Optimal Long-Run Fiscal Policy with Heterogeneous Agents[*]

Adrien Auclert[†] Michael Cai[‡] Matthew Rognlie[§] Ludwig Straub[¶]

March 2024

Preliminary

Abstract

We introduce a new approach to characterizing the steady state of dynamic Ramsey taxation problems, and apply this approach to standard calibrations of heterogeneous-agent models a la Aiyagari (1995). In many cases, we find that such Ramsey steady states do not exist, with our results instead pointing to the optimality of long-run immiseration. When Ramsey steady states do exist, they are associated with optimal long-run labor income taxes close to 100%. We show that these conclusions are related to unreasonably strong anticipatory effects of future tax changes.

[*]We thank Anmol Bhandari, Mike Golosov, Marcus Hagedorn, and Narayana Kocherlakota for helpful comments and suggestions. This research is supported by the National Science Foundation grant numbers SES-1851717 and SES-2042691, as well as the Harvard Chae Initiative.

[†]Stanford University and NBER. Email: aauclert@stanford.edu.

[‡]Northwestern University. Email: michaelcai@u.northwestern.edu.

[§]Northwestern University and NBER. Email: matthew.rognlie@northwestern.edu.

[¶]Harvard University and NBER. Email: ludwigstraub@fas.harvard.edu.

1 Introduction

Models of household behavior with uninsurable idiosyncratic income risk a la Bewley-Aiyagari-Huggett—heterogeneous-agent models for short—have become nothing short of a success story over the last thirty years. At the micro level, they are able to generate reasonable buffer-stock saving behavior (Zeldes 1989, Deaton 1992, Carroll 1997), reasonable consumption and income profiles over the life-cycle (Gourinchas and Parker 2002, Storesletten, Telmer and Yaron 2004), reasonable contemporary and intertemporal MPCs (Kaplan and Violante 2014, Auclert, Rognlie and Straub 2023), and reasonable labor supply choices (Blundell and MaCurdy 1999). At the macro level, heterogeneous-agent models are at the heart of modern analyses of growth and inequality (Castañeda, Díaz-Giménez and Ríos-Rull 2003) and business cycles (Krusell and Smith 1998, Kaplan, Moll and Violante 2018, Auclert, Rognlie and Straub 2020).

Relative to the prevalence of heterogeneous-agent models in macroeconomics today, little is known about their normative properties: How should a planner trade off capital and labor taxation? How much debt should a planner use to finance the government, taking into account the self-insurance benefits to households? How progressive should taxes be if the usual equity-efficiency trade-off is mitigated by precautionary saving? Does increased inequality mean higher debt levels are optimal? One reason why answers to these questions have remained elusive is the computational complexity of heterogeneous-agent models, especially when embedded in an optimal policy problem.

In this paper, we propose a new method to analyze optimal long-run fiscal policy in heterogeneous-agent economies. Our method allows us to characterize the *Ramsey steady state (RSS)* of heterogeneous-agent models—the long-run steady state of the optimal full-commitment Ramsey plan. Ramsey steady states have been a central object of interest in optimal taxation for decades, going back to the early work of Chamley (1986) and Judd (1985), which characterized Ramsey steady states in representative-agent and two-agent models, respectively.[1] Our method is made possible by recent advances in expressing macroeconomic models in the "sequence space" (Auclert, Bardóczy, Rognlie and Straub 2021). The method yields an intuitive RSS optimality condition that involves highly interpretable and estimable *discounted elasticities*, which generalize related objects in Piketty and Saez (2013) and Straub and Werning (2020).

We evaluate our RSS optimality condition numerically for standard parameterizations of heterogeneous-agent models. Our main finding is that Ramsey steady states in this class of models are very extreme. In many parameterizations with balanced growth preferences, a Ramsey steady state doesn't exist at all, with our results pointing to immiseration—labor taxes approaching 100% and real consumption falling to zero. In others, a Ramsey steady state does exist, but it involves near-immiseration tax rates above 90%. The only versions of the model that we find can lead to reasonable Ramsey steady states far from immiseration are versions with non-balanced growth preferences and no wealth effects on labor supply (Greenwood, Hercowitz and Huffman 1988,

[1] See Straub and Werning (2020) for a recent qualification of this work.

henceforth GHH).

Model and RSS condition. We begin our exploration with a baseline model à la Aiyagari (1995), with two modifications. We allow for general household preferences and assume production is linear in labor, abstracting away from capital at first to ease exposition. We assume the social planner can fully commit to paths of proportional labor income taxes and public debt at date 0. We allow for a general social discount factor, though focus on the case where social and private discount factors coincide.

