partial\mathcal{A}_t}{\partial p_u}=0.$$

Letting $u\to\infty$ and using the same results we relied on when simplifying the FOCs, this becomes

$$\mathcal{E}_r^{\mathcal{A}}\cdot\mathcal{E}_p^{\boldsymbol{r}}+\mathcal{E}_\tau^{\mathcal{A}}\cdot\mathcal{E}_p^{\boldsymbol{\tau}}+\mathcal{E}_p^{\mathcal{A}}=0.$$

Following the same steps with the government's budget constraint yields

$$\left(\mathcal{E}_r^{\mathcal{T}}-B\right)\mathcal{E}_p^{\boldsymbol{r}}+\mathcal{E}_\tau^{\mathcal{T}}\cdot\mathcal{E}_p^{\boldsymbol{\tau}}+\mathcal{E}_p^{\mathcal{T}}=0.$$

Thus, we can solve for $\mathcal{E}_p^{\boldsymbol{r}}$ and $\mathcal{E}_p^{\boldsymbol{\tau}}$ from

$$\begin{bmatrix}\mathcal{E}_r^{\mathcal{A}} & \mathcal{E}_\tau^{\mathcal{A}}\\ \mathcal{E}_r^{\mathcal{T}}-B & \mathcal{E}_\tau^{\mathcal{T}}\end{bmatrix}\begin{bmatrix}\mathcal{E}_p^{\boldsymbol{r}}\\ \mathcal{E}_p^{\boldsymbol{\tau}}\end{bmatrix}=\begin{bmatrix}-\mathcal{E}_p^{\mathcal{A}}\\ -\mathcal{E}_p^{\mathcal{T}}\end{bmatrix}. \tag{20}$$

Analogous derivations for a perturbation dB_u lead to a system for $\mathcal{E}_B^{\boldsymbol{r}}$ and $\mathcal{E}_B^{\boldsymbol{\tau}}$

$$\begin{bmatrix} \mathcal{E}_r^{\mathcal{A}} & \mathcal{E}_\tau^{\mathcal{A}} \\ \mathcal{E}_r^{\mathcal{T}} - B & \mathcal{E}_\tau^{\mathcal{T}} \end{bmatrix} \begin{bmatrix} \mathcal{E}_B^{\boldsymbol{r}} \\ \mathcal{E}_B^{\boldsymbol{\tau}} \end{bmatrix} = \begin{bmatrix} 1 \\ \beta(1+r) - 1 \end{bmatrix}. \tag{21}$$

A.2 Computing the RSS

The algorithm to solve for the RSS relies on (18) and (19). To operationalize these equations, it is necessary to compute the discounted sum of the long-run response of aggregate outcomes $\mathcal{Y} \in \{\mathcal{U}, \mathcal{A}, \mathcal{T}\}$ to changes in interest rates and changes in the level and the progressivity of the tax system. In principle, this could be done by computing the Jacobians for e.g. aggregate welfare, assets, and taxes around the steady state implied by some candidate fiscal policy, and then taking the discounted sum of some far out column. However, this turns out to be too costly since we need a pretty large horizon to get convergence of the discounted sum. Given this, Auclert et al. (2023a) propose the following procedure for the long-run responses. For concreteness, I focus on the discounted sum of the long-run response of aggregate asset demand to changes in interest rates $\mathcal{E}_r^{\mathcal{A}}$. The same procedure applies to the other responses.

- Iterate backward to find perturbation $d\boldsymbol{a}^s(\theta, a)$ to policy function when shock dr is $s = 0, 1, \ldots$ periods in the future. Then, sum across periods

$$d\boldsymbol{a}(\theta, a) = \sum_{s=0}^{\infty} \beta^{-s} d\boldsymbol{a}^s(\theta, a).$$

- Using $d\boldsymbol{a}(\theta, a)$ as the perturbation to asset policy function, iterate forward to find implied change in distribution D for $s = 1, 2, \ldots$ periods in the future. Take sum

$$dD(\theta, a) = \sum_{s=1}^{\infty} \beta^s dD^s(\theta, a).$$

- Finally, aggregate $dA = d\boldsymbol{a} \cdot D^{ss} + dD \cdot \boldsymbol{a}^{ss}$.[21] Then, $\mathcal{E}_r^{\mathcal{A}} = \frac{dA}{dr}$.

With this in hand, I solve for the RSS as follows:

[21] Here, D^{ss} and $\boldsymbol{a}^{ss}$ are the distribution and policy function in the steady state implied by the candidate fiscal policy.

1. Given a candidate level of public debt B and progressivity p, solve for the interest rate r and the level of taxes τ that ensure asset market clearing and government budget balance.

2. Compute the discounted sum of the long-run responses of aggregate assets $\mathcal{E}_r^{\mathcal{A}}$, $\mathcal{E}_\tau^{\mathcal{A}}$, $\mathcal{E}_p^{\mathcal{A}}$, aggregate utility $\mathcal{E}_r^{\mathcal{U}}$, $\mathcal{E}_\tau^{\mathcal{U}}$, $\mathcal{E}_p^{\mathcal{U}}$, and taxes $\mathcal{E}_r^{\mathcal{T}}$, $\mathcal{E}_\tau^{\mathcal{T}}$, $\mathcal{E}_p^{\mathcal{T}}$ using the procedure outline above.

3. Solve for the discounted sum of the general equilibrium long-run responses using (20) and (21).

4. Check whether (18) and (19) are satisfied. If not, update r and p and repeat steps 1-4.

In practice, to reduce the complexity of general equilibrium, it is better to iterate on r and p, and then use asset market clearing to read off the implied level of public debt B. This means that in step (1) above, we only need to solve for the level of taxes τ that ensures government budget balance. Moreover, I use a two-step procedure to avoid multi-dimensional root finding algorithms. So in step (4), I first check whether (18) is satisfied. If this is not the case, I adjust p and repeat the process but keeping r fixed. Because this is a one-dimensional problem, the updating for p can be done via Brent's method. Once (18) is satisfied at the candidate level of r, I check whether (19) is satisfied. If this is not the case, I adjust r and repeat the process, resolving for p along the way. As explained in the main text, I verify that there is a unique interior solution to (19), fixing p. I also verify that there is a unique solution to (18). See Appendix A.5 for a figure.

A.3 Alternative formulation of the RSS problem

To simplify the derivations and side-step the complexity introduced by the presence of general equilibrium responses in (15) and (16), we can proceed as follows. By Walras' Law, problem (6) is equivalent to

$$\max_{\{r_t, B_t, p_t, \tau_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) \quad \text{s.t} \quad \begin{cases} \mathcal{A}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = B_t, \\ G + \mathcal{C}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = \mathcal{Y}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) \end{cases} \quad \forall t.$$

Now, given choices of $\{r_t, \tau_t, p_t\}$, the asset market clearing condition pins down $\{B_t\}$. After dropping all redundant choice variables and the associated constraints, the problem reduces to

$$\max_{\{r_t, p_t, \tau_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) \quad \text{s.t} \quad G + \mathcal{C}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = \mathcal{Y}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right).$$

The Lagrangian for this problem is

$$\max_{\{r_t, \tau_t, p_t\}} \sum_{t=0}^{\infty} \beta^t \left\{ \mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) + \lambda_t^{GM} \{\mathcal{Y}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) - \mathcal{C}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) - G\} \right\}$$

The first-order conditions are

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial r_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial r_u} - \frac{\partial \mathcal{C}_t}{\partial r_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{22}$$

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial \tau_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial \tau_u} - \frac{\partial \mathcal{C}_t}{\partial \tau_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{23}$$

