Optimal Public Debt with Redistribution [*]

Matias Bayas-Erazo

November 21, 2023

Click here for most recent version

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 these two elements? In this paper, I consider a benevolent optimizing government and explore how both progressivity and indebtedness depend on the planner's preferences for redistribution. I compute the optimal mix of debt and progressivity in a standard heterogeneous-agent incomplete-markets economy. Somewhat surprisingly, I show 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. This is mainly due to a novel interest rate channel: redistributive taxation reduces the need to self-insure and thus makes government borrowing more expensive. I then back out implied preferences for redistribution in advanced economies and argue that they are inconsistent with both Utilitarian and Rawlsian criteria.

[*] 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, Joao Guerreiro, Fergal Hanks, Christian Hellwig, Diego Huerta, and Kwok Yan Chiu for comments that helped improve the paper. All remaining errors are my own.

1 Introduction

When markets are incomplete, public debt can provide liquidity and help agents self-insure against idiosyncratic income risk (Woodford, 1990; Aiyagari and McGrattan, 1998). At the same time, 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. What is the optimal mix of debt and progressivity? How does it depend on social preferences for redistribution?

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.

This paper aims to fill this gap and studies how the optimal mix of debt and progressivity depends on social preferences for redistribution. Specifically, I compute the optimal mix of debt and progressivity in a standard heterogeneous-agent incomplete-markets model (Bewley, 1977; Huggett, 1993; Aiyagari, 1994). Households face uninsurable idiosyncratic income risk and self-insure by holding safe assets, subject to borrowing constraints. The government controls the supply of safe assets and a progressive tax on labor income.

By reducing the variance of after-tax income, progressive taxes act as a form of social insurance against idiosyncratic income risk. However, they also distort the economy. Public debt also plays a key role in these economies by relaxing the constraints that agents face (Woodford, 1990). Specifically, by increasing liquidity, public debt can help agents self-insure, possibly reducing the need for a distortionary tax system. Understanding the interaction between public debt and progressive taxation is thus key to the design of optimal fiscal policy.

I explore two different concepts of long-run optimality. I start with the optimal steady state problem, computing the mix of debt and progressivity that maximizes welfare ignoring transitions (OSS problem). Next, I study fully dynamic optimal policy and solve for the limiting steady state of the Ramsey problem for this economy (RSS problem). In both cases, I find that planners with stronger preferences for redistribution favor more progressive tax systems, as one would expect. Somewhat surprisingly, they also favor *lower* levels of public debt.

To explain what is behind this result, I study the effects of progressive tax reforms. I show that increasing the progressivity of the tax system puts upward pressure on interest rates. This *interest rate channel of progressivity* arises because 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.[1] In addition, 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 also goes down. So both on the cost side and on the benefits side, more redistribution reduces the incentive to issue public debt. As a result, Rawlsian planners, who naturally favor redistributive tax systems, find it optimal to issue lower levels of public debt. Conversely, planners that prioritize efficiency over equity, what I call Efficiency planners, find it optimal to issue high levels of public debt and finance this with a regressive tax system. This is due to the fact that they are not concerned with the adverse redistribution effects of high debt issuance. Instead, they focus on the liquidity benefits of public debt and find debt to be a more effective way of providing social insurance.

This property of the optimal mix holds regardless of whether I compute the optimum by maximizing welfare in steady state or whether I do so by taking into account transitional dynamics. To show this, I solve for the Ramsey steady state, the limiting steady state of the fully dynamic optimal policy problem that takes into account transitions and lets both instruments vary over time. My approach to this problem builds on the sequence-space approach introduced by Auclert et al. (2023), extending it to allow for departures from utilitarian welfare criteria.

To the best of my knowledge, the existence of the Ramsey steady state in the standard heterogeneous-agent model with separable preferences remains an open question. Angeletos et al. (2022) study the Ramsey problem in a special class of incomplete-markets economies that feature a similar role for public debt, while abstracting from heterogeneity and wealth dynamics. They show that long-run satiation can be optimal, depending on the primitives of the economy. In line with this result, numerical investigations in Auclert et al. (2023) suggest that a version of the Friedman rule applies when the planner uses a utilitarian welfare criterion. Specifically, the planner finds it optimal to satiate the demand for public debt, issuing debt to the point where the interest rate equals the discount rate. However, this is not consistent with an interior steady state in this class of models: aggregate asset demand explodes as the interest rate approaches the discount rate from below.

To get around this, I depart from standard welfare criteria and consider a planner that uses generalized dynamic stochastic weights to conduct welfare assesments (Davila and Schaab, 2022; LeGrand and Ragot, 2023). These weights are allowed to depend on endogenous outcomes and are used to aggregate individual instantaneous utilities. They represent the value that the planner assigns to a marginal unit of consumption by a particular individual in a given period and untie society's concerns for fairness from individual lifetitime utilities. To the extent that a model with infinitely lived agents is meant to capture altruistically linked generations, this type of social planner can also be interpreted as one that distinguishes between the welfare of each generation (Phelan, 2006; Farhi and Werning, 2007). In particular, it allows for the possibility that the planner cares about inequality of future generations *directly*. Following this interpretation, I

[1] This force reminds of the one in Angeletos et al. (2022), where easing the underlying friction likewise leads to an increase in the cost of government borrowing.

refer to this type of planners as generational planners. Although non-standard, this class of social welfare functions nests the utilitarian criterion and allows me to speak about inequality-aversion while keeping the problem tractable.

I am able to find an interior Ramsey steady state for generational planners that are averse to wealth inequality. In other words, a Ramsey planner that puts relatively high weight on the instantaneous utility of the asset-poor no longer finds satiation optimal. Moreover, the optimal level of public debt decreases with the parameter that governs the planner's aversion to wealth inequality, whereas the optimal progressivity of the tax system increases. When I compare the optimal mix across both solution concepts quantitatively, I find that the OSS problem overestimates the costs of debt issuance and overestimates the benefits of progressive tax systems.

The substitutability/complementarity between public debt and progressive taxation depends on the optimality criterion that is used. For steady state planners, the two instruments can be substitutes or complements depending on the level of public debt. When the debt is relatively low, they are substitutes in the sense that the optimal progressivity of the tax system falls in response to exogenous increases in the level of public debt. If instead the debt happens to be high, optimal taxes become more progressive as the level of public debt increases. I argue that this is because the insurance/liquidity effect of public debt dominates for low values of debt and thus reduces the need to insure via distortionary labor income taxation. The adverse redistribution effects of increasing the public debt, triggered by the rise in interest rates that favors the rich, start to dominate as the level of debt continues to grow. To address the resulting inequality in consumption, planners that value equity find it optimal to rely more heavily on progressive income taxes.

The planners that takes into account transitions always see the two instruments as substitutes, regardless of the level of public debt and the taste for redistribution. The adverse redistribution effects only happen in the long run, which the Ramsey planner discounts appropriately. This is related to the fact that the OSS solution concept overestimates the costs of debt issuance (Angeletos et al., 2022).

To explore the robustness of these results, I consider a number of extensions. First, I analyze a version of the model with multiple safe assets. This version 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.

Second, 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). Among other things, 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 features increasing average tax rates but decreasing marginal tax rates.

Third, I consider the possibility of taxing savings. A linear tax makes no difference in the baseline model where debt is the only safe asset. In the model with multiple safe assets, it brings the results closer to the model where debt is the only safe asset. This is because taxing capital allows the government to effectively control the supply of the alternative assets. Specifically, I show that it is optimal to use the tax on savings to implement the golden rule in the OSS and the modified golden rule in the RSS.

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 throughout this paper. The results suggest that implied preferences for redistribution in the US and other advanced economies are inconsistent with standard Utilitarian and Rawlsian criteria– the covariance between welfare weights and capital/labor income is *positive*. At the same time, allowing for variation in risk across countries can explain the observed long-run mix of debt and progressivity in advanced economies.

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 steady-state welfare but ignoring transitional dynamics. Like most of the literature, they restrict attention to linear taxes and thus ignore the interaction with redistributive 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. However, due to the computational challenges associated with the fully dynamic optimal policy problem, he continues to ignore transitions. Angeletos et al. (2022) solve the full Ramsey problem in a stylized incomplete-markets economy, bypassing the computational challenges of the original problem. 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. (2023) 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. 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. 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 approach of Auclert et al. (2023) to allow for departures from utilitarian welfare criteria and show that an interior RSS exists whenever the planner has some aversion to inequality.

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).[2] 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 $\phi > 0$.

[2] 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.

Individual productivity θ evolves according to 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

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

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):

$$T_t(y) = y - \tau_t\, y^{1-p_t}, \tag{2}$$

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 5.2, 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 (2) 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

$$G + r_{t-1}B_{t-1} = B_t - B_{t-1} + \int T_t(\boldsymbol{y}_t(x))dD_t(x), \quad \text{for all } t. \tag{3}$$

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

$$\int \boldsymbol{y}_t(x)dD_t(x) = \int \boldsymbol{c}_t(x)dD_t(x) + G, \quad \text{for all } t. \tag{4}$$

Finally, the asset market clearing condition is

$$\int \boldsymbol{a}_t(x)dD_t(x) = B_t, \quad \text{for all } t. \tag{5}$$

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 5.1, when I consider a more general production technology and allow firms to issue claims to physical capital.

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%.[3] 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.[4] 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
σ_ϵ	variance of AR(1)	0.033	p	progressivity of taxes	0.181
EIS	curvature in u	1	τ	level of taxes	0.641
Frisch	curvature in v	$1/2$	ϕ	borrowing limit	0

3 OPTIMAL STEADY STATE

This section contains the main results for the optimal steady state. After stating the problem, I discuss how progressivity affects the real interest rate and social welfare. I then present the results for the optimal mix and study how it depends on the planner's taste for redistribution.

[3] 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 $A/Y = 1.4$.

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

It is important to note the different optimality concepts that the literature has used when computing optimal policy in dynamic models with heterogeneous agents. Due to the difficulties involved in tracking the wealth distribution, the methods for solving the dynamic optimal policy problem originally proposed by Aiyagari (1995) have been developed fairly recently (Acikgoz et al., 2023; Auclert et al., 2023; 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 first step both because it is computationally tractable and helps develop intuition. In addition, it allows me to sidestep some issues that arise in the problem that takes into account transitions.

Given this, I start by formulating an optimal steady state problem for a planner that controls both the progressivity of the tax schedule and the level of public debt. Section 4 presents the unrestricted Ramsey problem that takes into account transitions and lets both instruments vary over time.

3.1 Optimal steady state problem (OSS)

Given some exogenous spending needs G, 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 steady state welfare, subject to the government's budget constraint and equilibrium behavior. To formalize this, let $\mathcal{W}$ denote aggregate welfare in the stationary equilibrium, given an interest rate and fiscal policy. Throughout most of the paper, I will focus on a social welfare function (SWF) that departs from utilitarian welfare criteria by allowing a flexible form of aversion to inequality. In particular, I assume that the planner aggregates individual utilities using weights that may depend on an agent's current productivity and asset holdings. Formally,

$$\mathcal{W}(r, \tau, p) = \sum_{t=0}^{\infty} \beta^t \int_i \omega(\theta_i, a_i)\{u(c_i) - v(\ell_i)\}\, di, \tag{6}$$

where the weights $\omega(\theta, a) \propto \exp(-\alpha_\theta\, \theta - \alpha_a\, a)$, normalized such that $\int_i \omega(\theta_i, a_i)\, di = 1$.[5] 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. The 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. If $\alpha_a < 0$ or $\alpha_\theta < 0$, the planner instead prioritizes the welfare of individuals with higher asset or income levels, respectively. The absence of time subscripts inside the integral reflects the fact that all steady state planners behave as if the economy jumps immediately to the stationary equilibrium implied by the candidate fiscal policy. Therefore, from the perspective of the planner, policy functions and distributions are constant over time.

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

[5] This normalization ensures that the SWF does not penalize increases in asset holdings *per se*.

on weights that are endogenous. I offer two alternative microfoundations. The most literal one is that the planner uses generalized dynamic stochastic weights to conduct welfare assessments, in the sense of Davila and Schaab (2022). As that paper shows, the introduction of dynamic stochastic weights allows one to formalize new welfare criteria that society may find appealing. This unties 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 the social planner considers it fair to redistribute across these dimensions. An alternative microfoundation is that the social planner distinguishes between the welfare of each generation (Phelan, 2006; Farhi and Werning, 2007). To the extent that a model with infinitely lived agents is meant to capture altruistically linked generations of finitely lived agents, this SWF captures the preferences of a planner that may care about inequality of future generations *directly*.[6] More concretely, 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*.

To alleviate concerns, I also examine more standard welfare criteria. The purpose of this is to convince the reader that the 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 of the initial generation. Formally,

$$\mathcal{W}^{GU}(r, \tau, p) = \int \omega(\theta_i, a_i) V(\theta_i, a_i)\, di, \tag{7}$$

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 (6) may reflect a stronger notion of aversion to inequality than the ones behind equation (7). Unlike generational planners, generalized utilitarian planners are not paternalistic: welfare assesments are based on individual lifetime utilities.

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[7], social welfare for a "Bénabou planner" is given by

$$\mathcal{W}^{\alpha}(r, \tau, p) = \left(\int \bar{c}_i^{1-\frac{1}{\alpha}}\, di \right)^{\frac{1}{1-\frac{1}{\alpha}}}. \tag{8}$$

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

[6] 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 (6).

[7] 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 $V(\theta_i, a_i)$.

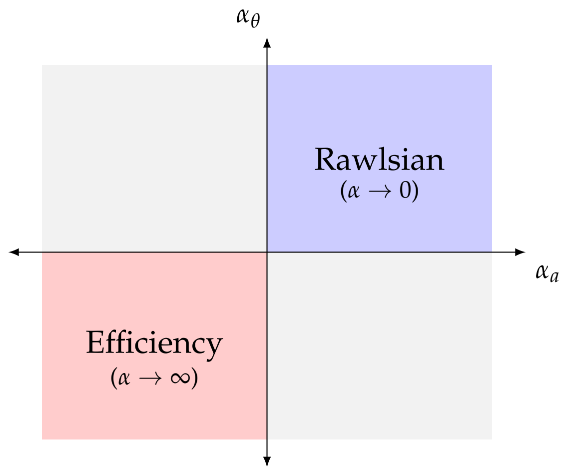

Figure 1: Preferences for redistribution

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. Like generalized utilitarian planners, Bénabou planners are not paternalistic.

Figure 1 illustrates the connection between the different welfare functions. As expected, one can construct weights for both the generational and generalized utilitarian planners in order to approximate the behavior of Bénabou planners. In Appendix A, I show that a Bénabou planner that is Rawlsian ($\alpha \to 0$) behaves like a planner that evaluates social welfare according to (6) or (7) with weights that decrease along the asset dimension ($\alpha_a > 0$). Similarly, when $\alpha \to \infty$, the Bénabou planner is like a generational or generalized utilitarian planner with weights that increase along the asset dimension ($\alpha_a < 0$). Relative to (8), the specifications in (6) and (7) offer more flexibility by allowing the planner to separate aversion to inequality into the asset and productivity dimensions. Having said this, aversion to inequality along the productivity dimension does not seem to play a prominent role. Thus, in what follows, I will set $\alpha_\theta = 0$ and focus on three benchmark cases: $\alpha_a > 0$, $\alpha_a = 0$, and $\alpha_a < 0$. Whenever $\alpha_a > 0$, I will say that the planner is Rawlsian and has a preference for redistribution. If $\alpha_a < 0$, I will call this an Efficiency planner and say that the planner has no preference for redistribution.

With these definitions in hand, let $\mathcal{A}(r, \tau, p)$ and $\mathcal{T}(r, \tau, p)$ denote aggregate asset holdings and aggregate tax revenues, respectively. Given a social welfare function, the OSS problem be written as

$$\max_{\{r,B,\tau,p\}} \mathcal{W}(r, \tau, p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(r, \tau, p) = B, \\ G + rB = \mathcal{T}(r, \tau, p) \end{cases} . \tag{9}$$

Angeletos et al. (2022) clarify how this concept of long run optimality overestimates the costs of debt issuance and thus underestimates the optimal long run level of public debt. However, it allows me to illustrate the forces at play in a transparent manner and to explore the robustness of the results to my choice of SWF. To understand what is being driven by the fact that this concept of long-run optimality ignores welfare effects along the transition, Section 4 studies the steady state of the dynamic Ramsey problem.

3.2 First-order effects of progressivity

Before solving the problem, I carry out a comparative static exercise that will isolate the mechanism behind the main result in this paper. With this in mind, consider a small permanent change in the progressivity of the tax schedule dp around the baseline economy. The goal is to understand how this reform affects equilibrium outcomes and illustrate the forces at play. Due to the nature of the OSS problem, I focus on steady state outcomes.