We show that in this economy aggregate household behavior can be summarized by *sequence-space functions*, that is, we can write aggregate household asset demand A_t at date t as a function of the entire sequences of interest rates and wages (after taxes) $A_t = \mathcal{A}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right)$; likewise, we can write aggregate labor supply N_t at date t as $N_t = \mathcal{N}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right)$. Similar sequence-space functions have proved very useful in the previous literature, for decompositions (e.g. Kaplan et al. 2018, Farhi and Werning 2019) as well as structuring and solving heterogeneous-agent models (Auclert et al. 2020, 2021, 2023). We use these sequence space functions to define *discounted elasticities*, such as $\epsilon^{N,w}$, defined as the response of the present value of labor supply to a fully anticipated one-time increase in wages, and similarly for other elasticities of sequence-space functions. These elasticities are potentially estimable in micro data and turn out to be useful "sufficient statistics" for household behavior.

Our main theoretical result is a pair of necessary optimality conditions that have to hold at the Ramsey steady state of our economy. The first optimality condition is assuming a converging multiplier λ_t on our implementability condition. It is best thought of as equating benefits and costs from the planner choosing to provide more liquidity (raising r, benefiting households) financed by higher labor taxes (lowering w, hurting households). The optimality condition then equates the *liquidity benefit* with disincentive effects of lower wages on *labor supply* and costs from *redistribution* of resources from the average worker to the average saver. Our second optimality condition also involves these three terms, but explicitly allows for exploding multipliers $\lambda_t \to \infty$, in the spirit of Straub and Werning (2020).

Immiseration and near-immiseration. We put our optimality conditions into action and evaluate them numerically for common heterogeneous-agent models in section 4. We begin with standard log-separable preferences and a conventional calibration to the U.S. economy. We check the optimality condition for all possible steady states of the model, finding that the benefit from marginally increasing interest rates and increasing labor income taxes is always strictly positive. Thus, a Ramsey steady state does not exist. Further, our results suggest that the economy instead tends towards immiseration, with labor taxes approaching 100% and real consumption converging to zero. We identify the discounted labor supply elasticity $\epsilon^{N,w}$ as the main force towards immiseration in the model: We find that $\epsilon^{N,w}$ is *negative*, implying that greater labor taxation actually *increases* labor supply. As agents rationally anticipate greater future labor taxes, they raise their hours in the present. This turns the traditional labor supply cost of high labor taxation into a benefit.

Our finding of immiseration raises several questions. Chief among them is how sensitive this finding is to the assumed preferences and the calibration of the model (section 5). We conduct an extensive robustness analysis and find that, as long as preferences are consistent with balanced growth as in King, Plosser and Rebelo (1988), there is either a Ramsey steady state very close to immiseration (labor taxes above 90%) or no Ramsey steady state at all, with results pointing to immiseration. This is true for a variety of different income processes, for different assumptions on initial government debt or government spending, and for models with lump-sum transfers and progressive taxes.

In section 6, we then turn to an economy with capital and capital income taxes (as in Aiyagari 1995). We derive a pair of two very similar RSS optimality conditions. We show that, again, no Ramsey steady state exists. Instead, our results point to immiseration again. On top of that, we find results suggesting that Lagrange multipliers diverge, breaking the modified golden rule of capital accumulation.

Non-balanced growth preferences and alternative household sides. A common alternative to balanced-growth preferences are additively separable and GHH preferences, both of which we consider in section 5. For additively separable preferences, we can always find Ramsey steady states when the EIS is greater than 1, with labor income taxes strictly below 100%. This is because an EIS greater than 1 dampens the negative wealth effect that higher labor taxes have on labor supply, pushing $\epsilon^{N,w}$ into positive territory. Vice versa, when the EIS is less than 1, labor taxes approach 100% again. GHH preferences also give rise to Ramsey steady states as wealth effects on labor supply are entirely absent with these preferences.

We end our paper by discussing several alternative household sides in section 8. We find that bond-in-utility models—often regarded as "tractable" version of heterogeneous-agent models—have conceptually similar predictions for Ramsey steady states as our full-blown Aiyagari model. Indeed, we also find immiseration in a standard log-separable version of a bond-in-utility model. We show how our optimality conditions are also useful in understanding why standard overlapping generations models and models with alternating income states (Woodford 1990) generally do not lead to immiseration.