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial p_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial p_u} - \frac{\partial \mathcal{C}_t}{\partial p_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots, \tag{24}$$

together with goods-market clearing period-by-period. To characterize the steady state, we take $u \to \infty$ and assume the multipliers converge: $\lambda_t^{GM} \to \lambda^{GM} > 0$ as $t \to \infty$. The first term in (22) clearly goes to $\mathcal{E}_r^{\mathcal{U}}$. Now consider the second term, $\sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \frac{\partial \mathcal{Y}_t}{\partial r_u}$. Since $\frac{\partial \mathcal{Y}_t}{\partial r_s}$ is a quasi-Toeplitz matrix with some symbol vector $\boldsymbol{y} = \{y_t\}$, the matrix $E_{t,s} \equiv \lambda_t^{GM} \frac{\partial \mathcal{Y}_t}{\partial r_s}$ is also quasi-Toeplitz with symbol vector $\lambda^{GM} \boldsymbol{y} = \{\lambda^{GM} y_t\}$. Therefore,

$$\lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \frac{\partial \mathcal{Y}_t}{\partial r_u} = \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} E_{t,u} = \lim_{u \to \infty} \sum_{m=-u}^{\infty} \beta^m E_{m+u,m},$$

where I have substituted $t = m + u$. Following the logic in Appendix A.1

$$\sum_{m=-u}^{\infty} \beta^m E_{m+u,m} \to \sum_{m=-\infty}^{\infty} \beta^m \lambda^{GM} y_m = \lambda^{GM} \sum_{m=-\infty}^{\infty} \beta^m y_m = \lambda^{GM} \mathcal{E}_r^{\mathcal{Y}},$$

as $u \to \infty$. Analogously, the third term goes to $\lambda^{GM} \mathcal{E}_r^{\mathcal{C}}$. This establishes that the first-order condition for the interest rate in the steady state becomes

$$\mathcal{E}_r^{\mathcal{U}} + \lambda^{GM} \left(\mathcal{E}_r^{\mathcal{Y}} - \mathcal{E}_r^{\mathcal{C}} \right) = 0, \tag{25}$$

Similar reasoning can be applied to (23) and (24) to obtain the steady-state conditions for the tax rate and the progressivity of the tax system:

$$\mathcal{E}_{\tau}^{\mathcal{U}} + \lambda^{GM}\left(\mathcal{E}_{\tau}^{\mathcal{Y}} - \mathcal{E}_{\tau}^{\mathcal{C}}\right) = 0, \tag{26}$$

$$\mathcal{E}_{p}^{\mathcal{U}} + \lambda^{GM}\left(\mathcal{E}_{p}^{\mathcal{Y}} - \mathcal{E}_{p}^{\mathcal{C}}\right) = 0. \tag{27}$$

Using (26) to eliminate λ^{GM},

$$\mathcal{E}_{r}^{\mathcal{U}} = \frac{\mathcal{E}_{\tau}^{\mathcal{U}}}{\mathcal{E}_{\tau}^{\mathcal{C}} - \mathcal{E}_{\tau}^{\mathcal{Y}}}\left(\mathcal{E}_{r}^{\mathcal{C}} - \mathcal{E}_{r}^{\mathcal{Y}}\right), \tag{28}$$

$$\mathcal{E}_{p}^{\mathcal{U}} = \frac{\mathcal{E}_{\tau}^{\mathcal{U}}}{\mathcal{E}_{\tau}^{\mathcal{C}} - \mathcal{E}_{\tau}^{\mathcal{Y}}}\left(\mathcal{E}_{p}^{\mathcal{C}} - \mathcal{E}_{p}^{\mathcal{Y}}\right). \tag{29}$$

Equations (28) and (29), together with the goods market-clearing condition in steady state, can be used to search for a candidate Ramsey steady state $\{r, \tau, p\}$. The optimal level of debt can then be read-off the asset-market clearing condition. The advantage of this formulation is that it does not require solving for general equilibrium in each iteration. Moreover, the discounted sums of the general equilibrium responses do not enter the equations and therefore the mysterious step relying on the convolution theorem is no longer required. The disadvantage is that it requires solving for a larger system. It turns out, however, that one can use homogeneity of aggregate consumption and output to reduce the dimensionality of the system (cf. Auclert et al., 2023c). I verify that the solution implied by this formulation is consistent with the one in Appendix A.1.

A.4 Proof of Proposition 1

Recall that, for any $u = 0, 1, \ldots$, the following must be true

$$\sum_{t=0}^{\infty}\sum_{s=0}^{\infty}\beta^{t-u}\frac{\partial \mathcal{U}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \sum_{t=0}^{\infty}\sum_{s=0}^{\infty}\beta^{t-u}\lambda_t\frac{\partial \mathcal{T}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \lambda_u - \beta\lambda_{u+1}(1+\boldsymbol{r}_u) - \sum_{t=0}^{\infty}\beta^{t-u}\lambda_t\frac{\partial \boldsymbol{r}_t}{\partial B_u}B_{t-1} = 0. \tag{30}$$

These conditions are necessary for the Ramsey plan to be optimal. If they are not satisfied, a small perturbation dB_u would increase social welfare, contradicting the optimality of the Ramsey plan. To derive (9), we let $u \to \infty$ and assume that $\lambda_t \to \lambda^{RSS} > 0$ and $B_t \to B^{RSS} < \infty$ as $t \to \infty$.

Using the quasi-Toeplitz property of Jacobians, together with the convolution theorem, as in Appendix A.1, verifies that the first term in (30) goes to $\mathcal{E}_r^{\mathcal{U}} \cdot \mathcal{E}_B^{\boldsymbol{r}}$. The third term goes to $\lambda^{RSS}\{1-\beta(1+\boldsymbol{r})\}$.

For the second term, define the matrix $M_{t,s} \equiv \lambda_t \frac{\partial \mathcal{T}_t}{\partial r_s}$. Since $\frac{\partial \mathcal{T}_t}{\partial r_s}$ is a quasi-Toeplitz matrix with some symbol vector $\boldsymbol{h} = \{h_t\}$, $M_{t,s}$ is also quasi-Toeplitz with symbol vector $\lambda^{RSS}\boldsymbol{h} = \{\lambda^{RSS} h_t\}$. Therefore,

$$\lim_{u\to\infty} \sum_{t=0}^{\infty}\sum_{s=0}^{\infty} \beta^{t-u}\lambda_t \frac{\partial \mathcal{T}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} = \lim_{u\to\infty} \sum_{t=0}^{\infty}\sum_{s=0}^{\infty} \beta^{t-u} M_{t,s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} = \lim_{u\to\infty} \sum_{m=-u}^{\infty} \beta^m \sum_{n=-u}^{\infty} M_{m+u,n+u}\frac{\partial \boldsymbol{r}_{n+u}}{\partial B_u},$$

where I have substituted $t = m+u$ and $s = n+u$. Given the quasi-Toeplitz property of $M_{t,s}$ and $\frac{\partial \boldsymbol{r}_t}{\partial B_s}$,

$$\sum_{m=-u}^{\infty} \beta^m \sum_{n=-u}^{\infty} M_{m+u,n+u}\frac{\partial \boldsymbol{r}_{n+u}}{\partial B_u} \to \sum_{m=-\infty}^{\infty} \beta^m \sum_{n=-\infty}^{\infty} \lambda^{RSS} h_{m-n} g_n = \lambda^{RSS} \sum_{m=-\infty}^{\infty} \beta^m \sum_{n=-\infty}^{\infty} h_{m-n} g_n,$$

where $\{g_t\}$ is the symbol vector of $\frac{\partial \boldsymbol{r}_t}{\partial B_s}$. The term inside the second summation is the convolution of two summable sequences. Applying the convolution theorem and following the logic in Appendix A.1,

$$\lambda^{RSS} \sum_{m=-\infty}^{\infty} \beta^m \sum_{n=-\infty}^{\infty} h_{m-n} g_n = \lambda^{RSS} \sum_{k=-\infty}^{\infty} \beta^k h_k \sum_{k=-\infty}^{\infty} \beta^k g_k = \lambda^{RSS} \mathcal{E}_r^{\mathcal{T}} \cdot \mathcal{E}_B^{\boldsymbol{r}}$$