Lemma 1 characterizes the response of individual value functions and the real interest rate to this change in the tax system. The first part shows that varying progressivity alters individual life-time utilities through partial equilibrium (direct) and general equilibrium (indirect) channels.

Lemma 1 *The response of individual outcomes dV to a small permanent change in progressivity dp solves*

$$dV(x) = u'(\boldsymbol{c}(x))\left(\boldsymbol{y}(x)^{1-p}d\boldsymbol{\tau} + a\, d\boldsymbol{r} - \mathbf{z}(x)\log \boldsymbol{y}(x)\right) + \beta \mathbb{E}_{\theta'|\theta}\left[dV(\boldsymbol{a}(x), \theta')\right]. \tag{10}$$

The general equilibrium responses $d\boldsymbol{r}$ and $d\boldsymbol{\tau}$ are pinned down by the following system:

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

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

To see this more clearly, iterate (10) forward to write the response 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 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 2a 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 2b 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 argue below, is positive.

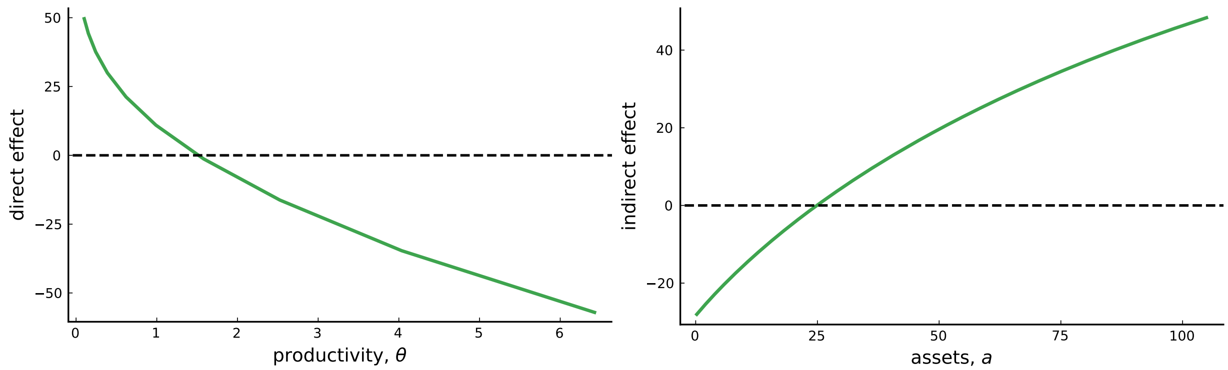

(a) Direct effect along the productivity dimension

(b) Indirect effect along the asset dimension

Figure 2: Direct and indirect effects across the state space

INTEREST RATE CHANNEL OF PROGRESSIVITY

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. 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 $\boldsymbol{dr} > 0$ in general. Figure 3 illustrates this mechanism by plotting the shift in aggregate asset demand following a progressive tax reform. The point at which the solid black line intersects with the grey line gives the equilibrium interest rate prior to the reform. The intersection between the blue dotted line and the grey line gives the interest rate after the reform. The difference between the two dashed lines is the interest rate channel of progressivity.

This is the first paper to isolate this mechanism in a quantitative heterogeneous-agent incomplete-markets model and work out the implications for optimal fiscal policy.[8] 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 reduce the demand for safe assets and lowers the upper bound on deficits in a model with perpetual youth. Taken together, their findings point to a potential conflict between progressive taxation and debt issuance. However, the importance of this mechanism in more realistic models remains unclear.

[8] 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. However, they do not study the implications for optimal fiscal policy.

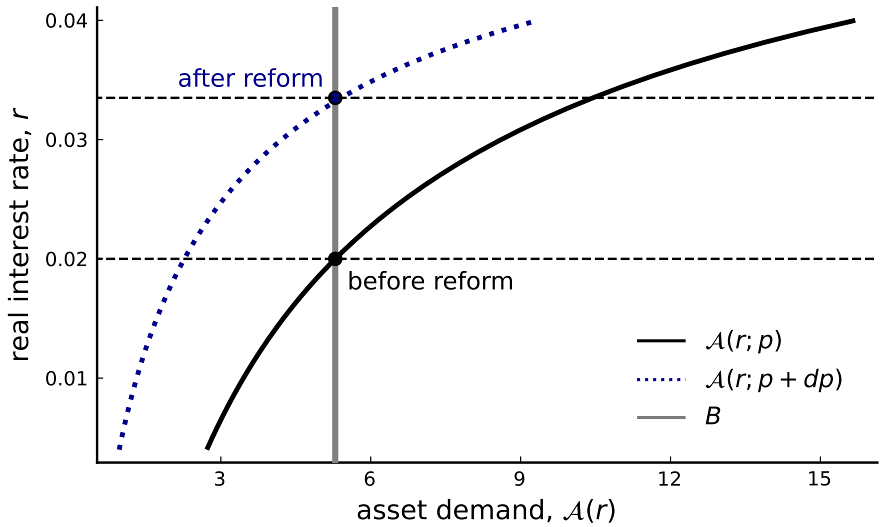

Figure 3: Equilibrium in the asset market before and after progressive reform

The quantitative relevance of the interest rate channel depends on the sensitivity of aggregate safe asset demand to changes in the real interest rate. If aggregate safe asset demand is sufficiently elastic, then a small change in interest rates suffices to clear the market. I'm unaware of empirical estimates for the elasticity of the risk-free rate to changes in the progressivity of the tax system that could be used to discipline the effect. However, there is a vast empirical literature estimating the (semi) elasticity of the risk-free rate to changes in the level of public debt. Mian et al. (2022) survey the empirical evidence and find that most estimates lie in between 1.2% and 2.2%. When I compute this elasticity in the calibrated model, I find that the semi-elasticity of the risk free rate to changes in the level of public debt is 0.6%, well below the lower bound of the empirical estimates. This suggests that the standard heterogeneous agent model features an interest rate elasticity of aggregate asset demand that is *too high* relative to the data.[9] Therefore, the interest rate channel would likely be amplified in models that are more in line with the empirical estimates.

In any case, the mechanism appears quantitatively strong in a standard calibration of this model. As evidence for this, I now show that the indirect effects of progressive tax reforms, driven by the interest rate response, can outweigh the direct effects. Indeed, 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

[9]This is not driven by the absence of capital in the baseline model. See Section 5.1.

question.

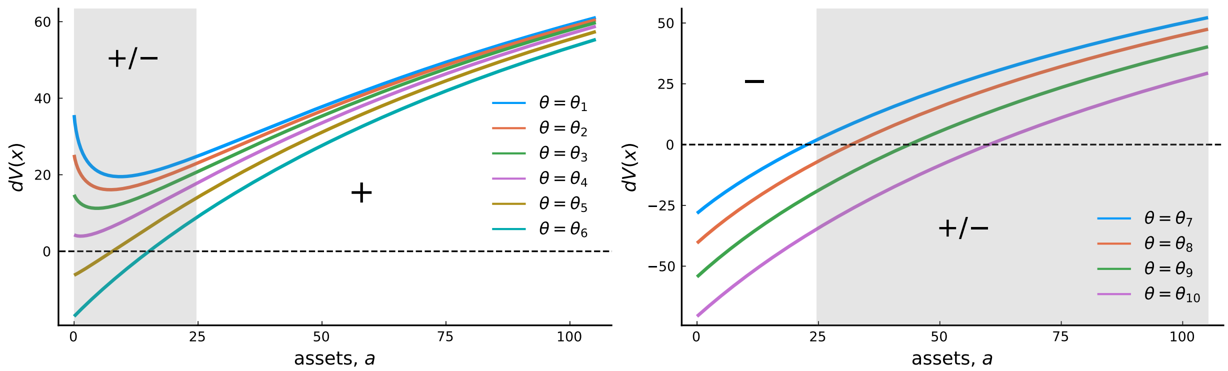

(a) Total effect at the bottom of the θ distribution

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

Figure 4: Individual responses across the state space

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. As expected from the discussion above, 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.

One can then aggregate these effects across the distribution to look at the response of social welfare. In the appendix, I show that one can use Lemma 1 to derive a simple expression for the change in social welfare for Bénabou planners under the assumption that the elasticity of intertemporal substitution is equal to one. For generational planners, it is not possible to derive an analogous expression because the envelope theorem does not apply, and thus the behavioral responses enter the formulas. However, one can still compute the change in social welfare numerically and decompose it into direct and indirect effects. To do so, I use the fact that the general equilibrium effects of the reform– mainly, the interest rate response and the change in the level of taxes– can be computed using the second part of Lemma 1. Figure 5 implements

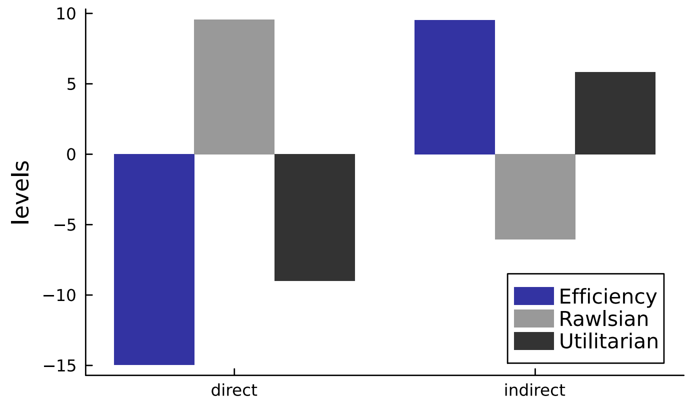

(a) Generational planners

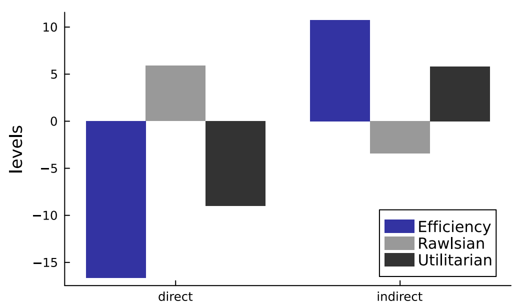

(b) Bénabou planners

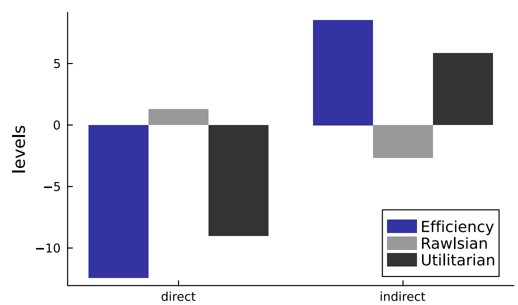

(c) Generalized utilitarian planners

Figure 5: Welfare decomposition of small permanent changes in progressivity

this decomposition around the baseline economy for the three benchmark generational planners: Rawlsian planner ($\alpha_a > 0$), Utilitarian planner ($\alpha_a = 0$), and Efficiency planner ($\alpha_a < 0$). For completeness, I also report the results for the welfarist planners in the bottom panels. The sign of the direct effect and the indirect effect depend on the planner's preference for redistribution. Both the Utilitarian and Efficiency planners see the indirect effect of progressive tax reforms as welfare-improving but dislike the direct effect. Inequality-averse planners benefit from the direct effect but lose from the indirect effect. Intuitively, if the planner uses weights that concentrate at the bottom of the distribution, the benefits of progressive tax reforms are only due to the direct effect because this is the one that favors the poor. It reduces the taxes paid at the bottom while increasing the taxes paid at the top, benefiting the agents that are most valuable to the planner. On the other hand, if the planner uses weights that increase along the asset distribution, the benefits of progressive tax reforms come from the indirect effect that helps the asset rich (higher interest rates) but hurts the asset poor (higher taxes). Finally, the qualitative similarities in the decomposition across the three pannels reassures that the details behind the SWF are not driving the results.

To sum up, progressive tax reforms lead to higher interest rates and this indirect effect is quantitatively relevant at the individual and the aggregate level. As I show below, this force plays a key role when it comes to understanding the optimal mix of debt and progressivity and leads inequality-averse planners to choose relatively low levels of public debt.

3.3 Optimal steady state mix

To solve the optimal steady state problem, I proceed in two steps. First, I consider restricted versions of (9) assuming that the progressivity is chosen optimally while the debt is held fixed, and viceversa. These *inner problems* help understand how exogenous movements in one instrument affect the optimal use of the other. This will inform the tradeoffs that shape the optimal mix and allows me to characterize the relationship between public debt and progressive income taxes. Then, I turn to the optimal mix and how it depends on social preferences for redistribution. The details behind the computations are in Appendix B.

Let us start with the optimal level of progressivity holding fixed the level of public debt. In the appendix, I show that the first order condition for the optimal choice of progressivity in this inner problem is

$$\frac{\partial \mathcal{W}}{\partial r}\frac{d\boldsymbol{r}}{dp}+\frac{\partial \mathcal{W}}{\partial \tau}\frac{d\boldsymbol{\tau}}{dp}+\frac{\partial \mathcal{W}}{\partial p}=0, \tag{11}$$

where the GE responses $\frac{d\boldsymbol{r}}{dp}$ and $\frac{d\boldsymbol{\tau}}{dp}$ can be computed using part two of Lemma 1. This optimality condition simply states that the aggregate direct and indirect effects of progressive tax reforms must offset each other at the optimum. From the discussion above, we know that the term that is due to the interest rate response tends to be positive. The second term, the part that is due to the increase in the level of taxes, tends to be negative. The sign of the sum of these first two terms– the indirect effect of progressive tax reforms– depends on the planner's taste for redistribution. We saw that, around the baseline economy, the overall indirect effect was positive for the Utilitarian and Efficiency planners and negative for the Rawlsian planner. The last term is the direct effect of progressive tax reforms, and tends to be positive for the Rawlsian planner but negative for both the Utilitarian and Efficiency planners. The results in Figure 5 already hint at the fact that a Rawlsian planner will favor a more progressive tax system than Utilitarian and Efficiency planners: the overall effect on welfare of a progressive tax reform around the baseline economy is positive for the former and negative for the latter. The fact that inequality-averse planners implement more progressive tax systems is not surprising nor new. More interesting is to focus on how the optimal level of p varies with B. I use (11) to compute the optimal level of progressivity numerically and study how it varies with the level of public debt.

Figure 6a summarizes this side of the relationship between the two instruments by tracing out a curve in the space of debt-to-GDP and progressivity for each of the three benchmark planners. These curves show how the optimal progressivity of the tax system varies with the debt-to-GDP ratio for the Rawlsian, Utilitarian, and Efficiency planners. The shape of the grey dashed line reflects the fact that the Rawlsian planner mostly sees the two instruments as complements. The intuition behind this is that an inequality-averse planner focuses on the adverse redistribution effect of issuing debt, which increases interest rates and thus tends to benefit the asset rich. To address the resulting inequality in consumption, such a planner finds it optimal to rely more heavily on progressive income taxes. The Efficiency planner (dotted blue line)

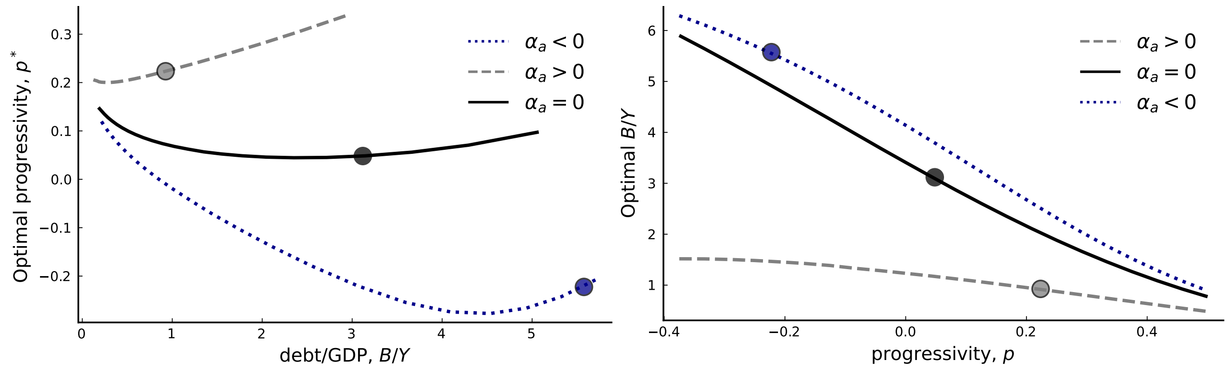

(a) Optimal p as a function of B/Y in the OSS

(b) Optimal B/Y as a function of p in the OSS

Figure 6: Relationship between debt and progressivity in the OSS

mostly considers the positive insurance/liquidity effect of issuing public debt. Public debt increases the supply of safe assets in the economy and helps individuals self-insure against idiosyncratic income risk. As a result, one can afford to provide less insurance via the tax system. A planner with no aversion to inequality will always find it optimal to do so– it increases allocative efficiency and insurance opportunities remain constant. For this reason, the Efficiency planner mostly sees debt as a substitute for progressivity. The Utilitarian planner (solid black line) has a mixture of both the Rawlsian and the Efficiency planner. When the debt-to-GDP ratio is low, there is a lot to gain from increasing the supply of safe assets and the insurance effect dominates. This means that there is less need for insurance through the tax system, resulting in diminished tax progressivity. For high values of debt, rising inequality makes tax-based redistribution more valuable, even at the cost of additional distortions.