Taking stock. There are two broad reactions a reader might have to our findings of the optimality of (near-)immiseration in a large class of otherwise very reasonable heterogeneous-agent models. The first reaction is that immiseration is an artifact of the full-commitment setup we are using. Indeed, with limited commitment or a social discount factor that exceeds the private one, we may no longer find (near-)immiseration to be the solution to the planning problem. However, this will not change the fact that households themselves would still *want* immiseration in the long run, only that the planner cannot deliver it to them.

The second reaction is that a negative elasticity $\epsilon^{N,w}$ clearly has the wrong sign relative to what it should be empirically. Indeed, if one were to modify household behavior to dampen the anticipatory effects of wage changes (similar to dampened anticipation effects in García-Schmidt and Woodford 2019 and Gabaix 2020), we suspect that $\epsilon^{N,w}$ would flip its sign to be positive, making immiseration no longer optimal. We view this as a fruitful avenue for future work.

Literature. Our paper is most closely related to the literature studying Ramsey taxation in heterogeneous agent models, specifically Aiyagari (1995), Acikgoz, Hagedorn, Holter and Wang (2018), Chien and Wen (2022), and Dyrda and Pedroni (2023). Relative to these papers, our paper develops a new method to characterize Ramsey steady states, and demonstrates how in many standard cases (though not all) immiseration is optimal. Since it is easiest to relate to these papers in detail after presenting our results, we describe them separately in section 9.

Our paper is also related to the classic paper by Aiyagari and McGrattan (1998), which assumes an infinitely patient planner, that is, a planner that only cares about the long-run steady state, ignoring the transition it takes to get there. We focus instead on the case where the social and private discount factors coincide. Boar and Midrigan (2022) solve a planning problem that does take transitions into account, but restrict attention to constant tax rates. LeGrand and Ragot (2023) derive analytical results in a Woodford (1990) alternating-income-states economy; when they simulate impulse responses in a full heterogeneous-agent model, they do so assuming that the planner has income-state-dependent welfare weights, effectively endowing the planner with different preferences from those of the agents. Finally, Dávila, Hong, Krusell and Ríos-Rull (2012) consider the optimality of individual saving behavior in a heterogeneous-agent economy. Their constrained efficient allocation corresponds to that chosen by a planner that is able to levy individual-specific capital taxes or subsidies, whose revenue is then rebated lump-sum back to the individual agents. We focus, instead, on capital and labor tax schedules that are the same for all agents, and allow the government to borrow and spend.

2 Heterogeneous agent economy

Our economy is a standard Aiyagari (1994, 1995) economy, with flexible labor supply and a general utility function. The only simplification we make in the first part of the paper is that we study an economy without capital. However, as we show in great detail in section 6, all our conclusions carry over to an economy with capital and capital taxes. Time is discrete, $t = 0, 1, 2, \ldots$, and there is no aggregate risk. All agents have perfect foresight.

2.1 Households

There is a continuum of households, labeled by $i \in [0, 1]$. Each household i draws an idiosyncratic productivity process e_{it} that follows a positive and recurrent Markov chain; e_{it} is iid across households with a mean normalized to 1. Household i's date-0 expected utility is given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_{it}, n_{it}) \right] \tag{1}$$

Here, $\beta \in (0,1)$ is the private discount factor, c_{it} denotes consumption and n_{it} labor supply. $u(c,n)$ is a general per-period utility function satisfying the usual Inada conditions.[2] Households are able to self-insure against their idiosyncratic productivity shocks by holding assets a_{it}, earning an interest rate r_t between periods $t-1$ and t. Their budget constraint is given by

$$c_{it} + a_{it} = (1 + r_t)\, a_{it-1} + w_t e_{it} n_{it} \tag{2}$$

where $w_t > 0$ is the date-t wage per effective unit of labor. Households are also subject to a standard borrowing constraint

$$a_{it} \geq 0 \tag{3}$$

Taken together, households choose $\{c_{it}, n_{it}, a_{it}\}$ in order to maximize (1) subject to (2) and (3).