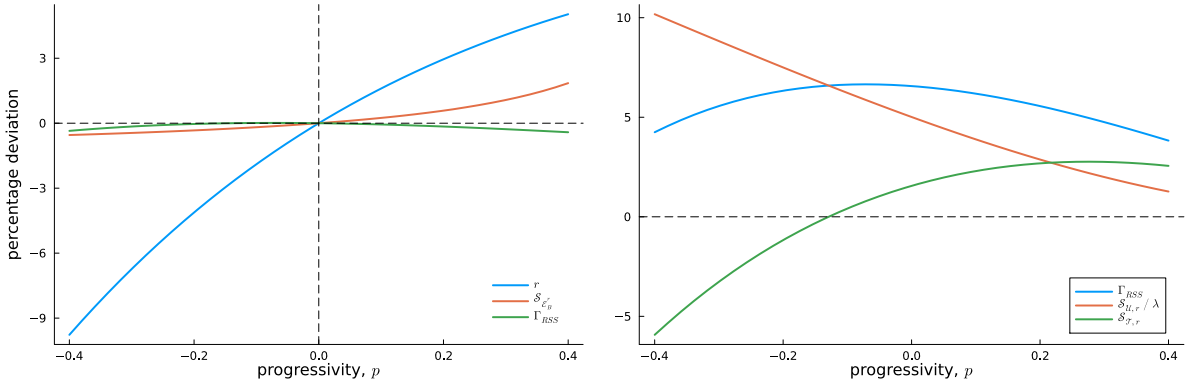
This establishes that the second term in (30) goes to $\lambda^{RSS}\mathcal{E}_r^{\mathcal{T}} \cdot \mathcal{E}_B^{\boldsymbol{r}}$. Analogously, fourth term in (30) goes to $\lambda^{RSS}\mathcal{E}_B^{\boldsymbol{r}} B^{RSS}$. Therefore, as $u \to \infty$, (30) becomes

$$\mathcal{E}_r^{\mathcal{U}} \cdot \mathcal{E}_B^{\boldsymbol{r}} + \lambda^{RSS}\mathcal{E}_r^{\mathcal{T}} \cdot \mathcal{E}_B^{\boldsymbol{r}} + \lambda^{RSS}\{1-\beta(1+\boldsymbol{r})\} - \lambda^{RSS}\mathcal{E}_B^{\boldsymbol{r}}\, B^{RSS} = 0. \tag{31}$$

Dividing through by λ^{RSS} gives (9).

A.5 ADDITIONAL RSS FIGURES

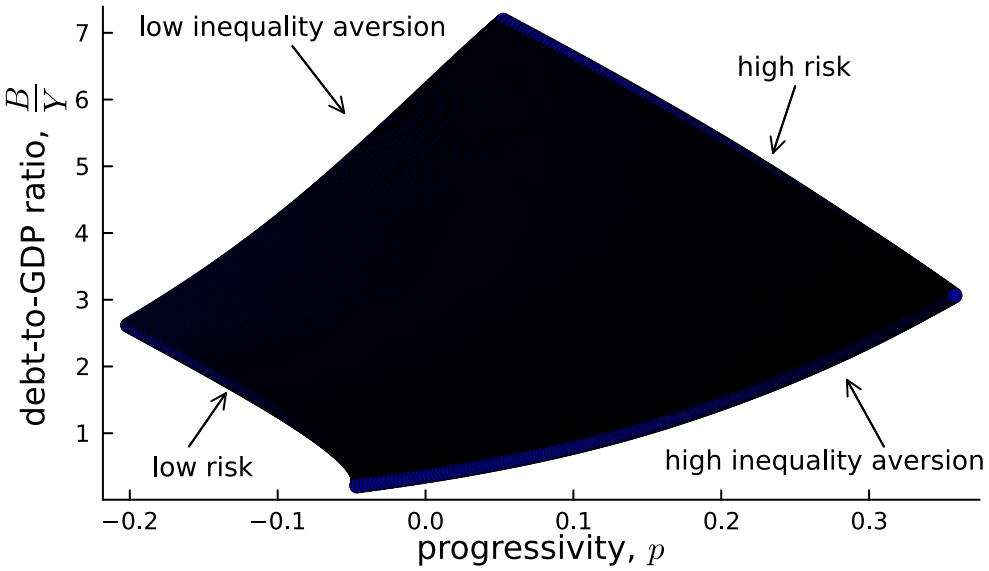
Figure 14: Sufficient statistics and progressivity in the RSS



(a) r, $\mathcal{E}_B^{\boldsymbol{r}}$, and Γ^{RSS} as progressivity varies

(b) Unpacking $\Gamma^{RSS} \equiv \frac{\mathcal{E}_r^{\mathcal{U}}}{\lambda^{RSS}} + \mathcal{E}_r^{\mathcal{T}}$

Figure 15: Optimal long-run mix of debt and progressivity and idiosyncratic risk



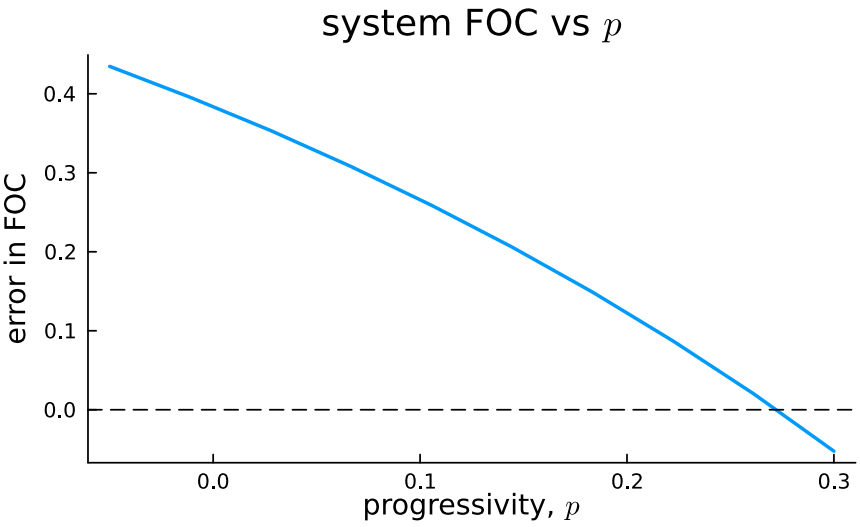


Figure 16: Verifying uniqueness of solution to (18)

B APPENDIX TO OPTIMAL STEADY STATE PROBLEM

This appendix provides additional details on the OSS problem and its solution. I start by deriving the optimality conditions for the OSS problem. I then describe the algorithm to compute the optimal level of public debt and progressivity. Finally, I discuss an alternative formulation of the OSS problem that allows for a sufficient-statistic representation of the optimality conditions that is similar to Proposition 1.

B.1 OPTIMALITY CONDITIONS FOR OSS PROBLEM

To take into account constraints in problem (11), for each (B, p), solve for *functions* $\boldsymbol{r}(B, p)$ and $\boldsymbol{\tau}(B, p)$. Then, the OSS problem reduces to an unconstrained maximization problem:

$$\max_{B,p} \mathcal{W}(\boldsymbol{r}(B, p), \boldsymbol{\tau}(B, p), p).$$

The first-order conditions with respect to p and B are, respectively:

$$\frac{\partial \mathcal{W}}{\partial r} \cdot \frac{\partial \boldsymbol{r}}{\partial p} + \frac{\partial \mathcal{W}}{\partial p} + \frac{\partial \mathcal{W}}{\partial \tau} \cdot \frac{\partial \boldsymbol{\tau}}{\partial p} = 0 \tag{32}$$

$$\frac{\partial \mathcal{W}}{\partial r} \cdot \frac{\partial \boldsymbol{r}}{\partial B} + \frac{\partial \mathcal{W}}{\partial \tau} \cdot \frac{\partial \boldsymbol{\tau}}{\partial B} = 0 \tag{33}$$

To solve for the GE derivatives in (32), one can use part one of Lemma 1 in the main text. The GE derivatives in (33) can be obtained from the following system:

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial r} & \frac{\partial \mathcal{A}}{\partial \tau} \\ \frac{\partial \mathcal{T}}{\partial r} - B & \frac{\partial \mathcal{T}}{\partial \tau} \end{bmatrix} \begin{bmatrix} \frac{\partial \boldsymbol{r}}{\partial B} \\ \frac{\partial \boldsymbol{\tau}}{\partial B} \end{bmatrix} = \begin{bmatrix} 1 \\ r \end{bmatrix}. \tag{34}$$

I use (32) and (33) to reduce the optimization problem in (11) to a (numerical) root-finding problem.[22]

B.2 Computing the OSS

The algorithm to compute the OSS proceeds as follows:

1. Given a candidate level of public debt B and progressivity p, solve for the interest rate r and the level of taxes τ that ensure asset market clearing and government budget balance.