I also study the relationship between debt and progressive taxes from the opposite perspective: choose debt optimally and hold the progressivity constant. Here, there is a similar optimality condition that can be used to compute the optimal level of public debt and study how it varies with p. Figure 6b summarizes the results that come out from this alternative inner problem. For a given progressivity of the tax system, the Efficiency planner always prefers more debt than the Utilitarian planner, who in turn prefers more debt than the Rawlsian planner. In other words, the dotted blue line is never below the solid black line, which in turn is never below the grey dashed line. This is intuitive: the Efficiency planner only cares about the insurance/liquidity effect of debt issuance, whereas the Utilitarian and Rawlsian planners also care about the adverse redistribution effect. More interestingly, from this perspective, debt and progressive taxes are unambigously substitutes: the optimal debt-to-GDP ratio decreases with the progressivity of the tax system, regardless of the planner's preferences for redistribution. This is mostly due to the interest rate channel that was isolated above. As the progressivity of the tax system increases, there is more insurance going on through the tax system, so people want to hold less debt. Because of this, the interest rate rises and the cost of servicing a given level of debt increases. This means that the planner will find it optimal to reduce the level of public debt and use the additional resources to finance a more progressive tax system.

The results for these inner problems highlight how different types of planners view these instruments as substitutes or complements, depending on whether debt or progressivity are held fixed. To summarize, if the progressivity of the tax system increases, the level of public debt should fall because of the interest rate channel. On the other hand, if the level of public debt falls, the optimal progressivity of the tax system may increase or decrease due to the competing insurance (+) and redistribution (−) effects of public debt. This sheds light on how the trade-offs along each dimension depend on each other. To understand how these tradeoffs are resolved jointly, I now turn to the optimal mix of the two instruments.

To compute the optimal mix of debt and progressivity, I solve the *outer problem*. This involves searching for the optimal level of public debt along the curves in Figure 6a, or equivalently, the optimal progressivity of the tax system along the curves in Figure 6b. The dots in both panels illustrate the optimal mix for the three benchmark planners.

As expected, aversion to inequality makes optimal labor income taxes more progressive: the optimal mix for the Rawlsian planner (grey dot) is more progressive than the optimal mix for the Utilitarian planner (black dot), which in turn is more progressive than the Efficiency planner (blue dot). In fact, a planner with no aversion to inequality favors a tax system that is regressive. But the key insight is that planners with a taste for redistribution also favor *lower* levels of public debt.

To understand this property of the optimal mix notice that, for any level of public debt, the Rawlsian planner always chooses a more progressive tax system than the Utilitarian planner– graphically, the grey curve is always above the black curve in Figure 6a. The interest rate channel I discussed in Section 3.2 implies that there is upward pressure on the interest rate, making it more costly to service a given stock of public debt. Indeed, suppose the Rawlsian planner where to choose the same debt-to-GDP ratio as the Utilitarian planner (i.e the grey dot is exactly on top of the black dot). Because of the interest rate channel, the Rawlsian planner faces a higher interest rate than the Utilitarian planner. To reiterate, 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 offer a higher interest rate. By reducing the level of public debt, the Rawlsian planner can reduce the marginal cost of debt issuance and use the extra resources to finance a more generous tax and transfer system. Despite the reduction in the supply of safe assets, this deviation improves welfare because inequality averse planners put a higher weight on the adverse redistribution effects of debt issuance and discount the insurance/liquidity effect. As a result, the optimal level of public debt decreases with the planner's aversion to inequality.

Table 2 reports the fiscal policy, macroeconomic aggregates, and distributional measures in the steady state of the calibrated economy (first column) and contrasts them with the optimal steady states of the three benchmark planners (remaining columns). The optimal mix for the Rawlsian planner is fairly close to the fiscal policy of the baseline economy: the optimal progressivity is slightly above the one for the US whereas the optimal debt-to-GDP ratio is somewhat below. The Utilitarian and Efficiency planners implement a tax system that is a lot less progressive than the US and use significantly higher levels of

Table 2: Optimal steady state mix and macroeconomic aggregates

	Baseline	**Rawlsian**		**Utilitarian**		**Efficiency**	
Var.		Optimal	Change	Optimal	Change	Optimal	Change
p	0.181	0.209	0.028	0.048	-0.133	-0.241	-0.422
τ	0.641	0.642	0.001	0.576	-0.065	0.383	-0.258
B/Y	1.4	1.21	-0.190	3.12	1.72	6.50	5.10
Y	0.945	0.934	-0.010	0.974	0.029	0.985	0.041
Z	0.586	0.580	-0.005	0.553	-0.032	0.427	-0.158
C/Y	0.648	0.644	-0.004	0.659	0.011	0.663	0.015
r	2%	1.9%	-0.094 pp	2.91%	0.914 pp	3.52%	1.52 pp
gini_C	0.257	0.252	-0.006	0.276	0.019	0.328	0.071
gini_Y	0.413	0.409	-0.004	0.440	0.027	0.499	0.086
gini_Z	0.344	0.329	-0.014	0.421	0.078	0.589	0.245
gini_A	0.672	0.682	0.010	0.604	-0.069	0.543	-0.123
HtM	0.159	0.199	0.040	0.061	-0.097	0.020	-0.139

Note: This table summarizes optimal policies and the change in macroeconomic aggregates in the optimal steady state.

public debt. Interestingly, the high level of public debt favored by the Efficiency planner achieves the lowest fraction of constrained households. This leads to the lowest inequality in asset holdings across the three benchmark planners. However, because of the reliance on regressive taxes, the optimal steady state distribution for an Efficiency planner features the highest inequality in consumption. The Rawlsian planner, by favoring a more progressive tax system, implements a steady state distribution that features the lowest inequality in consumption. But, by issuing relatively low public debt, the optimal steady state of an inequality-averse planner has a significant fraction of agents at the constraint. To address this, the planner strongly redistributes along the income dimension (see the large difference in the Gini coefficients for before-tax income and after-tax income).

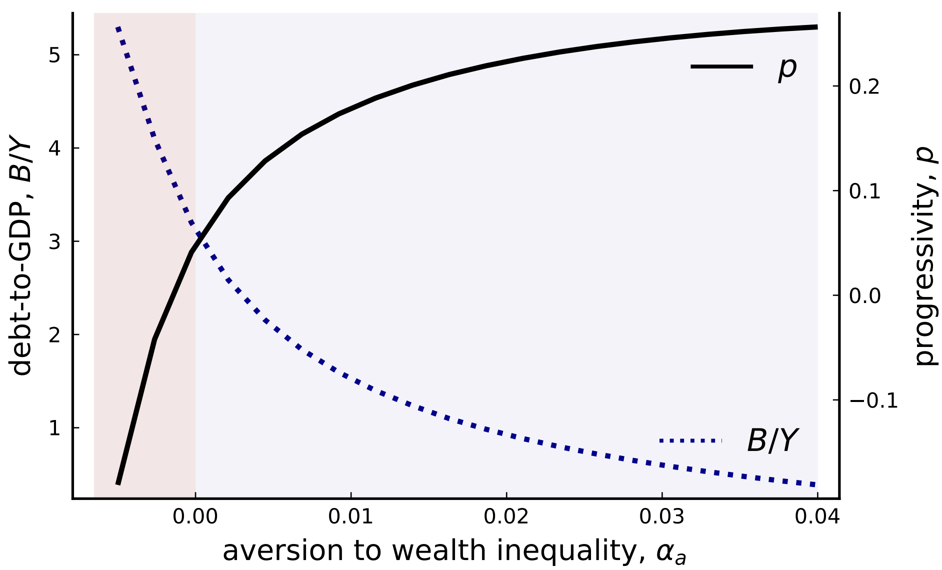

Figure 7: Optimal mix of debt and progressivity with generational planners

The fact that inequality-averse planners implement more progressive tax systems is expected, but their choice of relatively low levels of public debt is surprising. Figure 7 makes this relationship more explicit by varying the planner's aversion to wealth inequality and showing how the optimal mix of debt (left axis) and progressivity (right axis) changes. The colors are meant to capture the idea that Democrats are typically associated with some form of aversion to inequality, whereas Republicans are not. The results are contrary to what one would expect based on the view that Democrats (left-wing) administrations are more fiscally irresponsible than Republican (right-wing) administrations.[10] Again, the interaction between debt and progressive income taxes, driven by the interest rate channel, is key to understand why the optimal level of public debt decreases with the planner's aversion to inequality. We saw earlier that inequality-averse planners prefer relatively lower levels of public debt when p is fixed. But this alone is not enough

[10] One reason for this difference could be that the authorities in charge for making decisions about taxes (Congress) are not speaking to those in charge of the issuance and management of public debt (Treasury). Another possibility is that policy makers are unaware of the fact that the government's stance on progressive taxation influences its capacity to incur debt.

Table 3: Interaction between debt and progressivity in the OSS

	Rawlsian			**Utilitarian**			**Efficiency**		
Var.	optimal	fix p	fix B	optimal	fix p	fix B	optimal	fix p	fix B
p	0.209	0.181	0.212	0.048	0.181	0.045	-0.241	0.181	-0.197
τ	0.642	0.646	0.640	0.576	0.596	0.630	0.383	0.539	0.555
B/Y	1.21	1.28	1.254	3.12	2.27	2.13	6.50	3.25	3.68

Note: This table summarizes the relationship between debt and progressive income taxes in the optimal steady state.

to conclude that inequality-averse planners who also optimize over p will favor low public debt. Indeed, we know that inequality-averse planners will choose more progressive tax systems. Without knowing the relationship in Figure 6b, it is possible that an inequality-averse planner who optimizes jointly over B and p ends up favoring higher levels of public debt. However, the fact that the two instruments are substitutes (in the sense that the optimal level of B decreases as p increases) means that an inequality-averse planner who optimizes over B and p will unambigously favor lower levels of public debt.

Table 3 makes this clear by comparing the results from the joint problem to those from the inner problems where one instrument is held fixed. For each planner, the first column reproduces the optimal mix when both B and p can be adjusted. The second column reports the optimal mix when the progressivity is held fixed at the baseline level. Notice that optimizing only over debt and the level of taxes delivers an optimal debt-to-GDP ratio for the Efficiency planner that is significantly below the one in the optimal mix. The difference between the results for this inner problem and the joint problem shrinks as we move towards a utilitarian SWF and almost dissapears with an inequality-averse planner. This is because of the substitutability between the two instruments discussed above. When the progressivity is held fixed at the baseline value of $p = 0.181$, which is significantly *higher* than the value chosen by the Efficiency planner in the full optimum, the optimal level of public debt is *lower*. The third column optimizes over p holding the level of public debt constant. Across all planners, the optimal choice of progressivity in this inner problem is fairly close to the one that comes out from the joint problem. This suggests that the optimal p is not very sensitive to the choice of B. What really matters for the fact that inequality-averse planners favor lower levels of public debt is the sensitivity of the optimal choice of B to p.

The property of the optimal mix that was isolated in this section holds more generally. First, it is robust to my choice of welfare criteria. In Appendix A, I consider generalized utilitarian planners (7) and vary the parameter that governs the planner's concern for the asset poor. There, I also find that the optimal level of public debt falls and the optimal progressivity increases (see Figure 22). Moreover, consistent with my interpretation of generational planners, a generalized utilitarian planner who only cares about inequality of the initial generation favors a slightly less progressive tax system and somewhat higher levels of public debt for a given $\alpha_a > 0$. In addition, a similar pattern holds for Bénabou planners with the parameter that governs the planner's aversion to inequality. In other words, whenever social preferences are represented by a social welfare function that puts a relatively high weight on the wellbeing of the asset poor, the optimal mix involves low debt and high progressivity. The next section shows that this remains true in the steady state of the unrestricted Ramsey plan, which takes into account transitional dynamics and lets both instruments vary over time.

4 Ramsey steady state

This section introduces the Ramsey problem and studies its steady state. The fact that the Ramsey planner takes into account transitional dynamics makes the problem more complicated. I build on recent work by Auclert et al. (2023) to arrive at a simple characterization for the steady state of the Ramsey plan that is straightforward to implement numerically. After verifying that an interior Ramsey steady state exists, I turn to the optimal long run mix of debt and progressivity. I show that the key insight of the previous section– that inequality averse planners favor lower levels of public debt– carries over to the Ramsey steady state. The computational details are relegated to Appendix C.

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. In contrast to the OSS problem, this problem takes into account the transition path and allows both instruments to vary over time. I will focus on the limiting steady state of the Ramsey plan, leaving the analysis of transitional dynamics towards that steady state to future work. This is possible because, unlike complete-markets models, the optimal long-run policy in this class of models can be computed independently of initial conditions.[11] Formally, the dynamic Ramsey problem is

$$\max_{\{r_t, B_t, p_t, \tau_t\}} \sum_{t=0}^{\infty} \beta^t \mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) \quad \text{s.t} \quad \begin{cases} \mathcal{A}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = B_t, \quad \forall t \\ G + (1 + r_{t-1})B_{t-1} = B_t + \mathcal{T}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right), \ \forall t \end{cases} . \tag{12}$$

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 tax holdings and aggregate tax revenues in period t. Following the previous section, I assume that the Ramsey planner aggregates individual utilities as follows:

$$\mathcal{U}_t\left(\{r_s\}, \{\tau_s\}, \{p_s\}\right) = \int_i \omega_t(\theta_{it}, a_{it}) U(c_{it}, l_{it})\, di, \tag{13}$$

with the weights $\omega_t(\theta, a) \propto \exp\left(-\alpha_\theta \theta - \alpha_a a\right)$.[12] This departs from standard welfare criteria for the same reasons as above. To reiterate, I think of this as a planner that cares about the well being of future generations directly, not just indirectly, via the utility of the initial generation. It nests the utilitarian SWF when

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

[12] The weights now 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.

$\omega_t(\theta, a) = 1$ for all a, all θ, and all t. Here, it's not possible to use a more conventional SWF that depends on individual lifetime utilities with weights that are fixed over time, as in (7). Doing so would make the optimal long-run policy depend on the economy's initial conditions, while the tools we have to solve this problem rely heavily on the independence of the steady state from these initial conditions.

Allowing for departures from a utilitarian welfare criterion turns out to be important for addressing the complications that arise when searching for the steady state of the Ramsey plan in this class of models. At the time of this writing, 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 for a separable CRRA utility function. 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. (2023) 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 interior steady state in this class of models: aggregate asset demand asymptotes to infinity as $\beta(1 + r) \to 1$. The intuition behind this result is that, unlike steady state planners, the Ramsey planner takes into account the benefits of issuing public debt *along the transition*. In the short run, issuing public debt allows the planner to reduce distortionary taxation. This additional benefit calls for more and more liquidity, to the point that the interest rate approaches the discount rate and aggregate asset demand asymptotes to infinity.

Below, I show that this is not the case when the Ramsey planner puts a relatively high weight on the instantaneous utility of agents (or generations) with low asset holdings. In this case, an interior Ramsey steady state exists.

4.2 Existence of RSS with aversion to inequality

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

My approach to (12) builds on the sequence-space method introduced by Auclert et al. (2023), extending it to allow for departures from utilitarian welfare criteria. In Appendix C, I derive a set of first-order conditions for the steady state of the Ramsey problem and elaborate on the computational approach to solve them. Here, I plot the error in the first order condition for the optimal long-run choice of B 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. (2023): 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 high weight on the utility of generations with low asset holdings ($\alpha_a > 0$). In the appendix, I also verify that existence and uniqueness go through when the planner also optimizes over the progressivity of the tax system.