As shown in Auclert et al. (2023), aggregate (partial equilibrium) household behavior in this economy can be expressed entirely as a function of the sequences of interest rates and wages, $\{r_t, w_t\}_{t=0}^{\infty}$—the two prices in the economy. Conditional on $\{r_t, w_t\}_{t=0}^{\infty}$, households can solve for their consumption policies $c_t^*(e, a_-)$, labor supply policies $n_t^*(e, a_-)$, and savings policies $a_t^*(e, a_-)$. Those, then, perfectly determine the evolution of the wealth distribution $\Psi_t(e, a_-)$, starting from some arbitrary given initial distribution $\Psi_0(e, a_-)$. Date-t aggregate (partial equilibrium) household asset demand can then be written as

$$A_t = \mathcal{A}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \equiv \int a_t^*(e, a_-)\, d\Psi_t(e, a_-) \tag{4}$$

Aggregate (partial equilibrium) effective labor supply can be written as

$$N_t = \mathcal{N}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \equiv \int n_t^*(e, a_-) d\Psi_t(e, a_-) \tag{5}$$

Aggregate consumption can be written as

$$C_t = \mathcal{C}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \equiv \int c_t^*(e, a_-) d\Psi_t(e, a_-) \tag{6}$$

Finally, utilitarian flow utility can be written as

$$U_t = \mathcal{U}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \equiv \int u(c_t^*(e, a_-), n_t^*(e, a_-)) d\Psi_t(e, a_-) \tag{7}$$

[2]In particular: $u : (0, \infty) \times [0, \infty) \to \mathbb{R}$ is twice differentiable and strictly concave, with $\lim_{c \to 0} u_c(c,n) = \infty$ and $\lim_{c \to \infty} u_c(c,n) = 0$ for any $n \geq 0$; $\lim_{n \to 0} u_n(c,n) = 0$ and $\lim_{n \to \infty} u_n(c,n) = \infty$ for any $c > 0$.

Each one of these four "curly" functions, $\mathcal{A}, \mathcal{N}, \mathcal{C}$, and $\mathcal{U}$, is a *sequence-space function*, in that it maps the two sequences of prices into sequences of aggregate household behavior and household utility.

In order to prove our main results below, we have to impose regularity assumptions on the four curly functions $\mathcal{A}, \mathcal{N}, \mathcal{C}, \mathcal{U}$ and their derivatives. To state the regularity assumptions, we first define the concept of a β-quasi-Toeplitz matrix.

Definition 1. An infinite matrix $\mathbf{M} = [M_{t,s}] \in \mathbb{R}^{\mathbb{N}\times\mathbb{N}}$ is a *β-quasi-Toeplitz matrix with symbol vector* $\mathbf{a} = \{a_t\} \in \mathbb{R}^{\mathbb{Z}}$ if:

1. for any $t, s \geq 0$ we have
$$\lim_{u\to\infty} M_{t+u,s+u} = a_{t-s} \tag{8}$$

2. the tails of $\mathbf{a}$ decay to zero exponentially fast, and faster than β for negative indices, that is, for some $0 < \tilde{\beta} < \beta$, $0 < \tilde{\gamma} < 1$ and constants $C_1, C_2 > 0$, we have
$$|a_t| \leq \begin{cases} C_1 \tilde{\gamma}^t & t \geq 0 \\ C_2 \tilde{\beta}^{-t} & t < 0 \end{cases} \tag{9}$$

3. the distance between $M_{t,s}$ and a_{t-s} decays exponentially, too, that is, for some constants $C_3, C_4 > 0$
$$|M_{t,s} - a_{t-s}| \leq \begin{cases} C_3 \tilde{\gamma}^{t-s} & t-s \geq 0 \\ C_4 \tilde{\beta}^{-(t-s)} & t-s < 0 \end{cases} \tag{10}$$

A Toeplitz matrix is a matrix whose elements are constant along all sub-diagonals, that is, $M_{t,s}$ only depends on $t-s$. A quasi-Toeplitz matrix is one which has this property approximately, for large t, s, as in (8). A *β-quasi-Toeplitz matrix* is one in which entries $M_{t,s}$ decay at least at speed β as we make s much larger than t, as in (9) and (10). The reason β-quasi-Toeplitz matrices are useful for our analysis is that, as it turns out, they embody the structure of derivatives of sequence-space functions of stationary economic models, such as $\mathcal{A}, \mathcal{N}, \mathcal{C}, \mathcal{U}$ for our heterogeneous-agent household side.

For the results of this paper to be valid, we then impose the following regularity assumption.