2. Compute the partial equilibrium derivatives of aggregate welfare $\frac{\partial \mathcal{W}}{\partial r}$, $\frac{\partial \mathcal{W}}{\partial \tau}$, $\frac{\partial \mathcal{W}}{\partial p}$, aggregate tax revenues $\frac{\partial \mathcal{T}}{\partial r}$, $\frac{\partial \mathcal{T}}{\partial \tau}$, $\frac{\partial \mathcal{T}}{\partial p}$, and aggregate asset demand $\frac{\partial \mathcal{A}}{\partial r}$, $\frac{\partial \mathcal{A}}{\partial \tau}$, $\frac{\partial \mathcal{A}}{\partial p}$ via numerical differentiation.

3. Use Lemma 1 for the general equilibrium derivatives $\frac{\partial \boldsymbol{r}}{\partial p}$ and $\frac{\partial \boldsymbol{\tau}}{\partial p}$ and (34) for $\frac{\partial \boldsymbol{r}}{\partial B}$ and $\frac{\partial \boldsymbol{\tau}}{\partial B}$.

4. Check whether (32) and (33) are satisfied. If they are, stop. Otherwise, adjust B and p and repeat the process.

As in Appendix A.2, it is better to iterate on r and p, and then use asset market clearing to read off the implied level of public debt B. This means that in step (1) above, we only need to solve for the level of taxes τ that ensures government budget balance. Once more, to avoid multi-dimensional root-finding algorithms, I separate the problem into two steps. So in step (4), I first check whether (32) is satisfied. If this is not the case, I

[22] I check that the solution is unique, so there is no need to worry about multiple roots. Even if this were the case, one could easily rank them by evaluating the objective function because doing so is not computationally hard.

adjust p and repeat the process but keeping r fixed. Because this is a one-dimensional problem, the updating for p can be done via Brent's method.[23] Once (32) is satisfied at the candidate level of r, I check whether (33) is satisfied. If this is not the case, I adjust r and repeat the process, resolving for p along the way. Finally, note that step (3) does not require re-calculating the equilibrium, as the GE derivatives can be expressed in terms of PE derivatives.

B.3 Alternative OSS formulation

To avoid solving for general equilibrium in each iteration, one can also work with the goods market clearing condition. Indeed, by Walras' Law, the OSS problem in (11) is equivalent to:

$$\max_{\{r,\tau,p,B\}} \mathcal{W}(r,\tau,p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(r,\tau,p) = B, \\ \mathcal{C}(r,\tau,p) + G = \mathcal{Y}(r,\tau,p) \end{cases} . \tag{35}$$

Notice that, given an interest rate r and a CRP tax code $\{\tau, p\}$, the planner can always find a level of public debt B to ensure asset market clearing holds. After dropping this constraint and the associated choice variable, the problem reduces to

$$\max_{\{r,\tau,p\}} \mathcal{W}(r,\tau,p) \quad \text{s.t} \quad \mathcal{C}(r,\tau,p) + G = \mathcal{Y}(r,\tau,p).$$

The Lagrangian for this problem is

$$\max_{\{r,\tau,p\}} \mathcal{W}(r,\tau,p) + \lambda^{GM} \left\{ \mathcal{Y}(r,\tau,p) - \mathcal{C}(r,\tau,p) - G \right\}.$$

The associated first-order conditions are

$$\begin{aligned} \frac{\partial \mathcal{W}}{\partial \tau} + \lambda^{GM} \left\{ \frac{\partial \mathcal{Y}}{\partial \tau} - \frac{\partial \mathcal{C}}{\partial \tau} \right\} = 0, \\ \frac{\partial \mathcal{W}}{\partial r} + \lambda^{GM} \left\{ \frac{\partial \mathcal{Y}}{\partial r} - \frac{\partial \mathcal{C}}{\partial r} \right\} = 0, \\ \frac{\partial \mathcal{W}}{\partial p} + \lambda^{GM} \left\{ \frac{\partial \mathcal{Y}}{\partial p} - \frac{\partial \mathcal{C}}{\partial p} \right\} = 0, \end{aligned}$$

[23] Here, I also check that there is a unique value of p that solves (32).

together with the goods market clearing condition. Using the first one to eliminate λ^{GM}, I arrive to

$$\frac{\partial \mathcal{W}}{\partial r} = \frac{\frac{\partial \mathcal{W}}{\partial \tau}}{\frac{\partial \mathcal{C}}{\partial \tau} - \frac{\partial \mathcal{Y}}{\partial \tau}} \left\{ \frac{\partial \mathcal{C}}{\partial r} - \frac{\partial \mathcal{Y}}{\partial r} \right\} \tag{36}$$

$$\frac{\partial \mathcal{W}}{\partial p} = \frac{\frac{\partial \mathcal{W}}{\partial \tau}}{\frac{\partial \mathcal{C}}{\partial \tau} - \frac{\partial \mathcal{Y}}{\partial \tau}} \left\{ \frac{\partial \mathcal{C}}{\partial p} - \frac{\partial \mathcal{Y}}{\partial p} \right\} \tag{37}$$

Equations (36) and (37), combined with the goods market clearing condition, can be used to solve for the three unknowns $\{r, \tau, p\}$. As before, the advantage of this formulation is that it does not require solving for general equilibrium in each iteration. I use homogeneity of aggregate consumption and output to reduce the dimensionality of the system and verify that the solution implied by this formulation is consistent with the solution obtained via the algorithm in B.2.

B.4 Sufficient statistic representation of optimality conditions

I derive a simple "sufficient-statistic" representation for the optimal choice of debt B that is analogous to the one derived in Proposition 1. Fix p and τ, and zoom-in on the optimal choice of B. Let $\boldsymbol{r}(B, \tau, p)$ denote the interest rate that clears the asset market given the fiscal policy of the government. Then, letting λ denote the Lagrange multiplier on the government's budget constraint, the optimal choice of debt is

$$B^{OSS} = \arg\max_{B} \left\{ \mathcal{W}(\boldsymbol{r}(B, \tau, p), \tau, p) + \lambda \left(\mathcal{T}(\boldsymbol{r}(B, \tau, p), \tau, p) - \boldsymbol{r}(B, \tau, p) B - G \right) \right\}. \tag{38}$$

At an interior solution, the first order condition for this problem implies

$$\mathcal{W}_r \frac{\partial \boldsymbol{r}}{\partial B} + \lambda \left\{ \mathcal{T}_r \frac{\partial \boldsymbol{r}}{\partial B} - \frac{\partial \boldsymbol{r}}{\partial B} B^{OSS} - \boldsymbol{r} \right\} = 0. \tag{39}$$

Define the social marginal value of public debt, in dollar terms, as $\Gamma \equiv \frac{\mathcal{W}_r}{\lambda} + \mathcal{T}_r$. Similarly, let $\epsilon_B^r \equiv \frac{\partial \boldsymbol{r}}{\partial B} \frac{B^{OSS}}{1+\boldsymbol{r}}$ denote the elasticity of interest rates with respect to changes in the level of public debt *at the optimum*. Then, (39) can be written as

$$\Gamma \cdot \epsilon_B^r - \frac{\boldsymbol{r}}{1+\boldsymbol{r}} B^{OSS} - \epsilon_B^r \cdot B^{OSS} = 0.$$

Solving for B^{OSS} yields:

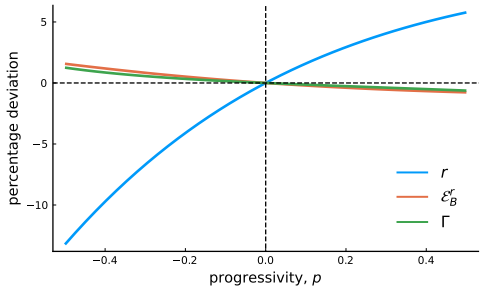
$$B^{OSS} = \frac{\epsilon_B^r}{\epsilon_B^r + \frac{\boldsymbol{r}}{1+\boldsymbol{r}}} \times \Gamma. \tag{40}$$

This equation shows that the optimal level of public debt depends on three objects: the social marginal value of public debt Γ, the *elasticity* of interest rates with respect to changes in the level of public debt $\mathcal{E}_B^r$, and the *level of interest rates* $\boldsymbol{r}$. This is similar to the sufficient-statistic representation derived in Proposition 1. The difference, however, is that the premium no longer appears in this equation. Intuitively, this is because debt only enters as a cost in the optimal steady state problem.