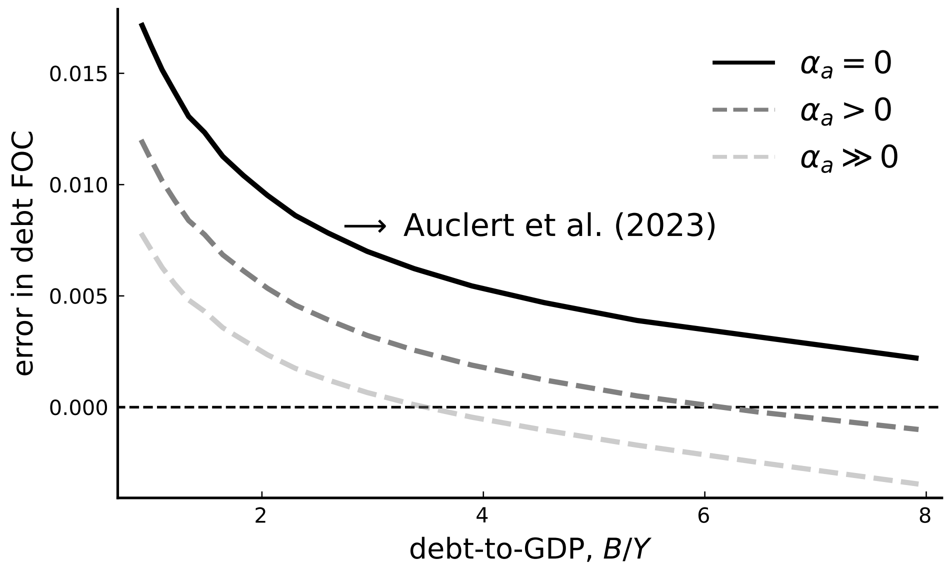

Figure 8: Existence of interior RSS with inequality-averse generational planners

To understand what drives this result, I look at the Jacobians of aggregate utility to changes in the interest rate and changes in the level of taxes: $\mathcal{J}^{\mathcal{U},r}_{t,s} = \frac{\partial \mathcal{U}_t}{\partial r_s}$ and $\mathcal{J}^{\mathcal{U},\tau}_{t,s} = \frac{\partial \mathcal{U}_t}{\partial \tau_s}$.[13] Intuitively, the β-discounted sum of a far-out column of $\mathcal{J}^{\mathcal{U},r}$, $\lim_{s\to\infty}\sum_{t=0}^{\infty}\beta^{t-s}\frac{\partial \mathcal{U}_t}{\partial r_s}$, captures the marginal benefit of increasing the level of public debt around the RSS. The β-discounted sum of a far out-column of $\mathcal{J}^{\mathcal{U},\tau}$, $\lim_{s\to\infty}\sum_{t=0}^{\infty}\beta^{t-s}\frac{\partial \mathcal{U}_t}{\partial \tau_s}$, captures the marginal cost of increasing the level of distortionary taxes around the RSS. At the optimum, the two should be equal (after correcting for GE effects). The figure below plots the far-out columns of these sequence-space Jacobians. It shows that the key difference between a utilitarian planner and an inequality-averse generational planner comes from how they evaluate the marginal benefit of increasing the level of public debt.

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.

The marginal cost of adjusting the level of distortionary taxes is fairly similar across both kinds of planners, reflecting the fact that *all* households bear the costs of debt. This explains why, holding fixed the shape of the tax schedule, the optimal level of public debt decreases with the planner's aversion to inequality. However, this does not yet imply that the optimal level of public debt is lower for inequality-averse

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

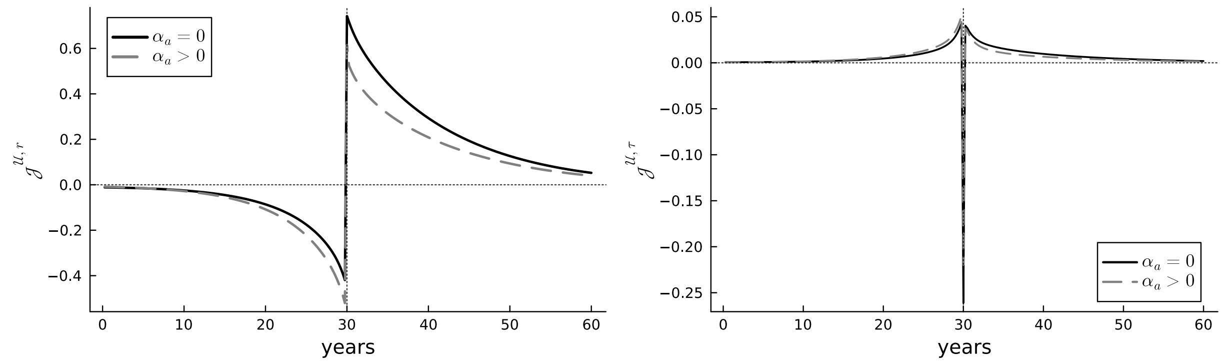

(a) Jacobian of $\mathcal{U}$ to changes in r

(b) Jacobian of $\mathcal{U}$ to changes in τ

Figure 9: Jacobians of $\mathcal{U}$ with and without aversion to inequality

planners. This is because the planner can also change the shape of the tax schedule and, if debt and progressive taxes are complements, this may lead to a higher level of public debt. Even though the interest rate channel implied that this was not the case for steady-state planners, here there is an additional benefit of issuing public debt: it allows the planner to reduce the level of distortionary taxation. In the next section, I show that the interest rate channel is strong enough to offset this force and that the optimal level of public debt is indeed lower for inequality-averse planners.

4.3 OPTIMAL LONG-RUN MIX

Now that we know that an interior RSS exists whenever the planner puts a higher weight on generations with low asset holdings, I turn to the optimal long-run mix of debt and progressivity. The property of the optimal mix uncovered in Section 3 carries over to the steady state of the Ramsey problem. However, the substitutability/complementarity between the two instruments changes. I argue that this is because long run considerations are now discounted appropriately, whereas short run considerations become more prominent.

To see this, I start with the two analogous *inner problems* where I optimize over one instrument while keeping the other constant. Figure 10a displays how the optimal long-run progressivity varies with the debt-to-GDP ratio by tracing out a curve in $(\frac{B}{Y}, p)$-space for the three benchmark planners.[14] It shows that planners who take into account transitional dynamics see the two instruments as substitutes, regardless of their taste for redistribution.

To understand why, observe that the adverse redistribution effects of a higher debt-to-GDP are mostly present in the long run, which the Ramsey planners discounts appropriately. Indeed, a higher debt to GDP ratio leads to higher interest rates in the long run, but during the transition, this need not be the case. Increasing the debt allows the government to reduce taxes in the short run, which boosts disposable income

[14]When B/Y is fixed, it is possible to compute the RSS for the Utilitarian and Efficiency planners.

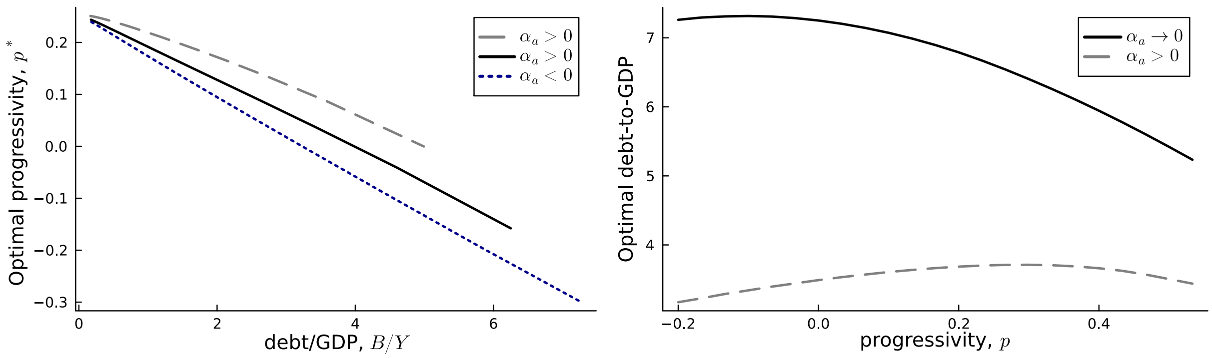

(a) Optimal p as a function of B/Y in the RSS

(b) Optimal B/Y as a function of p in the RSS

Figure 10: Relationship between debt and progressivity in the RSS

for households. As a result, they can save more, and thus the increase in interest rates required to clear the asset market is smaller in the short run than in the long run. In fact, it can even lower interest rates in the short run. This is illustrated in the left panel of Figure 11, which plots the interest rate response to a small one-time permanent change in the level of public debt around the baseline economy. Conversely, the liquidity/insurance benefits of public debt are present in both the short run and the long run. Thus, all planners find it optimal to rely less heavily on progressive income taxes for insurance purposes.

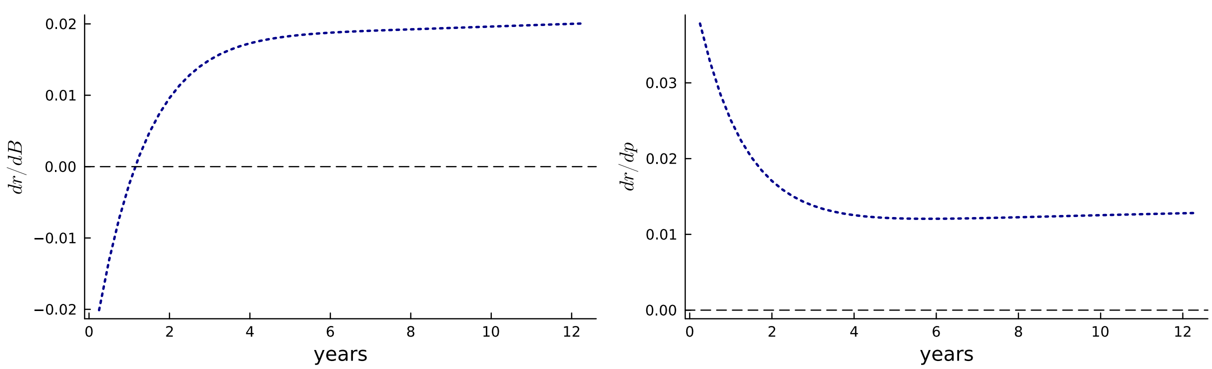

(a) Dynamic first order effects of B on r

(b) Dynamic first order effects of p on r

Figure 11: Dynamic first order effects of debt and progressivity on r

Figure 10b displays the results for the opposite inner problem, where the progressivity is held constant and the debt is chosen optimally. Here, it is only possible to construct the curve for the Rawlsian planner because an interior RSS only exists for inequality-averse planners. But to get a sense of how the relationship varies with aversion to inequality, I also display the results for an inequality-averse planner that is arbitrarily close to the Utilitarian planner ($\alpha_a \to 0$).

Two effects explain the non-monotonic behavior of the curves. First, the now-familiar interest rate effect means that as we move along the x-axis, the *path* of interest rates faced by all planners shifts upward (as shown in the right panel of Figure 11). This increases the cost of financing the public debt in both the short run and the long run, and pushes towards a lower level of debt-to-GDP. At the same time, as the tax system becomes more progressive, it also becomes more distortionary. Issuing debt can counteract some of these distortions by reducing the overall level of taxes, at least in the short run. This additional benefit of accumulating debt may increase the optimal debt-to-GDP ratio in the long run. For a strongly inequality-averse planner (dashed grey line), the two forces appear to offset each other. This results in an optimal debt-to-GDP ratio that is fairly unresponsive to changes in the progressivity of the tax system. However, for a weakly inequality-averse planner (solid black line), the optimal ratio decreases as the progressivity of the tax system increases. Due to the high levels of debt favored by this type of planner, there is not enough fiscal space to take advantage of the short-run benefits of issuing public debt. Instead, the escalating fiscal cost, driven by the interest rate channel, becomes the primary concern.

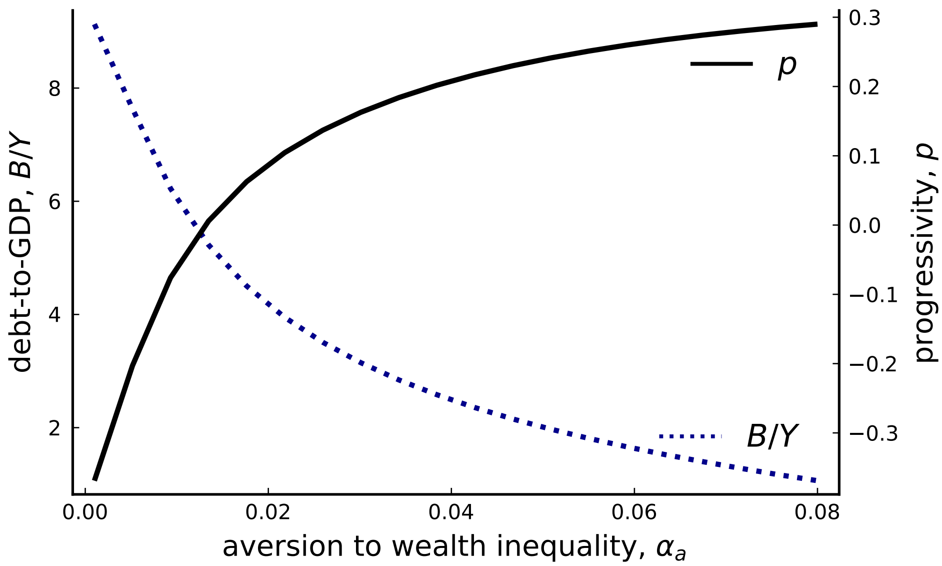

Figure 12: Optimal mix of debt and progressivity in the RSS

Turning to the optimal mix, Figure 12 illustrates how the optimal level of public debt (dotted blue line) and the optimal progressivity of the tax system (solid black line) vary with the planner's aversion to wealth inequality. Echoing the findings in Section 3, the optimal level of public debt decreases with α_a whereas the optimal progressivity of the tax system increases. In addition, one can see how the debt-to-GDP ratio starts to explode as $\alpha_a \to 0$ and the SWF becomes utilitarian, in line with the results in Auclert et al. (2023). A utilitarian Ramsey planner that takes into account the benefits of issuing public debt along the transition finds it optimal to satiate the demand for public debt. On the other hand, giving the planner a taste for redistribution moves the optimum away from the Friedman rule.

The reason why inequality-averse planners continue to favor lower levels of public debt goes back to the intuition behind the results for the OSS problem. The planner's aversion to inequality naturally leads to a more progressive tax system. Because of the interest rate channel, the planner faces a *path* of interest rates that is relatively higher. But, in contrast to the OSS, this logic is incomplete. Here, there is an additional force that works in the opposite direction: a higher p makes the tax system more distortionary and issuing public debt can offset some of these distortions. In other words, as p increases, the benefits of accumulating public debt also increase because this allows the planner to smooth tax distortions in the short run. However, given that the interest rate channel is active in both the short run and the long run, it is strong enough to counteract this force. Thus, planners that favor more progressive tax systems also end up favoring lower levels of public debt.

Despite this qualitative similarity, there are non-trivial quantitative differences between the optimal mix in the RSS and the OSS. Figure 13 compares the optimal mix of the OSS problem with the optimal mix of the RSS problem. The left panel plots the optimal level of public debt and confirms the observation in Angeletos et al. (2022). That paper points out that the OSS problem overestimates the costs of debt issuance and thus underestimates the optimal level of debt. Indeed, the RSS problem takes into account the benefit of increasing debt in the short run, which allows the planner to smooth tax distortions. The OSS problem only takes into account the fiscal cost of increasing the level of public debt. As a result, the optimal level of public debt for a planner that takes into account transitions is *higher*.

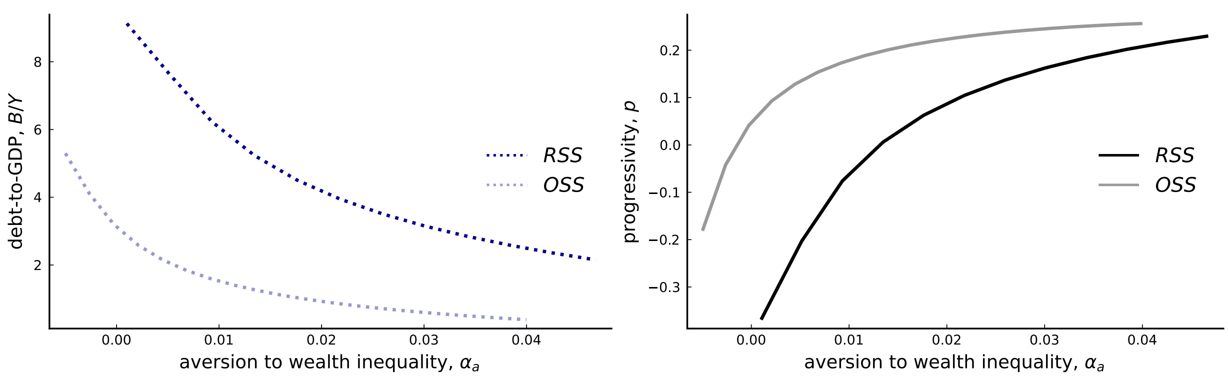

(a) Optimal debt-to-GDP, RSS vs OSS

(b) Optimal progressivity, RSS vs OSS

Figure 13: Optimal mix in the RSS and the OSS

While the OSS problem underestimates the optimal level of public debt, the opposite holds true for taxes. The right panel of Figure 13 plots the optimal progressivity of the tax system as a fuction of the planner's aversion to inequality and shows that the OSS problem *overestimates* the benefits of progressive tax systems. The intuition for this is that the RSS problem takes into account the increase in fiscal costs that are trigerred by progressive tax reforms *along the transition*. Because the interest rate channel operates even in the short run, a progressive tax reform today increases the costs of financing a given path of public debt across all periods. Therefore, a planner that takes into account transitions is less willing to increase the progressivity of the tax system.