Assumption 1. *Each one of the curly functions $\mathcal{X} \in \{\mathcal{A}, \mathcal{N}, \mathcal{C}, \mathcal{U}\}$ satisfies the following four assumptions:*

(i) *$\mathcal{X}$ is a well-defined continuous function with respect to the sup-norm.*

(ii) *Evaluated at convergent sequences $\{r_s, w_s\}_{s=0}^{\infty}$ with $r_s \to r^* < 1/\beta - 1$ and $w_s \to w^* > 0$, $\mathcal{X}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right)$ converges to a limit $\mathcal{X}^{ss}(r^*, w^*)$ that only depends on the limits r^*, w^*.*

(iii) *Evaluated at convergent sequences $\{r_s, w_s\}_{s=0}^{\infty}$ with $r_s \to r^* < 1/\beta - 1$ and $w_s \to w^* > 0$, $\mathcal{X}$ is Fréchet-differentiable in each of its two arguments, with bounded linear derivatives $\mathbf{X}^{(r)}\left(\{r_s, w_s\}_{s=0}^{\infty}\right) : \ell^\infty \to \ell^\infty$ and $\mathbf{X}^{(w)}\left(\{r_s, w_s\}_{s=0}^{\infty}\right) : \ell^\infty \to \ell^\infty$.*

(iv) Each of the derivatives is representable by a β-quasi-Toeplitz matrix whose symbol vector only depends on the limits r^*, w^*.

Assumption 1 makes explicit the four properties of the sequence-space functions $\mathcal{A}, \mathcal{N}, \mathcal{C}, \mathcal{U}$ that we need for our results. Property (i) simply requires continuity, property (ii) stationarity, and property (iii) requires differentiability. Property (iv) requires that the derivatives of the four sequence-space functions are β-quasi-Toeplitz. Properties (ii) and (iv) rule out all non-stationary household sides, such as ones with a representative agent or a spender-saver structure.

In appendix A.1, we show that assumption 1, including properties (ii) and (iv), is satisfied for a discretized version of the heterogeneous-agent model introduced in this section.[3] We believe that the assumption is satisfied for any standard, stationary model of household behavior (in section 8 we consider bonds-in-utility and overlapping generations models as tractable examples). In this sense, we regard assumption 1 as a regularity assumption which does not meaningfully restrict the set of models beyond requiring stationarity. We impose the assumption throughout this paper.

The limit $\mathcal{X}^{ss}$ introduced in property (ii) is exactly characterizing steady state behavior of variable $\mathcal{X}$ as function of a steady state interest rate and wage. We write $\mathcal{A}^{ss}(r, w)$ for steady state asset holdings as function of r and w; $\mathcal{N}^{ss}(r, w)$ for steady state labor supply; $\mathcal{C}^{ss}(r, w)$ for steady state consumption; and $\mathcal{U}^{ss}(r, w)$ for steady state flow utility.

We consider several different household utility functions in this paper. The main ones we focus on are balanced growth compatible à la King et al. (1988), henceforth KPR,

$$u(c, n) = \frac{\left(ce^{-v(n)}\right)^{1-\sigma} - 1}{1 - \sigma} \tag{11}$$

for some $\sigma > 0$, such as the standard log-separable utility function,

$$u(c, n) = \log c - v(n) \tag{12}$$

with $v(n) \geq 0$ denoting the disutility from labor supply.

2.2 Production and government policy

Production is perfectly competitive and linear in effective labor. The unique output good at date t is produced using the technology

$$Y_t = N_t \tag{13}$$

This implies that the pre-tax wage w^* is identical to 1, $w^* = 1$.

The government is financing an exogenous amount of government spending $G > 0$ using a mix

[3] We believe our arguments in appendix A.1 also carry over to a non-discretized version of our model.

of proportional labor income taxes $\{\tau_t\}$ and government debt $\{B_t\}$. After-tax wages are simply

$$w_t = 1 - \tau_t \tag{14}$$

and tax revenue from labor income taxation is given by $\tau_t N_t$. The government budget constraint is

$$G + (1 + r_t) B_{t-1} = B_t + \tau_t N_t \tag{15}$$

We discuss how our results carry over to environments with endogenous government spending in appendix C.1.

2.3 Competitive equilibrium and implementability

We define equilibrium as follows.

Definition 2. A *competitive equilibrium* in our economy is a collection of quantities $\{Y_t, N_t, B_t, C_t, A_t\}_{t=0}^{\infty}$, tax rates $\{\tau_t\}_{t=0}^{\infty}$, and prices $\{r_t, w_t\}_{t=0}^{\infty}$, such that:

1. The after-tax wage is given by (14).
2. Households optimize given prices: consumption C_t is given by (6), assets A_t are given by (4), effective labor supply N_t is given by (5).
3. Output is given by (13).
4. The government budget constraint (15) holds.
5. The asset market clears, $A_t = B_t$. The goods market clears, $C_t + G = Y_t$.

A competitive equilibrium is a *steady state equilibrium* if all quantities, tax rates, and prices are constant.

Our approach will, essentially, follow the dual approach to Ramsey taxation. This lets us derive simple implementability conditions, stated directly in terms of the sequence-space functions defined above.