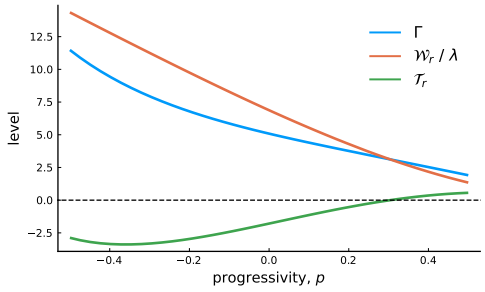
A higher interest rate unambigously decreases the optimal level of public debt, holding everything else equal. Somewhat more subtle, a higher elasticity $\mathcal{E}_B^r$ increases B^{OSS} when interest rates are positive. This is not obvious, since this object affects both the costs and benefits of public debt. But at an interior optimum, if $\boldsymbol{r} \geq 0$, it must be that $\Gamma \geq B^{OSS}$. Therefore, a more elastic interest rate increases the marginal benefit of public debt by more than it increases the fiscal cost, pushing towards higher B^{OSS}. Finally, when the social marginal value of public debt Γ increases, the optimal level of public debt also increases, holding everything else fixed.

This discussion implies that the progressivity of the tax system affects the optimal level of public debt through three channels: $\boldsymbol{r}$, $\mathcal{E}_B^r$, and Γ. The interest-rate channel, emphasized throughout the paper, implies that a more progressive tax system increases $\boldsymbol{r}$ and hence lowers the optimal level of public debt. Pushing in the same direction, a more progressive p lowers the marginal *private value* of B, as it helps agents insure against risk and thus reduces the value of public liquidity. This reinforces the effects driven by the interest rate channel. With income effects, however, the response of the marginal *social value* Γ, that includes revenue effects, can be less obvious. The effect of p on $\mathcal{E}_B^r$ is more complex. But this elasticity is very low in the model (compared to its data counterpart) and not too responsive to changes in p. So this channel does not seem to play a major role.

Figure 17: Sufficient statistics and progressivity in the OSS



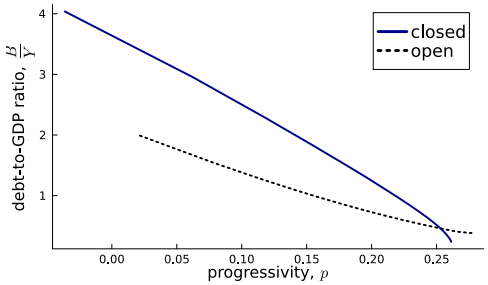
(a) r, $\mathcal{E}_B^r$, and Γ as progressivity varies



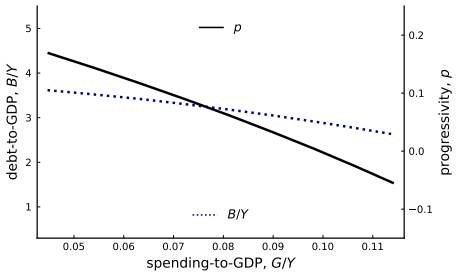
(b) Unpacking Γ as progressivity varies

B.5 ADDITIONAL OSS FIGURES

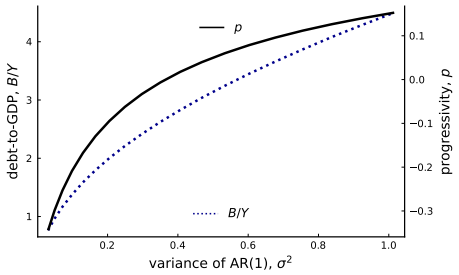
Figure 18: Comparatives statics with respect to spending, income risk, and openness



(a) Optimal mix in open and closed economies



(b) Optimal mix and government spending



(c) Optimal mix and income risk

C APPENDIX TO MODEL WITH CAPITAL

This appendix provides details on the model with capital. Appendix C.1 defines the equilibrium for this economy and describes the calibration of the model with capital. Appendix C.2 presents the dynamic Ramsey problem in the model with capital and proves the optimality of the modified golden rule. Finally, Appendix C.4 derives the full set of optimality conditions and details the computational procedure to solve for the Ramsey steady state in the model with capital.

C.1 EQUILIBRIUM DEFINITION AND CALIBRATION FOR MODEL WITH CAPITAL

Definition 2 An equilibrium in the model with capital is a sequence of prices $\{\bar{r}_t, \bar{w}_t, r_t, w_t\}$ and tax schedules $\{p_t, \tau_t, \tau_t^k\}$, a sequence of policy functions $\{\boldsymbol{c}_t(x), \boldsymbol{a}_t(x), \boldsymbol{y}_t(x)\}$, aggregates $\{L_t, K_t, B_t\}$, and distributions $\{D_t\}$ such that, given the initial distribution D_0 and the initial level of capital K_0, the following conditions hold:

(i) $\{\boldsymbol{c}_t(x), \boldsymbol{a}_t(x), \boldsymbol{y}_t(x)\}$ are optimal given $\{\bar{r}_t, \bar{w}_t\}$,

(ii) D_t is consistent with the policy functions and the Markov process for productivity,

(iii) after-tax wages $\bar{w}_t = \tau_t w_t^{1-p_t}$ and after-tax interest rates $\bar{r}_t = (1 - \tau_{kt}) r_t$,

(iv) before-tax wages $w_t = F_L(K_{t-1}, L_t)$ and before-tax interest rates $r_t = F_K(K_{t-1}, L_t) - \delta$,

(v) the government's budget constraint is satisfied

$$G + (1 + r_t) B_{t-1} = B_t + \mathcal{T}_t(\{\bar{r}_s\}, \{\bar{w}_s\}, \{p_s\}) + \tau_t^k \mathcal{A}_{t-1}(\{\bar{r}_s\}, \{\bar{w}_s\}, \{p_s\}),$$

(vi) the asset market, the labor market, and the goods market clear

$$\begin{aligned} \mathcal{A}_t(\{\bar{r}_s\}, \{\bar{w}_s\}, \{p_s\}) &= B_t + K_t, \\ \mathcal{L}_t(\{\bar{r}_s\}, \{\bar{w}_s\}, \{p_s\}) &= L_t, \\ \mathcal{C}_t(\{\bar{r}_s\}, \{\bar{w}_s\}, \{p_s\}) + K_t + G &= F(K_{t-1}, L_t) + (1 - \delta) K_{t-1}. \end{aligned}$$

□

The model with multiple safe assets is once again calibrated to the US economy. The capital share α is chosen in order to generate a capital-to-GDP ratio of 2.5, in line with US data. The discount factor β is then chosen to match a real interest rate of 2%, given the capital-to-GDP ratio of 2.5 and a debt-to-GDP ratio of 0.615. This is based on the average US federal debt in the data from 1995 to 2007.[24] The parameters of the income process and the borrowing limit are unchanged relative to Section 2. The depreciation rate for capital is 2% at the quarterly frequency. The value for the tax on capital income is taken from Trabandt and Uhlig (2011), whose estimation for the US in 2007 yields 36%. Government spending and the progressivity of the tax system are unchanged. Finally, the level of labor income taxes τ adjusts in order to ensure that the government budget constraint holds. Table 2 summarizes the parameter values that come out from this procedure.