To summarize, the key property of the optimal mix that I isolate in this paper holds across both concepts of long-run optimality that have been explored by the literature. At the same time, there are important quantitative differences: the OSS problem underestimates the optimal level of public debt but overestimates the optimal progressivity of the tax system. This is mainly due to the fact that long run considerations are now discounted appropriately and short run considerations become more prominent. I now show that these conclusions 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.

5 EXTENSIONS

5.1 MULTIPLE SAFE ASSETS

In the baseline model, the only supply of bonds outside the household sector comes from the government. This section presents 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. Even though this leads to some quantitative differences, the main result of the paper goes through: planners with stronger preferences for redistribution continue to favor lower levels of public debt. I continue to ignore taxes on savings, so the fiscal instruments are the same as those in Section 3.[15] For simplicity, I focus on the OSS problem and leave the analysis of the RSS in the model with capital to Section 5.3, where the planner is allowed to tax capital income.

For simplicity, assume that capital and government bonds are perfect substitutes. This means that in equilibrium the interest rate on both assets must be the same. In addition, firm optimality implies that the 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

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

Notice that this introduces an elastic supply of safe assets. The interest-rate elasticity of the supply of safe assets depends on the curvature of the production function. When production is linear in capital, the supply of capital is perfectly elastic and the real interest rate is pinned down by the firm's production technology. Despite being somewhat unrealistic, this model economy is useful: it helps understand the role of the interest rate response by shutting down this channel. The results for this "AK" economy are in Appendix D.2. When the interest rate is exogenous, there is no longer a beneficial role for public debt. It only crowds out capital, and all planners find it optimal to issue no debt. Moreover, the curves that summarize the relationship between the optimal progressivity and the debt-to-GDP ratio become flat. This makes sense; the forces isolated in the baseline model are entirely absent in that version of the model.

[15] See Appendix D for more details on the OSS problem in the model with capital.

Here, I work with a standard Cobb-Douglas production function $F(K, L) = K^{\alpha} L^{1-\alpha}$, which generates an imperfectly elastic supply of capital. The details behind the calibration of this version of the model are in Appendix D.1. The results are summarized in Figure 14.

Looking at the results for the inner problems, the interaction between the two instruments echoes the findings in Section 3. When the debt-to-GDP varies exogenously, the Rawlsian planner ($\alpha_a > 0$) sees the two instruments as complements: the optimal progressivity of the tax system increases as the level of public debt increases. The planners with no preference for redistribution see them as substitutes, as they focus on the insurance/liquidity effects of public debt. The benefits of issuing public debt are less prominent because now there is additional asset that households can use to smooth their consumption. This is why the optimal progressivity of the tax system becomes less responsive to changes in the level of public-debt for both the Utilitarian and Efficiency planners. On the other hand, when the progressivity of the tax system varies exogenously, the two instruments are unambiguously substitutes in the sense that the optimal level of debt decreases in response to a more progressive tax system regardless of the planner's taste for redistribution. Again, this is driven by the interest rate channel, which remains quantitatively strong in this version of the model.

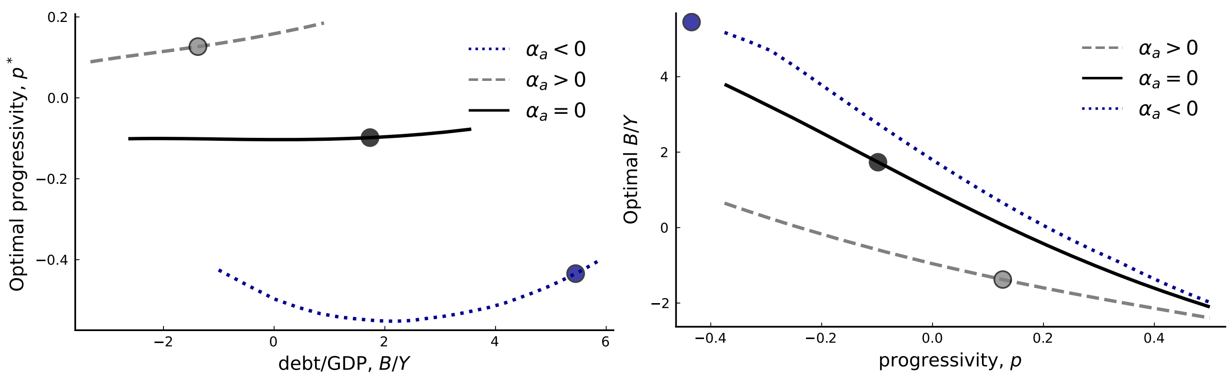

(a) Optimal p as a function of B/Y in the OSS with K

(b) Optimal B/Y as a function of p in the OSS with K

Figure 14: Relationship between debt and progressivity in the model with capital

Turning to the optimal mix, illustrated by the dots in the figures, planners that care about redistribution continue to favor lower levels of public debt. However, there are differences in terms of magnitudes. Across all kinds of planners, the optimal mix becomes less progressive and features lower levels of public debt. The changes along the progressivity dimension are due to the fact that the presence of capital means that there is an additional asset that households can use to self-insure. Along the debt dimension, there is an additional cost to issuing debt– it crowds out capital– and thus, the optimal level of debt is lower. In fact, a Rawlsian planner would prefer to issue no debt at all.

5.2 Alternative labor income tax schedules

I now discuss the properties of the optimal mix with two alternative labor income tax schedules. First, I show that the results hold when I consider a simpler tax system that still captures a form of progressive taxation: linear taxes with a lump-sum intercept. Then, I present the results for a three-parameter version of (2). This CRP_+ tax system introduces a negative intercept into the tax system analyzed so far and gives the planner more flexibility. Some recent papers have argued that this improves the empirical fit to the overall tax and transfer system in the United States.[16] Incorporating lump-sum transfers into the analysis does not change the observation that planners that care about redistribution favor lower levels of debt. However, it does have implications for the optimal shape of average and marginal taxes. As I show below, the optimal tax system now features progressive average taxes but regressive marginal taxes.

Linear taxation with lumpsum transfers

The simplest way to capture a motive for redistribution is to assume that taxes are given by

$$T(y) = (1-\tau)y - T_0,$$

with $T_0 \geq 0$. Werning (2007) studies tax systems of this form in a model with complete markets, which prevents him from relating the optimal level of public debt to redistribution. Flodén (2001) does so in an incomplete-markets model that is similar to the one used here. However, he focuses on how the beneficial effects of debt vanish when transfers are used optimally instead of relating the optimal level of public debt to redistribution. In addition, when debt is the only safe asset in the economy, I find that there is a role for public debt even if transfers are used optimally.

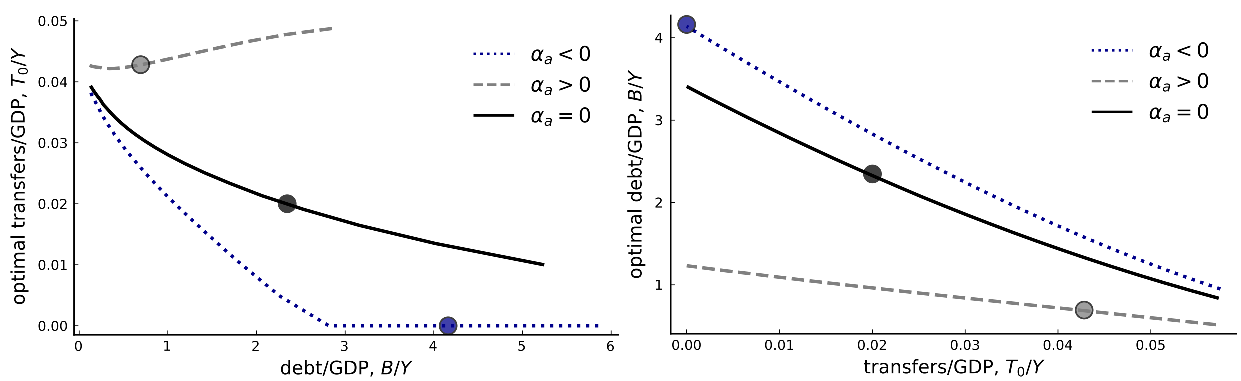

(a) Optimal T_0/Y as a function of B/Y

(b) Optimal B/Y as a function of T_0/Y

Figure 15: Relationship between debt and lumpsum transfers in the OSS

Figure 15 summarizes the results with linear taxes and a lump-sum intercept. The interaction between the two instruments, debt and transfers, is similar to the one uncovered in Section 3 where p indexed the

[16] 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.

progressivity of the tax system. Focusing on the right panel, across all kinds of planners, the optimal level of debt decreases as transfers increase. This is due to a similar interest rate channel. Transfers, by acting as a form of social insurance against idiosyncratic income risk, lower the aggregate demand for safe assets. This makes it more costly to finance a given stock of public debt.

Because of this, planners that care about redistribution continue to favor lower levels of public debt. The Rawlsian planner naturally favors the use of lump-sum transfers because this allows one to separate average from marginal taxes in a way that favors the poor. Because of the interest rate channel, they face a higher interest rate and thus find it optimal to issue lower levels of public debt. Efficiency planners would like to set $T_0 < 0$ in order to fund higher levels of public debt, since they see this as a more effective way to provide insurance.

CRP_+ TAX SYSTEM

Let us now turn to the results with the richer set of tax instruments. Suppose the tax system is

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

where, once again, $T_0 \geq 0$. Ferriere et al. (2022) show that these tax systems can deliver welfare gains that are almost as large as in the second-best allocation in a static economy.

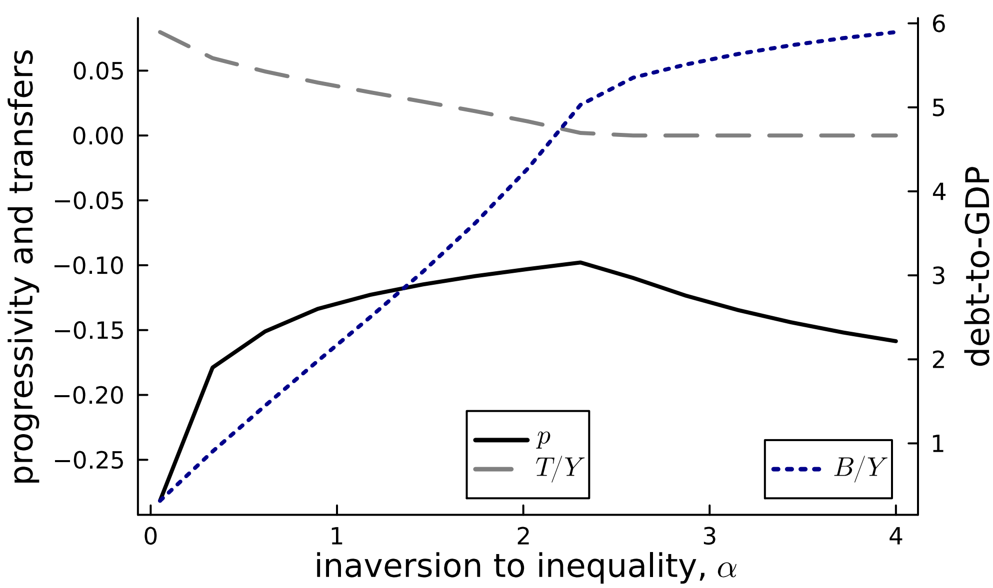

Figure 16: Optimal mix of debt and progressivity with transfers

Figure 16 summarizes the results with CRP_+ taxes by displaying how the optimal use of the three instruments– B (blue dotted line), T_0 (grey dashed line), and p (black solid line)– varies with the parameter that governs the planner's aversion to inequality. As above, inequality-averse planners rely on lump-sum transfers to separate average from marginal taxes and redistribute towards the poor. However, unlike

Section 3, these planners favor a negative p, which means that marginal tax rates decrease with income. In other words, the optimal mix features progressive average taxes but regressive marginal taxes.

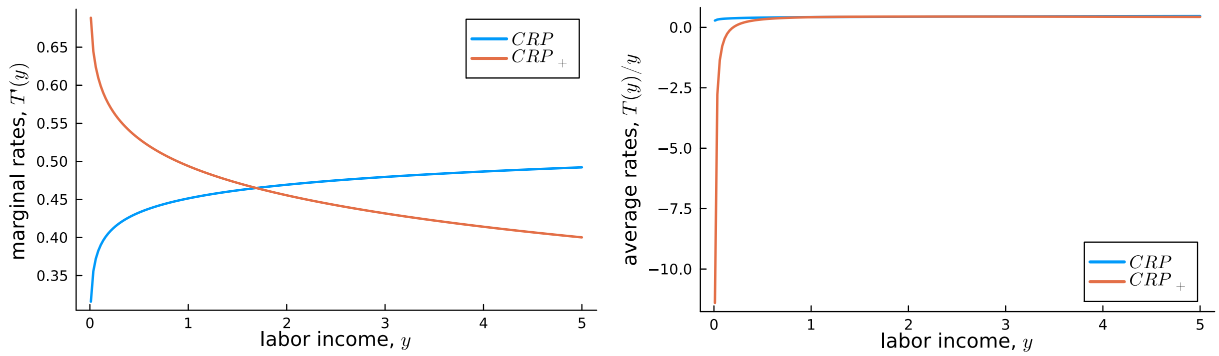

(a) Marginal taxes with and without lumpsum transfers (b) Averages taxes with and without lumpsum transfers

Figure 17: Optimal mix of debt and progressivity in the models with capital

Figure 17 zooms in on the difference between average and marginal taxes in the baseline model and the model with CRP$_+$ tax systems for a Utilitarian planner. Notice the difference in the shape of average and marginal taxes. Like Ferriere et al. (2022) point out, this allows the planner to achieve redistribution while preserving efficiency. This means that here, p is no longer a sufficient statistic for the progressivity of the tax system. So even though $p < 0$ for inequality-averse planners, this does not mean that they favor regressive tax systems. Indeed, a more global measure of progressivity, such as the change in the Gini coefficient as one moves from before-tax to after-tax income distributions, increases with α_a.

5.3 Taxes on savings

I now briefly discuss what happens if the planner has the option to tax savings through a linear tax on capital income τ_k. Of course, this makes no difference in the model where debt is the only safe asset in the economy. In the model with multiple safe assets, it allows the planner to control the capital-labor ratio of the economy, and thus the total supply of safe assets in the economy. In this sense, it brings the results closer to those of the baseline model.

In the appendix, I show the planner chooses τ_k in order to implement the golden rule ($F_K = \delta$) in the OSS and the modified golden rule ($F_K = \delta + \beta^{-1} - 1$) in the RSS.[17] The result for the RSS 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 a higher-than efficient capital stock to help agents self-insure because debt can be used.

Given that the capital-labor ratio has been determined by the production side of the economy, one then

[17] See Appendix E for the OSS and Appendix F for the RSS.

essentially solves a reparametrized version of the baseline model to obtain the results. From this perspective, it is not surprising that the qualitative properties of the results are the same as in the baseline model. Figure 18 illustrates this by plotting how the three instruments– debt (blue dotted line, right axis), progressivity (black solid line, left axis), and capital taxes (grey dashed line, left axis)– vary with the planner's aversion to inequality in the RSS. Focusing on the relationship between τ_k and inequality aversion, inequality-averse planners favor higher levels of capital taxes. This allows them to implement the modified golden rule while restricting the level of public debt to be low.