Proposition 1 (Implementability). *$\{r_t, w_t\}_{t=0}^{\infty}$ are part of a competitive equilibrium if and only if*

$$\mathcal{C}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) + G = \mathcal{N}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \tag{16}$$

which holds if and only if

$$G + (1 + r_t)\,\mathcal{A}_{t-1}\left(\{r_s, w_s\}_{s=0}^{\infty}\right) = \mathcal{A}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) + (1 - w_t)\,\mathcal{N}_t\left(\{r_s, w_s\}_{s=0}^{\infty}\right) \tag{17}$$

All proofs are relegated to the appendix. This result may at first seem surprising. How is the goods market clearing condition (16) on its own sufficient for equilibrium? The reason for this is that optimal household behavior is already substituted into the equation. This means that the sequences of household consumption $\{C_t\}$ and labor supply $\{N_t\}$ that clear the goods market also come from an optimizing household side. Setting government debt B_t equal to household optimal asset demand $\mathcal{A}_t(\{r_s, w_s\})$ then naturally clears the asset market. The only remaining equilibrium condition is the government budget constraint, but that one must hold by Walras' law. The logic behind the sufficiency of the government budget constraint (17) in proposition 1 is similar.

2.4 Infinitely anticipated shocks and discounted elasticities

An important partial equilibrium thought experiment in our analysis is that of an "infinitely anticipated shock". Consider the households in our economy entering period $t = 0$ in a steady state, consistent with constant interest rate r and after-tax wage w. Then, imagine that, at some far-away date $s \gg 0$, the interest rate is announced to increase by some small amount, for one period. Everything else remains the same. How does the time path of aggregate household asset demand A_t respond?

Using the sequence-space function for asset demand (4), we see that the date $s + h$ response, in percentage points of steady state assets, is exactly given by the derivatives

$$\frac{\partial \log \mathcal{A}_{s+h}}{\partial r_s} \tag{18}$$

where h runs from $-s$ to ∞. Interestingly, as explained in Auclert et al. (2021, 2023), for s sufficiently large, the derivatives start to become independent of s, and only depend on the time horizon h relative to the date of the anticipated shock. For our baseline calibration below (see section 4.1), figure 1(a) plots the derivatives $\partial \log \mathcal{A}_{s+h}/\partial r_s$. We can see that households basically do not respond more than 30 years prior to the shock, then build up assets to earn the additional interest rate and subsequently decumulate assets back down to the steady state.

To give a second example, figure 1(b) plots the derivatives of labor supply to an increase in labor income taxes (equivalent to a decrease in after-tax wages)

$$-\frac{\partial \log \mathcal{N}_{s+h}}{\partial \log w_s} \tag{19}$$

with a large date s. We see, once more, that households barely respond 30 years before the shock. Then, they slowly increase labor supply to offset the future income loss, sharply reduce their labor supply in the period of the tax increase due to a traditional substitution effect, and then return to working harder and offsetting the income loss. As we will see below, the anticipatory increase in labor supply to higher future labor income taxes will play a prominent role in this paper.

One way to summarize responses to infinitely anticipated shocks, such as the ones in (18) and (19), is by discounting them relative to the period of the shock. For a general discount factor