Table 2: Parameter values in model with multiple safe assets

Parameter	*Description*	*Value*	*Parameter*	*Description*	*Value*
β	discounting	0.995	G/Y	spending-to-GDP	0.088
ρ	persistence of AR (1)	0.966	K/Y	capital-to-GDP	2.5
			B/Y	debt-to-GDP	0.615
σ_ϵ	variance of AR(1)	0.033	p	progressivity of taxes	0.181
EIS	curvature in u	1			
$Frisch$	curvature in v	1/2	τ	level of taxes	0.620
α	capital share	0.25	τ_k	capital income tax	0.36
δ	depreciation rate	0.02	ϕ	borrowing limit	0

C.2 RAMSEY PROBLEM IN THE MODEL WITH CAPITAL AND τ_k

In the model with capital, the dynamic Ramsey problem can be written as

[24] This is based on Dyrda and Pedroni (2022) and is close the values in LeGrand and Ragot (2023) and Aiyagari and McGrattan (1998).

$$\max_{\{\bar{r}_t,\bar{w}_t,r_t,p_t,B_t,K_t,L_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}\right) \quad \text{s.t} \begin{cases} \mathcal{A}_t\left(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}\right) = B_t + K_t, \\ \mathcal{C}_t\left(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}\right) + K_t + G = F(K_{t-1},L_t) + (1-\delta)K_{t-1}, \\ w_t = F_L(K_{t-1},L_t),\ r_t = F_K(K_{t-1},L_t) - \delta, \\ \mathcal{L}_t(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}) = L_t \end{cases}$$

Once again, $\bar{r}_t = (1-\tau_{kt})r_t$ and $\bar{w}_t = \tau_t w_t^{1-p_t}$ denote the after-tax interest rate and the after-tax wage. Notice that we drop the government budget constraint because of Walras' Law. As in the main text, $\mathcal{U}_t$, $\mathcal{A}_t$, $\mathcal{C}_t$, and $\mathcal{L}_t$ are *sequence-space* functions that map sequences of after-tax interest rates, after-tax wages and progressivity into aggregates at time t. Since these functions aggregate optimal individual behavior using the distribution of agents, the Ramsey plan is implementable, provided that the constraints listed above are satisfied.

C.3 PROVING THE OPTIMALITY OF THE MODIFIED GOLDEN RULE

Given choices of $\{\bar{r}_t, \bar{w}_t, p_t, K_t\}$, asset market clearing condition pins down $\{B_t\}$. Similarly, given these choices, the labor market clearing condition pins down the sequence of labor demand $\{L_t\}$. The firm's optimality conditions then pin down the sequence of pre-tax wages $\{w_t\}$ and pre-tax interest rates $\{r_t\}$. Thus, after dropping all the redundant constraints and the associated constraints, the RSS problem reduces to

$$\max_{\{\bar{r}_t,\bar{w}_t,p_t,K_t,\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}\right) \quad \text{s.t} \quad \mathcal{C}_t + K_t + G = F(K_{t-1},\mathcal{L}_t) + (1-\delta)K_{t-1}$$

The Lagrangian for this problem is

$$\max_{\{\bar{r}_t,\bar{w}_t,p_t,K_t,\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left\{ \mathcal{U}_t\left(\{\bar{r}_s\},\{\bar{w}_s\},\{p_s\}\right) + \lambda_t^{GM}\left\{F(K_{t-1},\mathcal{L}_t) + (1-\delta)K_{t-1} - \mathcal{C}_t - K_t - G\right\}\right\}$$

The first-order condition with respect to capital is

$$\lambda_t^{GM} = \beta\lambda_{t+1}^{GM}\left[F_K\left(K_t,\mathcal{L}_{t+1}\left(\cdot\right)\right) + (1-\delta)\right], \quad \text{for } t = 0,1,\ldots. \tag{41}$$

Using homogeneity of degree one of the production function and defining the capital labor ratio $k_t \equiv \frac{K_t}{\mathcal{L}_{t+1}(\cdot)}$, we can rewrite this first order condition as

$$\lambda_t^{GM} = \beta\lambda_{t+1}^{GM}\left[F_K\left(k_t,1\right) + (1-\delta)\right], \quad \forall t.$$

If quantities and multipliers converge to an interior steady state, this condition becomes

$$1 = \beta \left[F_K(k, 1) + (1 - \delta) \right]$$

This establishes the optimality of the modified golden rule. This result was first proved by Aiyagari (1995) under the assumption that government spending is endogenous. Here, I verify that it goes through in the model with exogenous government spending and regardless of the planner's preference for redistribution. As discussed by Straub and Werning (2020), two situations would prevent the applicability of this result: (i) nonconvergence to an interior steady state and (ii) nonconvergence of the multipliers. These are ruled out by assumption.

C.4 Additional first-order conditions

In addition to the first-order condition with respect to capital (41), we also have:

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial \bar{r}_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial \bar{r}_u} - \frac{\partial \mathcal{C}_t}{\partial \bar{r}_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{42}$$

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial \bar{w}_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial \bar{w}_u} - \frac{\partial \mathcal{C}_t}{\partial \bar{w}_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{43}$$

$$\sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial p_u} + \sum_{t=0}^{\infty} \beta^{t-u} \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial p_u} - \frac{\partial \mathcal{C}_t}{\partial p_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{44}$$

If quantities and multipliers converge to an interior steady state, the same logic as in Appendix A.3 verifies that when $u \to \infty$, these conditions become

$$\mathcal{E}_{\bar{r}}^{\mathcal{U}} + \lambda^{GM} \left(w^{MGR} \mathcal{E}_{\bar{r}}^{\mathcal{L}} - \mathcal{E}_{\bar{r}}^{\mathcal{C}} \right) = 0, \tag{45}$$

$$\mathcal{E}_{\bar{w}}^{\mathcal{U}} + \lambda^{GM} \left(w^{MGR} \mathcal{E}_{\bar{w}}^{\mathcal{L}} - \mathcal{E}_{\bar{w}}^{\mathcal{C}} \right) = 0, \tag{46}$$

$$\mathcal{E}_{p}^{\mathcal{U}} + \lambda^{GM} \left(w^{MGR} \mathcal{E}_{p}^{\mathcal{L}} - \mathcal{E}_{p}^{\mathcal{C}} \right) = 0, \tag{47}$$

where $w^{MGR} = F_L(k^{MGR}, 1)$ is the pre-tax wage implied by the modified golden rule. One can use the three conditions above, after imposing the MGR, together with the goods market clearing condition to solve for the unknowns $\{\bar{r}, \bar{w}, p, \lambda^{GM}\}$. Appendix C.5 details the computational procedure.

C.5 COMPUTING THE RSS IN THE MODEL WITH CAPITAL

To operationalize (45)-(47) one needs to compute the discounted sum of the asymptotic response of aggregate outcomes $\mathcal{Y} \in \{\mathcal{U}, \mathcal{L}, \mathcal{C}\}$ to changes in the instruments of the planner:

$$\lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{Y}_t}{\partial \bar{r}_u}, \quad \lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{Y}_t}{\partial \bar{w}_u}, \quad \lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{Y}_t}{\partial p_u}.$$

This could be done by computing the Jacobians for welfare, labour supply, and consumption, around the steady state implied by some candidate $(\bar{r}, \bar{w}, p)$ and then taking the discounted sum of some far out column. Even if one uses the methods in Auclert et al. (2021), this turns out to be somewhat costly since we need a pretty large horizon to get convergence of the discounted sum. Therefore, I follow the same approach as in Appendix A.2 to compute the discounted long-run responses.