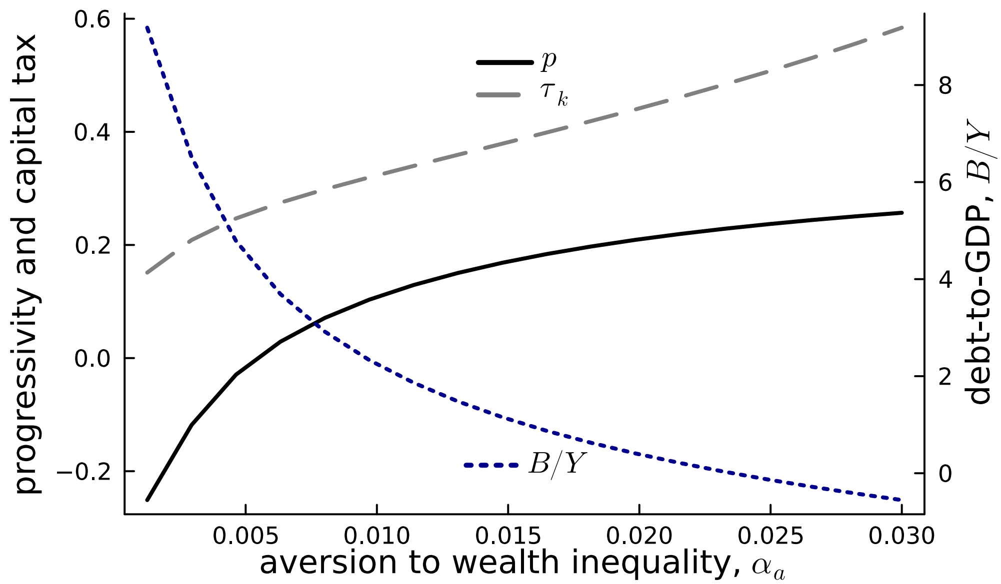

Figure 18: Optimal mix of debt and progressivity in the RSS with capital taxes

6 Inverting the optimium

This section presents an exercise inspired by the *inverse optimal taxation problem* (Bourguignon and Spadaro, 2012; Heathcote and Tsujiyama, 2021). The basic idea is to use 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 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. The takeway is that implied social preferences for redistribution appear inconsistent with both Utilitarian and Rawlsian criteria.

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

6.1 UTILITARIAN PLANNERS & US FISCAL POLICY

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 OSS with a utilitarian SWF analyzed in Section 3.[19] The 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 (0.048). The fact that a Utilitarian planner favors less progressivity is consistent with previous findings in the literature.

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, in the OSS problem without capital taxes (Section 5.1), the optimal level of debt is closer to what we see in the data, but the gap between optimal progressivity and estimated progressivity *increases*. Second, in the problem with capital taxes analyzed in Section 5.3, the situation reverts back to the case without capital. Allowing for more flexible forms of labor income taxation (i.e. lumpsum transfers as in 5.2) 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 Bénabou planner that is more inequality-averse than Utilitarian, 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 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, respectively. The exercise consists in finding the α_a and α_θ such that the solution to the first order conditions of the optimal policy problems 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 19 presents the results when I perform the inversion using the OSS (left panel) and RSS (right panel) solution concepts. The structure of the weights is similar across both panels, consistent with the fact that the qualitative properties of the solution in the optimal steady state and the Ramsey steady state are the same.

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

[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 is also well-below the optimal B.

inequality-averse planners favor lower levels of debt. But if we choose an aversion to inequality that makes the optimal debt-to-GDP in the OSS and RSS consistent with the data, from Figure 7 and Figure 12, 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 implies that the covariance between welfare weights and both asset and labor income is *positive*.

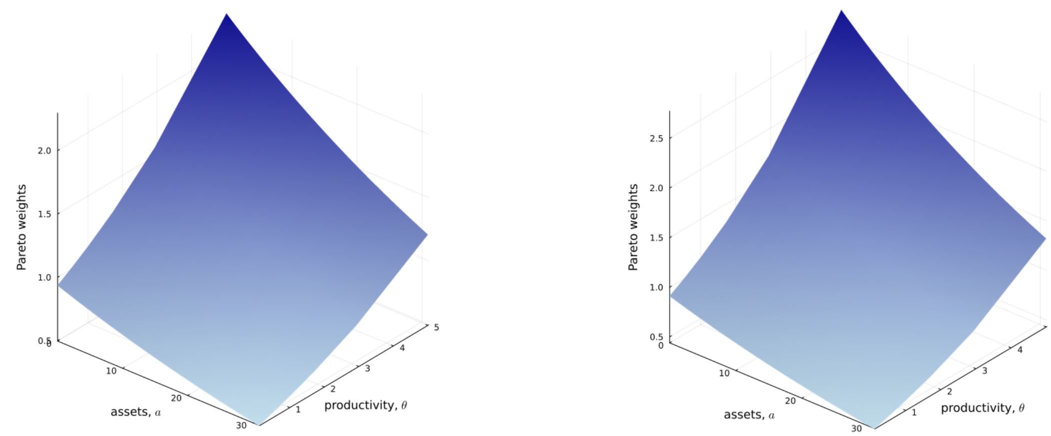

(a) Pareto weights for the US in the OSS problem

(b) Pareto weights for the US in the RSS problem

Figure 19: Inferred Pareto 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 20 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.

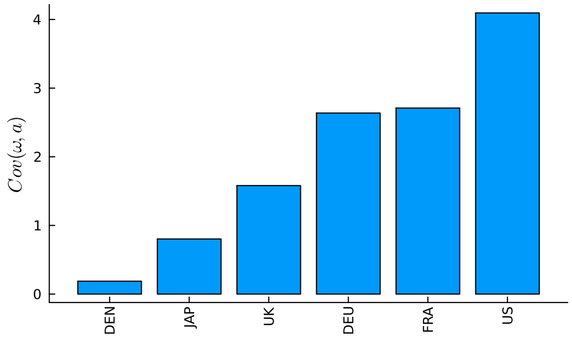

(a) Implied $Cov(\omega, a)$ in the OSS problem

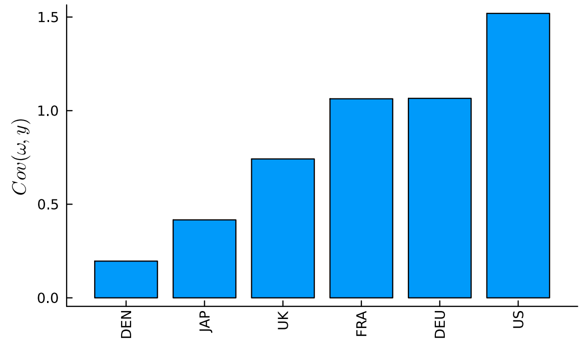

(b) $Cov(\omega, y)$ in the OSS problem

Figure 20: Inferred Pareto weights in selected advanced economies

7 Conclusion

This paper explores the optimal long-run mix of debt and progressivity in a standard heterogeneous-agent model. The key insight is that inequality-averse planners should favor lower levels of public debt. This is driven by a novel interest rate channel that the analysis identifies; namely, progressive income taxes reduce the need to self-insure against idiosyncratic risk, thereby reducing the aggregate demand for safe assets and increasing the fiscal cost of issuing public debt.

This property of the optimal mix appears robust to the presence of multiple safe assets and to restrictions imposed on the tax system. In addition, after comparing two different concepts of long-run optimality, I find that the qualitative properties of the optimal mix are unaffected by the inclusion or exclusion of transitions. However, there are important quantitative differences: the optimal steady state problem underestimates the long-run value of public debt and overestimates the progressivity.

Turning to a technical aspect, I extend the sequence-space approach to optimal policy introduced by Auclert et al. (2023) to accommodate departures from utilitarian welfare criteria. Allowing for some form of aversion to wealth inequality helps overcome the complications that arise when searching for a Ramsey steady state in this class of models.

In terms of policy implications, the results here provide useful insights for the design of fiscal policy by explaining how a government's stance on progressive taxation influences its capacity to incur debt. This underscores the need for coordinated decision-making concerning the level of public debt and the progressivity of the tax schedule.

Finally, the analysis in the paper abstracts from optimal policy *along the transition* and political economy considerations. Both of these are important directions for future research.

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. (2023). Optimal Long-Run Fiscal Policy with Heterogeneous Agents.

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.

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.

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*.

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.

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 Social welfare functions

In this appendix, I compare the different SWF introduced in Section 3. The purpose of this is to verify the connections illustrated in Figure 1 and show that the results are not driven by my choice of SWF.

A.1 Details on generational planners

To illustrate the basic idea behind generational planners, think of an economy with two dynasties A and B and two generations $t = 0, 1$. For simplicity, there is no risk. Following (6), a generational planner in this environment evaluates welfare according to

$$\mathcal{W} = \omega_0 u_0^A + (1-\omega_0) u_0^B + \beta\omega_1 u_1^A + \beta\omega_1 u_1^B.$$

where u_j^i is short-hand notation for the instantaneous utility of an agent from dynasty i and generation j. Here, ω_j denotes the weight on generation j of dynasty A. Notice that the expression above can be rewritten as

$$\mathcal{W} = \omega_0 \underbrace{\left(u_0^A + \beta u_1^A\right)}_{=V_0^A} + (1-\omega_0) \underbrace{\left(u_0^B + \beta u_1^B\right)}_{=V_0^B} + \beta\left(\omega_1 - \omega_0\right) u_1^A + \beta\left(\omega_0 - \omega_1\right) u_1^B.$$

This shows that a generational planner is not paternalistic with respect to the initial generation. Moreover, it satisfies the Pareto principle with respect to the welfare of the initial generation. But because the *effective* weights $\omega_1 - \omega_0$ on future generations can be negative, they do not satisfy the Pareto principle with respect to the welfare of the future generations. If dynasty A accumulates relative more assets than dynasty B, then $\omega_1 < \omega_0$ and the planner discounts the welfare of the second generation of dynasty A and increases the weight on the second generation of dynasty B.

A.2 Bénabou planners

Proposition 1 uses Lemma 1 to derive an expression for the change in social welfare in response to a small permanent change in the progressivity of the tax system with Bénabou planners. To simplify the algebra, it does so under the assumption that the elasticity of intertemporal substitution is equal to one.

Proposition 1 *Assume the $EIS = 1$. The response of social welfare $d\mathcal{W}$ to a small permanent change in progressivity dp is given by*

$$d\mathcal{W} = \frac{1}{1-\beta}\mathbb{E}_x\left[u'(\boldsymbol{c}(x))\left(\boldsymbol{y}(x)^{1-p} d\boldsymbol{\tau} - \boldsymbol{z}(x)\log\boldsymbol{y}(x) + a\, d\boldsymbol{r}\right)\right] + Cov(\gamma, dV) + \Lambda, \tag{A.1}$$

with the weights $\gamma(x) = \left(\frac{\bar{c}(x)}{\bar{C}}\right)^{1-\frac{1}{\alpha}}$. □

For intuition, consider the response of social welfare with a utilitarian welfare criterion (i.e. when $\alpha = 1$). In this case, the expression in (A.1) becomes

$$d\mathcal{W} = \frac{1}{1-\beta}\mathbb{E}_x\left[u'(\boldsymbol{c}(x))\left(\boldsymbol{y}(x)^{1-p} d\boldsymbol{\tau} - \boldsymbol{z}(x)\log\boldsymbol{y}(x) + a\, d\boldsymbol{r}\right)\right] + \Lambda.$$

This says that the effect on welfare is given by the sum of the average direct and indirect effects in the cross-section and an additional *distributional* effect Λ that captures the fact that the perturbation in the tax system changes the stationary distribution of the economy. In the general case ($\alpha \neq 1$), there is an additional correction that captures the planner's preferences for redistribution. This correction depends on the covariance between the weights γ and the responses of individual outcomes dV. Thus, the structure of these weights governs the differences in welfare assesments across planners.

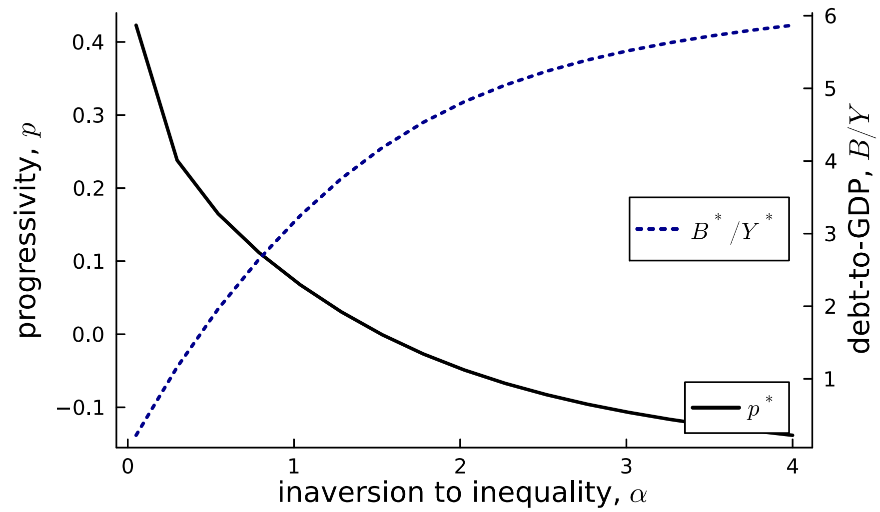

Figure 21: Optimal mix with Bénabou planners

A.3 Generalized utilitarian planners

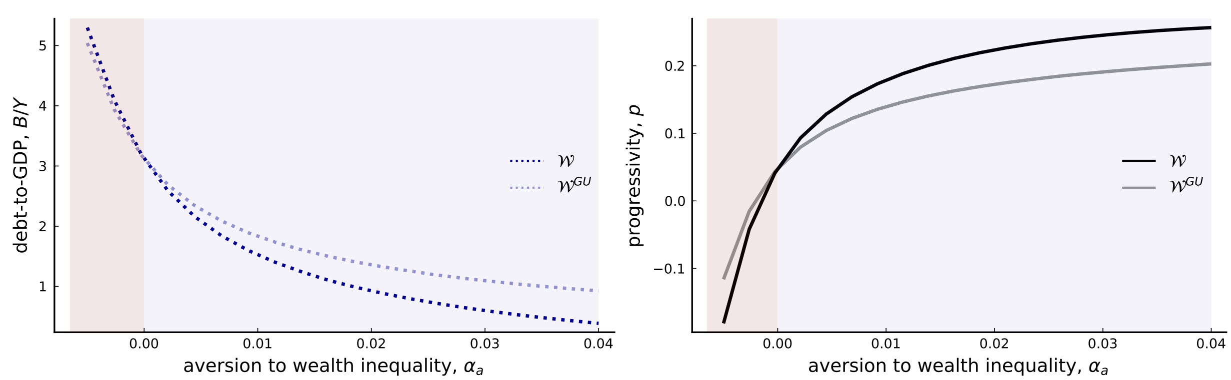

(a) Optimal debt-to-GDP, $\mathcal{W}$ vs $\mathcal{W}^{GU}$

(b) Optimal progressivity, $\mathcal{W}$ vs $\mathcal{W}^{GU}$

Figure 22: Optimal mix with generational and generalized utilitarian planners

B Appendix to OSS 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 allow for a sufficient-statistic representation of the optimality conditions.

B.1 Optimality conditions for OSS problem

To take into account constraints in problem (9), 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:

$$\frac{\partial \mathcal{W}}{\partial r} \cdot \frac{\partial \boldsymbol{r}}{\partial p} + \frac{\partial \mathcal{W}}{\partial p} + \frac{\partial \mathcal{W}}{\partial \tau} \cdot \frac{\partial \boldsymbol{\tau}}{\partial p} = 0 \tag{B.1}$$

$$\frac{\partial \mathcal{W}}{\partial r} \cdot \frac{\partial \boldsymbol{r}}{\partial B} + \frac{\partial \mathcal{W}}{\partial \tau} \cdot \frac{\partial \boldsymbol{\tau}}{\partial B} = 0 \tag{B.2}$$

To solve for the GE derivatives in (B.1), one can use part two of Lemma 1 in the paper. The GE derivatives in (B.2) can be obtained from the following system:

$$\begin{bmatrix} \frac{\partial \mathcal{A}}{\partial r} & \frac{\partial \mathcal{A}}{\partial \tau} \\ \frac{\partial \mathcal{T}}{\partial r} - B & \frac{\partial \mathcal{T}}{\partial \tau} \end{bmatrix} \begin{bmatrix} \frac{\partial \boldsymbol{r}}{\partial B} \\ \frac{\partial \boldsymbol{\tau}}{\partial B} \end{bmatrix} = \begin{bmatrix} 1 \\ r \end{bmatrix}. \tag{B.3}$$

I use (B.1) and (B.2) to reduce the optimization problem in (9) to a (numerical) root-finding problem.[20]

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 (B.3) for $\frac{\partial \boldsymbol{r}}{\partial B}$ and $\frac{\partial \boldsymbol{\tau}}{\partial B}$.

4. Check whether (B.1) and (B.2) are satisfied. If they are, stop. Otherwise, adjust B and p and repeat the process.

[20] I check that the solution is unique, so 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.