Figure 1: Household responses to infinitely anticipated shocks

(a) Asset demand response to interest rate increase

(b) Labor supply response to labor tax increase

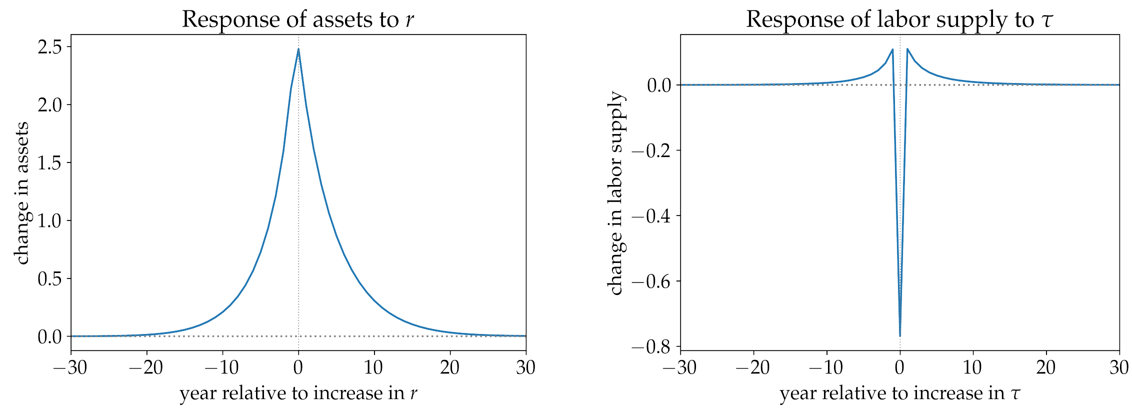

$\delta \in [\beta, 1]$, we define

$$\epsilon^{A,r}(\delta) \equiv \lim_{s\to\infty} \sum_{h=-s}^{\infty} \delta^h \frac{\partial \log \mathcal{A}_{s+h}}{\partial r_s} \qquad \epsilon^{N,\tau}(\delta) \equiv -\epsilon^{N,w}(\delta) \equiv -\lim_{s\to\infty} \sum_{h=-s}^{\infty} \delta^h \frac{\partial \log \mathcal{N}_{s+h}}{\partial \log w_s} \tag{20}$$

Analogously, we define derivatives of all the other sequence space functions as well, such as $\epsilon^{A,w}(\delta), \epsilon^{U,r}(\delta), \epsilon^{C,w}(\delta)$, etc. Whenever we write a derivative with respect to tax rates τ, we interpret that as the negative of the elasticity with respect to after-tax wages, as in (20). Our discounted derivatives have intellectual forebears in the work of Piketty and Saez (2013) and Straub and Werning (2020). We show in appendix A.2 that the discounted elasticities defined here are indeed well-defined in our model, and describe an efficient algorithm to compute them in appendix B.

The discount factor δ is left general for the moment. It will later take the role of the social discount factor, so a natural baseline will be to assume it equals the private discount factor, $\delta = \beta$. For the numerical examples shown in figure 1, we find $\epsilon^{A,r}(\beta) = 36$ and $\epsilon^{N,\tau}(\beta) = 0.10$ when using the private discount factor. Notably, the discounted labor supply elasticity becomes positive here, simply due to the anticipatory increase in labor supply before the tax increase.

While we focus on a specific description of household behavior in this section—one based on heterogeneous agents and uninsurable idiosyncratic income risk—similar discounted elasticities can be defined and computed for any other stationary model of household behavior, including bonds-in-utility and life-cycle models. We explore alternative household sides in section 8.

3 Sequence-Space Approach to the Ramsey Steady State

We are now ready to set up the full-commitment Ramsey problem and derive our main formal result: a new necessary condition characterizing the long-run steady state of the Ramsey plan, the

Ramsey steady state (RSS). Throughout this section, we work with a general social discount factor $\delta \in [\beta, 1)$. In the next section, we will also consider the limit $\delta \to 1$, which corresponds to the solution concept of the "optimal steady state"—a planner entirely focused on maximizing steady state welfare.

3.1 Ramsey problem

The Ramsey planner maximizes utilitarian welfare each period, discounted at δ,

$$\sum_{t=0}^{\infty} \delta^t \mathcal{U}_t \left(\{r_s, w_s\}\right) \tag{21}$$

subject to either of the two implementability conditions in proposition 1. We work with the government budget constraint in the main body of the paper, implying that the constraint we use is (17). The two sequences the planner maximizes over are the sequences of interest rates and after-tax wages, $\{r_s, w_s\}_{s=0}^{\infty}$. Intuitively, we can think of $w_s = 1 - \tau_s$ as being controlled directly by the planner, and r_s as being endogenously determined by the government budget constraint.

While the objective (21) is utilitarian, this is not crucial for our results. The reason is that our model does not have any permanent types. Any non-utilitarian welfare function with some weight κ_i on utility of individual i will, in the long run, be proportional to a utilitarian one, i.e. $\int_i \kappa_i u(c_{it}, n_{it}) di$ will end up being proportional to $\mathcal{U}_t$, simply because, absent permanent types, individuals are mixing and will eventually all look alike. This logic doesn't apply with permanent types, as we briefly explore in section 8.1.

3.2 Ramsey steady state

We define a Ramsey steady state as follows.

Definition 3. A steady-state equilibrium consisting of quantities Y, N, B, C, A, a tax rate τ, and prices r, w is called a *Ramsey steady state* of the economy if there exists a solution $\{r_s, w_s\}_{s=0}^{\infty}$ of the Ramsey problem (a *Ramsey plan*) such that $\mathcal{C}_t \left(\{r_s, w_s\}_{s=0}^{\infty}\right) \to C$, $\mathcal{A}_t \left(\{r_s, w_s\}_{s=0}^{\infty}\right) \to A$, $\mathcal{N}_t \left(\{r_s, w_s\}_{s=0}^{\infty}\right) \to N$, $w_t \to w$, and $r_t \to r$.