With these in hand, I solve for the unknowns $\{\bar{r}, \bar{w}, p, \lambda^{GM}\}$ as follows. First, I reduce the dimensionality of the system by using (46) to substitute out λ^{GM}. These leaves (45) and (47), together with goods-market clearing, to solve for $\{\bar{r}, \bar{w}, p\}$. Using homogeneity of aggregate consumption and aggregate labor supply, I can solve for $\bar{w}$ from the goods market clearing condition and, reducing the system to two equations. I then follow the same two-step procedure as in Appendix A.2 to solve for $\{\bar{r}, p\}$.

D APPENDIX: COMPUTING THE OPTIMAL TRANSITION

This appendix describes the algorithm to compute the optimal transition path. The approach combines the sequence-space methods used to characterize the Ramsey steady state with policy gradient techniques adapted from machine learning.

D.1 OVERVIEW

The Ramsey problem is stated in Section 4. As shown in Appendix A.3, by Walras' Law this problem can be reformulated with goods market clearing as the only constraint, after dropping the government budget constraint and the associated choice variable B_t. The transition algorithm works with this equivalent formulation. The key computational

challenge is the high dimensionality of the problem, as the planner must optimize over policy sequences of length T.

To solve this, I combine insights from two existing approaches. Auclert et al. (2023a) introduce sequence-space Jacobians to characterize the Ramsey steady state, but do not provide a general method for global optimization over transition paths. Dyrda and Pedroni (2022) optimize over parameterized paths in the time domain, but use computationally expensive brute-force methods.

I integrate these methods by parameterizing policy paths in the time domain, following Dyrda and Pedroni (2022), and using the sequence-space structure to compute analytical gradients of an augmented Lagrangian. The key extension relative to Auclert et al. (2021) is computing the Jacobians around arbitrary transition paths rather than a steady state.

D.2 Augmented Lagrangian formulation

Let $x = \{r_t, p_t, \tau_t\}_{t=0}^{T}$ denote the vector of all policy variables stacked over time, and let $\boldsymbol{g}(x) = \{\mathcal{Y}_t - \mathcal{C}_t - G\}_{t=0}^{T}$ denote the goods-market residuals. The augmented Lagrangian is:

$$\mathcal{L}_{\rho_{AL}}(x, \lambda) = -W(x) + \lambda^\top \boldsymbol{g}(x) + \frac{\rho_{AL}}{2} \|\boldsymbol{g}(x)\|^2, \tag{48}$$

where $W(x) = \sum_{t=0}^{T} \beta^t \mathcal{U}_t$ is aggregate welfare, $\lambda \in \mathbb{R}^T$ is the vector of Lagrange multipliers, and $\rho_{AL} > 0$ is a penalty parameter.

The gradient of the augmented Lagrangian with respect to the policy path is:

$$\nabla_x \mathcal{L}_{\rho_{AL}} = -\nabla_x W + (\nabla_x \boldsymbol{g})^\top (\lambda + \rho_{AL} \boldsymbol{g}), \tag{49}$$

which can be expressed in terms of sequence-space Jacobians as:

$$\nabla_x \mathcal{L}_{\rho_{AL}} = -\boldsymbol{J}_U^\top \boldsymbol{\beta} + (\boldsymbol{J}_Y - \boldsymbol{J}_C)^\top (\lambda + \rho_{AL} \boldsymbol{g}), \tag{50}$$

where $\boldsymbol{J}_U$, $\boldsymbol{J}_Y$, and $\boldsymbol{J}_C$ are the $T \times 3T$ Jacobians of aggregate utility, output, and consumption with respect to the policy path x (i.e., the entries of $\boldsymbol{J}_U$ are $\frac{\partial \mathcal{U}_t}{\partial x_s}$), and $\boldsymbol{\beta} = (1, \beta, \ldots, \beta^T)^\top$ is the vector of discount factors.

D.3 POLICY PARAMETERIZATION

To reduce dimensionality, I parameterize policy paths as:

$$x_t^j(\theta) = \bar{x}^j + e_t \sum_{k=1}^{K_j} \theta_k^j \phi_k(t), \tag{51}$$

where $\bar{x}^j$ is the terminal Ramsey steady-state value for instrument $j \in \{r, p, \tau\}$, $\{\phi_k\}_{k=1}^K$ are Chebyshev basis functions, and $e_t = (1 - t/T)^\kappa$ is an envelope that forces convergence to the terminal state. In practice, I use $K_j \approx 3$–5 basis functions per instrument. The baseline results in Section 4 use $\kappa = 10$. Section D.6 below examines sensitivity to this choice.

The gradient with respect to parameters follows from the chain rule:

$$\nabla_\theta \mathcal{L}_{\rho_{AL}} = \left(\frac{\partial x}{\partial \theta}\right)^\top \nabla_x \mathcal{L}_{\rho_{AL}} = \Phi^\top (e \odot \nabla_x \mathcal{L}_{\rho_{AL}}), \tag{52}$$

where Φ is the $T \times K$ basis matrix and e is the envelope vector.

To ensure policy paths remain within sensible bounds (e.g., $r_t > r_{\min}$), I apply a smooth transformation:

$$x_t^j = x_{\min}^j + \frac{x_{\max}^j - x_{\min}^j}{2}\left(1 + \tanh(z_t^j)\right), \tag{53}$$

where z_t^j is the unconstrained path. The gradient is adjusted accordingly using the derivative of the tanh transformation.

D.4 EFFICIENT JACOBIAN COMPUTATION

While Auclert et al. (2021) show how to compute Jacobians efficiently around a steady state using the "fake news" algorithm, expanding this to arbitrary transition paths requires care. The standard algorithm relies on the fact that, at the steady state, expectations of future marginal values are time-invariant. Along a transition path, these objects change over time.

I generalize the approach by noting that the recursive structure of the household problem still applies: a shock at time s does not affect optimal decisions at time $t > s$. Thus, to compute the s-th column of the Jacobian, the backward iteration need only start from the unperturbed continuation value at $t = s + 1$.

Parallelization. A key ingredient for feasibility is the parallelization of the Jacobian construction. Since the effect of a shock at time s on outcomes is independent of shocks at other times (for the purpose of constructing the Jacobian column), I compute each column $s \in \{0, \ldots, T\}$ in parallel. This leads to near-linear scaling with the number of cores and reduces the total computation time by a factor of 5 on standard workstations.

D.5 Optimization procedure

The outer loop uses the standard augmented Lagrangian method:

1. Initialize $\theta = 0$, $\lambda = 0$, $\rho_{AL} = 1$
2. **Inner loop**: Minimize $\mathcal{L}_{\rho_{AL}}(\theta, \lambda)$ over θ using L-BFGS with the analytical gradient (50). Continue until stationarity tolerance is met or a maximum number of iterations is reached.
3. **Multiplier update**: $\lambda \leftarrow \lambda + \rho_{AL} \cdot g(x^*)$
4. **Penalty update**: If constraint violation $\|g\|_\infty$ has not decreased sufficiently, increase the penalty: $\rho_{AL} \leftarrow 5\rho_{AL}$
5. Repeat steps 2–4 until $\|g\|_\infty < \varepsilon_g$ and $\|\nabla_\theta \mathcal{L}_{\rho_{AL}}\|_\infty < \varepsilon_\theta$