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. In addition, to avoid multi-dimensional root-finding algorithms, it helps to separate the problem into two steps. So in step (4), I first check whether (B.1) 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.[21] Once (B.1) is satisfied at the candidate level of r, I check whether (B.2) 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 (9) is equivalent to:

$$\max_{\{r,\tau,p,B\}} \mathcal{W}(r,\tau,p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(r,\tau,p) = B, \\ \mathcal{C}(r,\tau,p) + G = \mathcal{Y}(r,\tau,p) \end{cases} . \tag{B.4}$$

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

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

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

$$\frac{\partial \mathcal{W}}{\partial r} = \frac{\frac{\partial \mathcal{W}}{\partial \tau}}{\frac{\partial \mathcal{C}}{\partial \tau} - \frac{\partial \mathcal{Y}}{\partial \tau}} \left\{ \frac{\partial \mathcal{C}}{\partial r} - \frac{\partial \mathcal{Y}}{\partial r} \right\} \tag{B.5}$$

$$\frac{\partial \mathcal{W}}{\partial p} = \frac{\frac{\partial \mathcal{W}}{\partial \tau}}{\frac{\partial \mathcal{C}}{\partial \tau} - \frac{\partial \mathcal{Y}}{\partial \tau}} \left\{ \frac{\partial \mathcal{C}}{\partial p} - \frac{\partial \mathcal{Y}}{\partial p} \right\} \tag{B.6}$$

[21] Here, I also check that there is a unique value of p that solves (B.1).

(B.5) and (B.6), combined with the goods market clearing condition, can be used to solve for the three unknowns $\{r, \tau, p\}$. The advantage of this formulation is that it does not require solving for general equilibrium in each iteration. 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. I 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. 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

$$B^{OSS} = \arg\max_B \{\mathcal{W}(\boldsymbol{r}(B, \tau, p), \tau, p) + \lambda (\mathcal{T}(\boldsymbol{r}(B, \tau, p), \tau, p) - \boldsymbol{r}(B, \tau, p)B - G)\}. \tag{B.7}$$

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

$$\mathcal{W}_r \frac{\partial \boldsymbol{r}}{\partial B} + \lambda \left\{ \mathcal{T}_r \frac{\partial \boldsymbol{r}}{\partial B} - \frac{\partial \boldsymbol{r}}{\partial B} B^{OSS} - \boldsymbol{r} \right\} = 0. \tag{B.8}$$

Define the social marginal value of public debt, in dollar terms, as $\Gamma \equiv \frac{\mathcal{W}_r}{\lambda} + \mathcal{T}_r$. Similarly, let $\mathcal{E}_B^r \equiv \frac{\partial \boldsymbol{r}}{\partial B} \frac{B^{OSS}}{1+\boldsymbol{r}}$ denote the elasticity of interest rates with respect to changes in the level of public debt *at the optimum*. Then, (B.8) can be written as

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

Solving for B^{OSS} yields

$$B^{OSS} = \frac{\mathcal{E}_B^r}{\mathcal{E}_B^r + \frac{\boldsymbol{r}}{1+\boldsymbol{r}}} \times \Gamma. \tag{B.9}$$

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 $\boldsymbol{r}$.

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 can, in principle, affect the optimal level of public debt through three channels: $\boldsymbol{r}$, $\mathcal{E}_B^r$, and Γ. The interest-rate channel, emphasized in the paper, implies that a more progressive tax system increases 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 need for liquidity. This would only re-inforce the effects driven by the interest rate channel. With income effects, however, the response of 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.

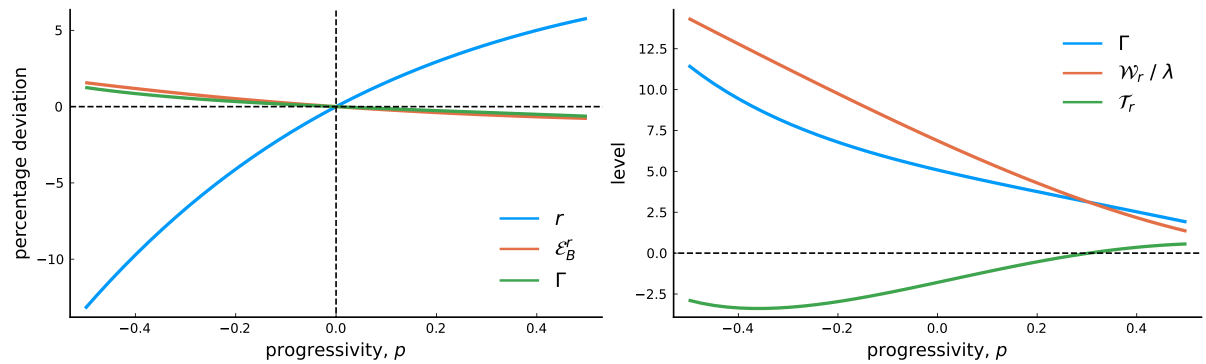

(a) r, $\mathcal{E}_B^r$, and Γ as progressivity varies

(b) Unpacking Γ as progressivity varies

Figure 23: Sufficient statistics and progressivity in the OSS

B.5 Additional OSS figures

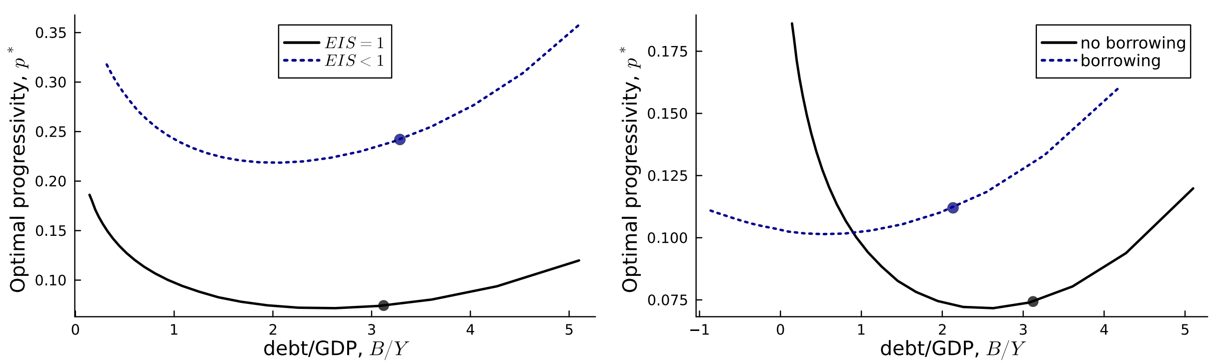

(a) Optimal mix with low and high EIS

(b) Optimal mix with and without borrowing

Figure 24: Comparatives statics with respect to the EIS and borrowing constraints

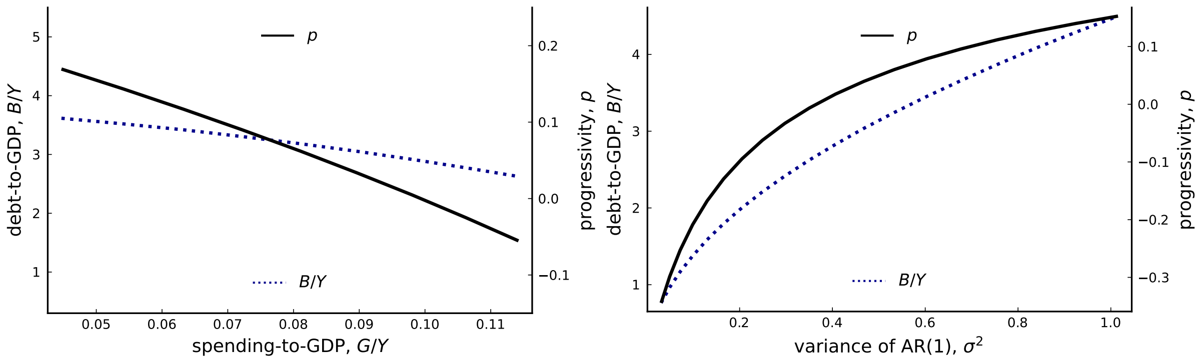

(a) Optimal mix and government spending

(b) Optimal mix and idiosyncratic income risk

Figure 25: Comparatives statics with respect to government spending and idiosyncratic income risk

C Appendix to RSS problem

This appendix provides additional details on the RSS 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. (2023). 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 that simplify the derivations and allow for a sufficient-statistic representation of the optimal long-run value of public debt.

C.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

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

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

$$\sum_{t=0}^{\infty} \beta^t \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial r_s} \frac{\partial \boldsymbol{r}_s}{\partial p_u} + \sum_{t=0}^{\infty} \beta^t \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial \tau_s} \frac{\partial \boldsymbol{\tau}_s}{\partial p_u} + \sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial p_u} = 0 \tag{C.1}$$

$$\sum_{t=0}^{\infty} \beta^t \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial r_s} \frac{\partial \boldsymbol{r}_s}{\partial B_u} + \sum_{t=0}^{\infty} \beta^t \sum_{s=0}^{\infty} \frac{\partial \mathcal{U}_t}{\partial \tau_s} \frac{\partial \boldsymbol{\tau}_s}{\partial B_u} = 0 \tag{C.2}$$

To simplify (C.1) and (C.2), let $u \to \infty$. Now, for any sequence space function $F_t(\{X_s\})$, define the discounted sum $\mathcal{S}_{F,X}$ of the long-run response

$$\mathcal{S}_{F,X} \equiv \lim_{u\to\infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial F_t}{\partial X_u}.$$

Then, use the *quasi-Toeplitz* property of Jacobians in stationary models and apply the convolution theorem to write the composition of Jacobians as the product of discounted sums.[22] This yields the two *scalar equations*

$$\mathcal{S}_{\mathcal{U},r} \cdot \mathcal{S}_{\boldsymbol{r},p} + \mathcal{S}_{\mathcal{U},\tau} \cdot \mathcal{S}_{\boldsymbol{\tau},p} + \mathcal{S}_{\mathcal{U},p} = 0 \tag{C.3}$$

$$\mathcal{S}_{\mathcal{U},r} \cdot \mathcal{S}_{\boldsymbol{r},B} + \mathcal{S}_{\mathcal{U},\tau} \cdot \mathcal{S}_{\boldsymbol{\tau},B} = 0 \tag{C.4}$$

The discounted sum of the general equilibrium long-run responses $\mathcal{S}_{\boldsymbol{r},p}$, $\mathcal{S}_{\boldsymbol{\tau},p}$, $\mathcal{S}_{\boldsymbol{r},B}$, and $\mathcal{S}_{\boldsymbol{\tau},B}$ can be computed from the system of equations derived below.

SYSTEM FOR GE RESPONSES

I start by deriving a system of equations for $\mathcal{S}_{\boldsymbol{r},p}$ and $\mathcal{S}_{\boldsymbol{\tau},p}$. 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 FOC, this becomes

$$\mathcal{S}_{\mathcal{A},r} \cdot \mathcal{S}_{\boldsymbol{r},p} + \mathcal{S}_{\mathcal{A},\tau} \cdot \mathcal{S}_{\boldsymbol{\tau},p} + \mathcal{S}_{\mathcal{A},p} = 0.$$

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

$$\left(\mathcal{S}_{\mathcal{T},r} - B\right) \mathcal{S}_{\boldsymbol{r},p} + \mathcal{S}_{\mathcal{T},\tau} \cdot \mathcal{S}_{\boldsymbol{\tau},p} + \mathcal{S}_{\mathcal{T},p} = 0.$$

Thus, we can solve for $\mathcal{S}_{\boldsymbol{r},p}$ and $\mathcal{S}_{\boldsymbol{\tau},p}$ from

$$\begin{bmatrix} \mathcal{S}_{\mathcal{A},r} & \mathcal{S}_{\mathcal{A},\tau} \\ \mathcal{S}_{\mathcal{T},r} - B & \mathcal{S}_{\mathcal{T},\tau} \end{bmatrix} \begin{bmatrix} \mathcal{S}_{\boldsymbol{r},p} \\ \mathcal{S}_{\boldsymbol{\tau},p} \end{bmatrix} = \begin{bmatrix} -\mathcal{S}_{\mathcal{A},p} \\ -\mathcal{S}_{\mathcal{T},p} \end{bmatrix}. \tag{C.5}$$

Analogous derivations for a perturbation dB_u lead to a system for $\mathcal{S}_{\boldsymbol{r},B}$ and $\mathcal{S}_{\boldsymbol{\tau},B}$

$$\begin{bmatrix} \mathcal{S}_{\mathcal{A},r} & \mathcal{S}_{\mathcal{A},\tau} \\ \mathcal{S}_{\mathcal{T},r} - B & \mathcal{S}_{\mathcal{T},\tau} \end{bmatrix} \begin{bmatrix} \mathcal{S}_{\boldsymbol{r},B} \\ \mathcal{S}_{\boldsymbol{\tau},B} \end{bmatrix} = \begin{bmatrix} 1 \\ \beta(1+r) - 1 \end{bmatrix}. \tag{C.6}$$

[22] See Auclert et al. (2023) for details.

C.2 Computing the RSS

The algorithm to solve for the RSS relies on (C.3) and (C.4). 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. (2023) 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 $\mathcal{S}_{\mathcal{A},r}$. 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 $dA = d\boldsymbol{a} \cdot D^{ss} + dD \cdot \boldsymbol{a}^{ss}$.[23] Then, $\mathcal{S}_{\mathcal{A},r} = \frac{dA}{dr}$.

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

1. Given a candidate r and p, solve for the level of taxes τ that ensures that the government budget constraint holds.

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

3. Solve for the discounted sum of the general equilibrium long-run responses using (C.5) and (C.6).

4. Check whether (C.3) and (C.4) are satisfied. If not, update r and p and repeat steps 1-4.

I use a two-step procedure to avoid multi-dimensional root finding algorithms, as in Appendix B.2. As explained in the main text, I verify that there is a unique interior solution to (C.4), fixing p. I also verify that there is a unique solution to (C.3). See Appendix C.5 for a figure.

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

C.3 Alternative formulation of the RSS problem

To simplify the derivations, we can follow the approach in Appendix B.3. By Walras' Law, problem (12) is equivalent to

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

Now, given choices of $\{r_t, \tau_t, p_t\}$, the asset market clearing condition pins down $\{B_t\}$. Therefore, after dropping all redundant choice variables and the associated constraints, the RSS 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 for this problem are

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial r_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial r_u} - \frac{\partial \mathcal{C}_t}{\partial r_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{C.7}$$

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial \tau_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial \tau_u} - \frac{\partial \mathcal{C}_t}{\partial \tau_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{C.8}$$

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial p_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(\frac{\partial \mathcal{Y}_t}{\partial p_u} - \frac{\partial \mathcal{C}_t}{\partial p_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots, \tag{C.9}$$

together with goods-market clearing period-by-period. Again, the key is to let $u \to \infty$. Then, in the limiting steady state of the Ramsey plan, these become:[24]

$$\mathcal{S}_{\mathcal{U},r} + \lambda^{GM} \left(\mathcal{S}_{\mathcal{Y},r} - \mathcal{S}_{\mathcal{C},r} \right) = 0, \tag{C.10}$$

$$\mathcal{S}_{\mathcal{U},\tau} + \lambda^{GM} \left(\mathcal{S}_{\mathcal{Y},\tau} - \mathcal{S}_{\mathcal{C},\tau} \right) = 0, \tag{C.11}$$

$$\mathcal{S}_{\mathcal{U},p} + \lambda^{GM} \left(\mathcal{S}_{\mathcal{Y},p} - \mathcal{S}_{\mathcal{C},p} \right) = 0. \tag{C.12}$$

Using the second one to eliminate λ^{GM},

$$\mathcal{S}_{\mathcal{U},r} = \frac{\mathcal{S}_{\mathcal{U},\tau}}{\mathcal{S}_{\mathcal{C},\tau} - \mathcal{S}_{\mathcal{Y},\tau}} \left(\mathcal{S}_{\mathcal{C},r} - \mathcal{S}_{\mathcal{Y},r} \right), \tag{C.13}$$

$$\mathcal{S}_{\mathcal{U},p} = \frac{\mathcal{S}_{\mathcal{U},\tau}}{\mathcal{S}_{\mathcal{C},\tau} - \mathcal{S}_{\mathcal{Y},\tau}} \left(\mathcal{S}_{\mathcal{C},p} - \mathcal{S}_{\mathcal{Y},p} \right). \tag{C.14}$$

(C.13) and (C.14), 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. I verify that the solution implied by this formulation is consistent with the one in Appendix C.1.

[24] Use the definition of the discounted long-run response and the quasi-Toeplitz property of Jacobians in stationary models. The mysterious step involving the convolution theorem is no longer needed.