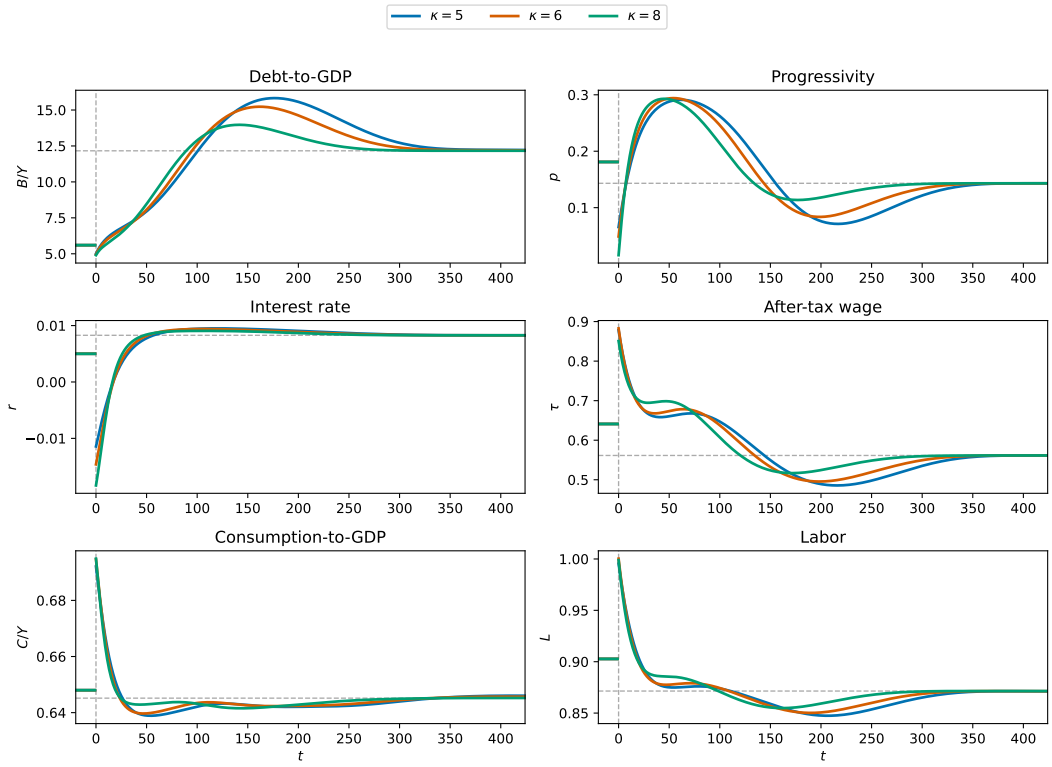
D.6 Sensitivity analysis

A key hyperparameter is κ, the curvature of the envelope function $e_t = (1 - t/T)^\kappa$ that anchors the policy path to the steady state. The baseline results in Section 4 use $\kappa = 10$.

Figure 19 plots the optimal transition paths under $\kappa \in \{5, 6, 8\}$. The qualitative patterns are the same across specifications: the planner flattens the tax schedule on impact, interest rates fall, and debt gradually rises towards the long-run target. There are quantitative differences but the overall shape of the transition is stable.

Figure 20 reports welfare gains across κ, measured as a permanent consumption-equivalent variation (CEV) relative to the status quo. Welfare is increasing in κ but the gains flatten out for $\kappa \geq 8$, suggesting that the baseline choice of $\kappa = 10$ is in the region where the approximation is not binding.

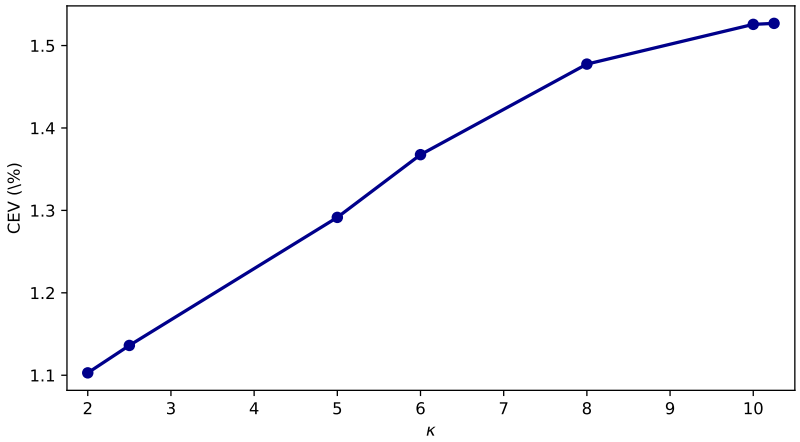
Figure 19: Sensitivity of optimal transition to envelope parameter κ



Note: Policy paths for $\{r_t, p_t, \tau_t\}$ are parameterized over $T_{\text{free}} = 425$ quarters using a Chebyshev time basis (degree 4 per instrument) and an envelope $e_t = (1 - t/T_{\text{free}})^{\kappa}$. Each line corresponds to a converged run that varies only κ (all other settings held fixed). The transition appends a constant tail of length $T_{\text{tail}} = 375$ at terminal policies. Pre-reform periods at the initial steady state are shown to the left of $t = 0$; dashed lines indicate Ramsey steady-state values.

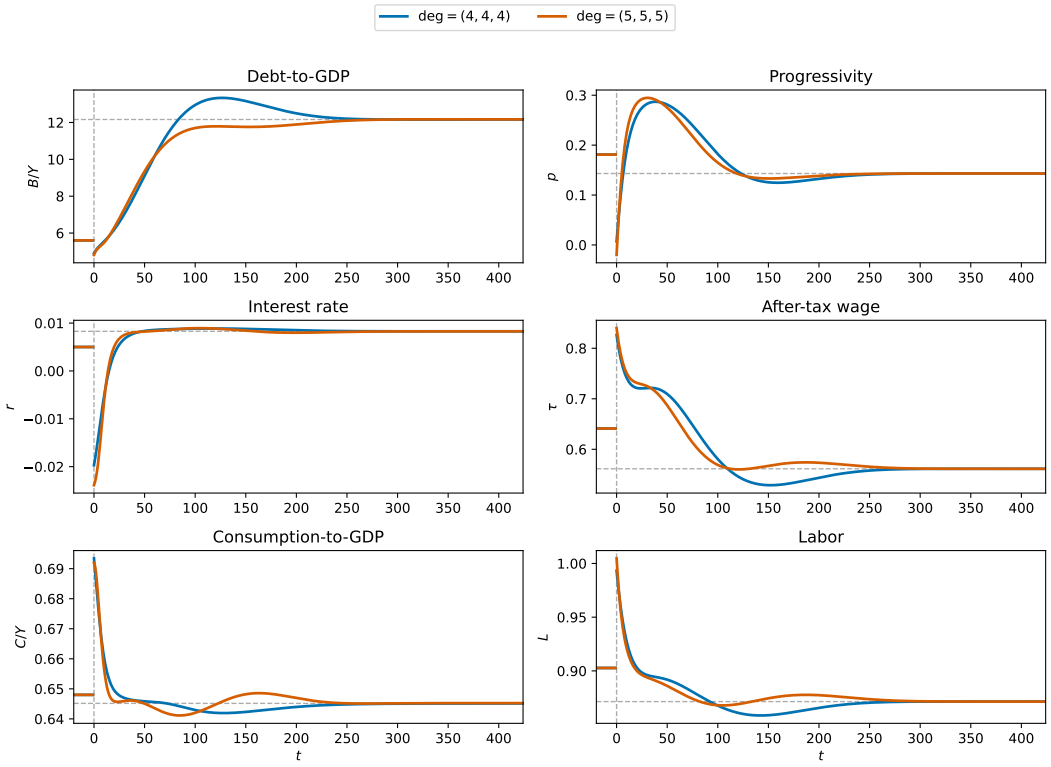
Sensitivity to basis degree. Figure 21 compares the optimal transition paths under degree-4 and degree-5 Chebyshev parameterizations, holding $\kappa = 10$ fixed. The paths are close, suggesting that the baseline results are not sensitive to increases in the dimensionality of the approximation.

Figure 20: Welfare sensitivity to envelope parameter κ (CEV)



Note: Welfare is reported as a permanent consumption-equivalent variation (CEV) relative to the status quo policy that holds instruments fixed at their initial steady-state values. For each κ, welfare is computed from the converged transition with $T_{\text{free}} = 425$ quarters and a constant tail of length $T_{\text{tail}} = 375$ at terminal policies.

Figure 21: Sensitivity of optimal transition to Chebyshev degree



Note: Both runs use $\kappa = 10$ and $T_{\text{free}} = 425$ with a tail of $T_{\text{tail}} = 375$. The degree-5 run is warm-started from the degree-4 solution. Dashed lines indicate Ramsey steady-state values.