C.4 Sufficient-statistic representation for optimal choice of B in the RSS

The optimal level of public debt in the Ramsey steady state B^{RSS} can also be expressed in terms of the objects that appeared in Appendix B.4. To see this, note that, for any $u = 0, 1, \ldots$, the following must be true

$$\sum_{t=0}^{\infty}\sum_{s=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \sum_{t=0}^{\infty}\sum_{s=0}^{\infty} \beta^{t-u}\lambda_t \frac{\partial \mathcal{T}_t}{\partial r_s}\frac{\partial \boldsymbol{r}_s}{\partial B_u} + \lambda_u - \beta\lambda_{u+1}(1+\boldsymbol{r}_u) - \sum_{t=0}^{\infty} \beta^{t-u}\lambda_t \frac{\partial \boldsymbol{r}_t}{\partial B_u} B_{t-1} = 0. \tag{C.15}$$

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 (C.15) does not hold, then a small perturbation dB_u would increase social welfare, contradicting the optimality of the Ramsey plan. Now, letting $u \to \infty$ and restricting attention to the steady state, (C.15) becomes

$$\mathcal{S}_{\mathcal{U},r}\, \mathcal{S}_{\boldsymbol{r},B} + \lambda\, \mathcal{S}_{\mathcal{T},r}\, \mathcal{S}_{\boldsymbol{r},B} + \lambda\{1 - \beta(1+\boldsymbol{r})\} - \lambda\, \mathcal{S}_{\boldsymbol{r},B}\, B^{RSS} = 0$$

Define the marginal social value of public debt and the discounted long-run elasticity around the RSS as $\Gamma^{RSS} \equiv \frac{\mathcal{S}_{\mathcal{U},r}}{\lambda} + \mathcal{S}_{\mathcal{T},r}$ and $\mathcal{S}_{\mathcal{E}_B^r} \equiv \frac{\mathcal{S}_{\boldsymbol{r},B} B^{RSS}}{1+\boldsymbol{r}}$, respectively. Rearranging terms and solving for B^{RSS},

$$B^{RSS} = \frac{\mathcal{S}_{\mathcal{E}_B^r}}{\beta - \frac{1}{1+\boldsymbol{r}} + \mathcal{S}_{\mathcal{E}_B^r}} \times \Gamma^{RSS} \tag{C.16}$$

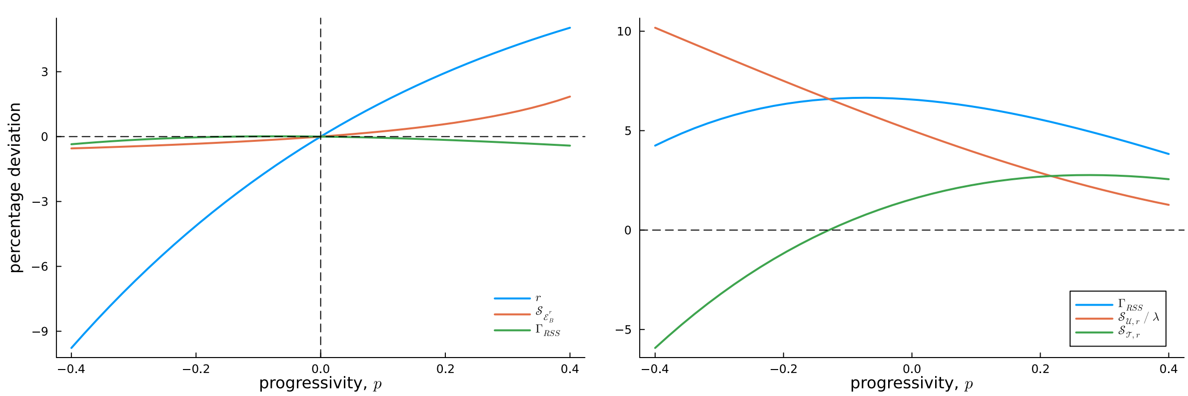

(a) r, $\mathcal{S}_{\mathcal{E}_B^r}$, and Γ^{RSS} as progressivity varies (b) Unpacking Γ^{RSS} as progressivity varies

Figure 26: Sufficient statistics and progressivity in the RSS

C.5 Additional RSS figures

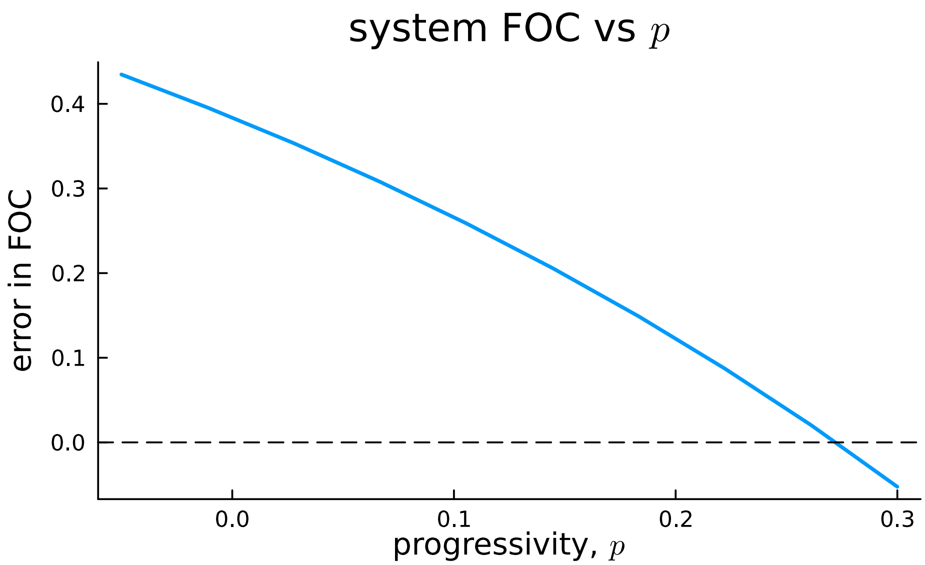

Figure 27: Verifying uniqueness of solution to (C.3)

D OSS PROBLEM IN THE MODEL WITH CAPITAL

In the model with capital, the OSS problem is

$$\max_{\{r,w,\tau,p,B,K,L\}} \mathcal{W}(r,\tau,p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(r,\tau,p) = B + K, \\ G + rB = \mathcal{T}(r,\tau,p) + \tau_k \, r\mathcal{A}(r,\tau,p), \\ w = F_L(K,L), \; r = F_K(K,L) - \delta, \\ \mathcal{L}(r,\tau,p) = L \end{cases} . \tag{D.1}$$

To be written

D.1 Calibration for model with capital

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.[25] The parameters of the income process 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 4 summarizes the parameter values.

Table 4: 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
σ_ϵ	variance of AR(1)	0.033	B/Y	debt-to-GDP	0.615
EIS	curvature in u	1	p	progressivity of taxes	0.181
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

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

D.2 Results in AK economy

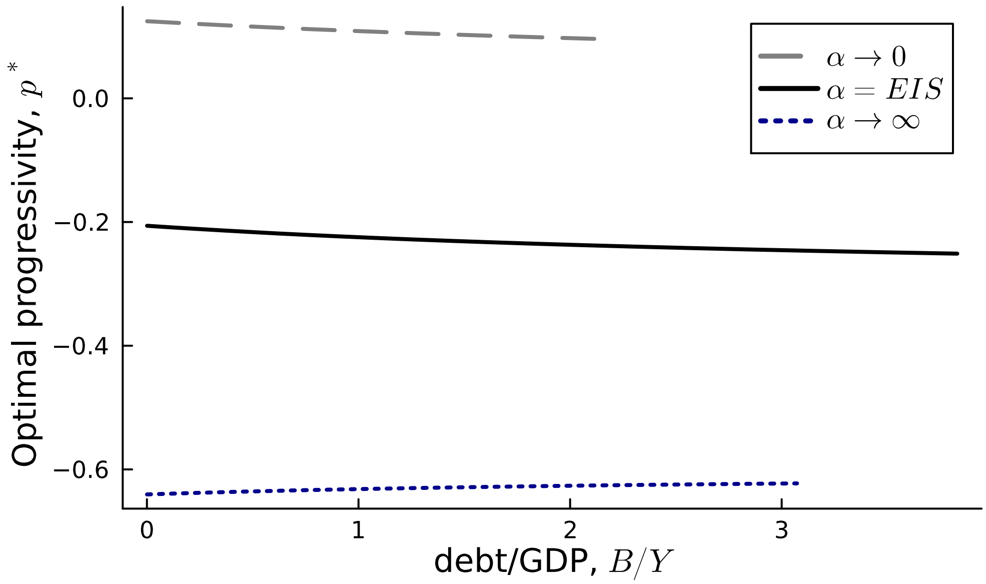

Figure 28: Optimal p as a function of B/Y in the AK economy

E OSS PROBLEM IN THE MODEL WITH CAPITAL AND τ_k

In the model with capital, the OSS problem can be written as

$$\max_{\{\bar{r},\bar{w},r,p,B,K,L\}} \mathcal{W}(\bar{r},\bar{w},p) \quad \text{s.t} \quad \begin{cases} \mathcal{A}(\bar{r},\bar{w},p) = B + K, \\ \mathcal{C}(\bar{r},\bar{w},p) + G = F(K,L) - \delta K, \\ w = F_L(K,L),\ r = F_K(K,L) - \delta, \\ \mathcal{L}(\bar{r},\bar{w},p) = L \end{cases} . \tag{E.1}$$

Here, $\bar{r} \equiv (1-\tau_k)(r+\delta)$ and $\bar{w} \equiv \tau w^{1-p}$ denote the after-tax interest rate and the after-tax wage, respectively. It is convenient to work with the goods market clearing condition and ignore the budget constraint of the government, which can be dropped because of Walras' Law. The only variables that enter the household's problem directly are the after-tax interest rate, the after-tax wage, and the progressivity of the tax system. This is why e.g. aggregate assets $\mathcal{A}$ are a function of these variables only.

E.1 PROVING THE OPTIMALITY OF THE GOLDEN RULE

Choosing $\bar{r}, \bar{w}, p$ and K pins down the level of public debt B through the asset market clearing condition. Similarly, given these choice variables, the labor market clearing condition pins down labor demand L. The firm's optimality conditions then pin down the pre-tax interest rate r and the pre-tax wage w. Thus, after dropping all redundant choice variables and the associated constraints, the problem reduces to

$$\max_{\{\bar{r},\bar{w},p,K\}} \mathcal{W}(\bar{r},\bar{w},p) \quad \text{s.t} \quad \begin{cases} \mathcal{C}(\bar{r},\bar{w},p) + G = F(K,\mathcal{L}(\bar{r},\bar{w},p)) - \delta K \end{cases} .$$

Letting λ^{GM} denote the Lagrange multiplier on the goods market clearing condition and defining the capital labor ratio $k \equiv \frac{K}{\mathcal{L}(\cdot)}$, the first order condition with respect to K can be written as

$$\lambda^{GM}\left[F_K(k,1) - \delta\right] = 0 \implies F_K(k,1) = \delta.$$

This establishes the optimality of the golden rule in the OSS problem for the model with capital.

E.2 ADDITIONAL FIRST ORDER CONDITIONS

The first order conditions with respect to $\bar{r}$, $\bar{w}$, and p are given by

$$\frac{\partial \mathcal{W}}{\partial \bar{r}} + \lambda^{GM}\left\{w^{GR}\frac{\partial \mathcal{L}}{\partial \bar{r}} - \frac{\partial \mathcal{C}}{\partial \bar{r}}\right\} = 0, \tag{E.2}$$

$$\frac{\partial \mathcal{W}}{\partial \bar{w}} + \lambda^{GM}\left\{w^{GR}\frac{\partial \mathcal{L}}{\partial \bar{w}} - \frac{\partial \mathcal{C}}{\partial \bar{w}}\right\} = 0, \tag{E.3}$$

$$\frac{\partial \mathcal{W}}{\partial p} + \lambda^{GM}\left\{w^{GR}\frac{\partial \mathcal{L}}{\partial p} - \frac{\partial \mathcal{C}}{\partial p}\right\} = 0, \tag{E.4}$$

where $w^{GR} = F_L(k^{GR}, 1)$ is the pre-tax wage implied by the golden rule. Using the second one to eliminate λ^{GM}, these become

$$\frac{\partial \mathcal{W}}{\partial \bar{r}} = \frac{\frac{\partial \mathcal{W}}{\partial \bar{w}}}{\frac{\partial \mathcal{C}}{\partial \bar{w}} - w^{GR}\frac{\partial \mathcal{L}}{\partial \bar{w}}} \left\{ \frac{\partial \mathcal{C}}{\partial \bar{r}} - w^{GR}\frac{\partial \mathcal{L}}{\partial \bar{r}} \right\}, \tag{E.5}$$

$$\frac{\partial \mathcal{W}}{\partial p} = \frac{\frac{\partial \mathcal{W}}{\partial \bar{w}}}{\frac{\partial \mathcal{C}}{\partial \bar{w}} - w^{GR}\frac{\partial \mathcal{L}}{\partial \bar{w}}} \left\{ \frac{\partial \mathcal{C}}{\partial p} - w^{GR}\frac{\partial \mathcal{L}}{\partial p} \right\}. \tag{E.6}$$

I use (E.5) and (E.6), together with the goods market clearing condition in steady state, to solve for the optimal steady state in the model with capital. To reduce the dimensionality of the problem, I use the fact that, when the elasticity of intertemporal substitution equals unity, aggregate steady-state consumption scales with the after-tax wage and aggregate labor supply is independent of it.

F RAMSEY PROBLEM IN THE MODEL WITH CAPITAL AND τ_k

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

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

Once again, $\bar{r}_t \equiv (1-\tau_{kt})r_t$ and $\bar{\varpi}_t \equiv \tau_t w_t^{1-p_t}$ denote the after-tax interest rate and the after-tax wage. Notice that we drop the government budget constraint because of Walras' Law. Here, $\mathcal{U}_t$, $\mathcal{A}_t$, $\mathcal{C}_t$, 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.

F.1 PROVING THE OPTIMALITY OF THE MODIFIED GOLDEN RULE

Given choices of $\{\bar{r}_t,\bar{\varpi}_t,p_t,K_t\}$, 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

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

The Lagrangian for this problem is

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

The first-order condition with respect to capital is

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

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[26]

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

This establishes the optimality of the modified golden rule.[27]

[26] As discussed by Straub and Werning (2020), two situations prevent applicability of this result: (i) nonconvergence to an interior steady state; or (ii) nonconvergence of the multipliers.

[27] This result was first proved by Aiyagari (1995) under the assumption that government spending is endogenous.

F.2 Additional first-order conditions

In addition to the first-order condition with respect to capital (F.1), we also have:

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial \bar{r}_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial \bar{r}_u} - \frac{\partial \mathcal{C}_t}{\partial \bar{r}_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{F.2}$$

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial \bar{\varpi}_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial \bar{\varpi}_u} - \frac{\partial \mathcal{C}_t}{\partial \bar{\varpi}_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{F.3}$$

$$\sum_{t=0}^{\infty} \beta^t \frac{\partial \mathcal{U}_t}{\partial p_u} + \sum_{t=0}^{\infty} \beta^t \lambda_t^{GM} \left(w_t \frac{\partial \mathcal{L}_t}{\partial p_u} - \frac{\partial \mathcal{C}_t}{\partial p_u} \right) = 0, \quad \text{for } u = 0, 1, \ldots \tag{F.4}$$

If quantities and multipliers converge to an interior steady state, in the RSS these conditions become

$$\lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial \bar{r}_u} + \lambda^{GM} \left(w^{MGR} \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{L}_t}{\partial \bar{r}_u} - \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{C}_t}{\partial \bar{r}_u} \right) = 0, \tag{F.5}$$

$$\lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial \bar{\varpi}_u} + \lambda^{GM} \left(w^{MGR} \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{L}_t}{\partial \bar{\varpi}_u} - \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{C}_t}{\partial \bar{\varpi}_u} \right) = 0, \tag{F.6}$$

$$\lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{U}_t}{\partial p_u} + \lambda^{GM} \left(w^{MGR} \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{L}_t}{\partial p_u} - \lim_{u \to \infty} \sum_{t=0}^{\infty} \beta^{t-u} \frac{\partial \mathcal{C}_t}{\partial p_u} \right) = 0, \tag{F.7}$$

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 $\{\bar{r}, \bar{\varpi}, p, \lambda^{GM}\}$. Appendix F.3 details the computational procedure.

F.3 Computing the RSS in the model with capital

To operationalize (F.5)-(F.7) 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:

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

In principle, this could be done by computing the Jacobians for welfare, labour supply, and consumption, around the steady state implied by some candidate $(\bar{r}, \bar{\varpi}, p)$ 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 too costly since we need a pretty large horizon to get convergence of the discounted sum. Therefore, I proceed as explained in Appendix B.2.