This means that a Ramsey steady state always has an associated Ramsey plan which converges to it. The following lemma shows that really all we need is that r_t and w_t converge.

Lemma 1. *Under assumption 1, a steady-state equilibrium with prices r, w is a Ramsey steady state if and only if there is a Ramsey plan $\{r_s, w_s\}_{s=0}^{\infty}$ with $r_t \to r$ and $w_t \to w$.*

This lemma is helpful as it states that we only need convergence of the two price sequences; convergence of any other equilibrium objects follows.

Before we state our main result, we define two additional objects. The first is the a unit-less measure of liquidity in any steady state of the economy, defined as the ratio of available assets for self insurance to aggregate after-tax income

$$\ell \equiv \frac{A}{wN}$$

Comparing liquid assets to income is a typical way to assess liquidity in the literature (e.g. Kaplan, Violante and Weidner 2014). Here, it makes sense theoretically, as the amount of income risk households are subject to scales with aggregate after-tax income wN. Thus, ℓ is essentially a comparison of the available liquid assets divided by the amount of income risk.

The second object we define is the effective marginal rate of substitution between a one time interest rate increase and a one time wage increase around some steady state,

$$\text{mrs}(\delta) \equiv \frac{\epsilon^{U,w}(\delta)}{\epsilon^{U,r}(\delta)}$$

The numerator as well as the denominator are discounted elasticities, for example, $\epsilon^{U,w}(\delta) = \lim_{s\to\infty} \sum_{h=-s}^{\infty} \delta^h \frac{\partial \mathcal{U}_{s+h}}{\partial \log w_s}$; $\epsilon^{\mathcal{U},r}(\delta)$ is defined similarly. When $\delta = \beta$, we show in appendix A.5 that $\text{mrs}(\beta)$ can be rewritten as

$$\text{mrs}(\beta) = \frac{1}{\ell} \cdot \frac{\int u_c(c_{it}, n_{it}) \frac{e_{it} n_{it}}{N} di}{\int u_c(c_{it}, n_{it}) \frac{a_{it-1}}{A} di} \tag{22}$$

The formula follows directly from the envelope theorem. The numerator is the marginal utility of the average labor income recipient; the denominator is the marginal utility of the average asset income recipient. Since the labor income recipients are typically poorer in the models we study, we typically expect $\ell \cdot \text{mrs}(\beta)$ to be greater than one.

3.3 Ramsey steady state condition

A useful way to start thinking about the Ramsey steady state is to derive first-order conditions of the Ramsey problem. For example, denoting by λ_t the date-t current-value Lagrange multiplier of the government budget constraint, we can write the first order condition with respect to r_s, for some period $s \geq 0$, as

$$\sum_{h=-s}^{\infty} \delta^h \frac{\partial \mathcal{U}_{s+h}}{\partial r_s} + \sum_{h=-s}^{\infty} \delta^h \lambda_{s+h} \left(\frac{\partial \mathcal{A}_{s+h}}{\partial r_s} + (1 - w_{s+h}) \frac{\partial \mathcal{N}_{s+h}}{\partial r_s} - (1 + r_{s+h}) \frac{\partial \mathcal{A}_{s+h-1}}{\partial r_s} \right) - \lambda_s \mathcal{A}_{s-1} = 0 \tag{23}$$

Many terms in this expressions look similar to the discounted elasticities we introduced in section 2.4. For instance, as one takes the limit $s \to \infty$, the first term immediately converges to $\epsilon^{U,r}(\delta)$. If the multiplier λ_t were to be constant, or converging to a constant, the three terms in parentheses would also all converge to discounted elasticities as $s \to \infty$. A similar first order condition can be written down for the wage w_s at some period s.

We next present a series of necessary conditions for a pair of prices (r, w) to be part of a Ramsey steady state. The main difference between the conditions will be assumptions made on the Lagrange multipliers λ_t of the associated Ramsey plan.

Converging multiplier. We begin with the case where the Lagrange multiplier converges to some nonzero constant.

Proposition 2. *If a pair of prices (r, w) is part of a Ramsey steady state of a Ramsey plan with converging Lagrange multipliers λ_t, then two conditions have to hold:*

1. *The steady state government budget constraint has to hold,*

$$G + r\mathcal{A}^{ss}(r, w) = (1 - w)\mathcal{N}^{ss}(r, w) \tag{24}$$

2. *The following RSS optimality condition has to hold*

$$\underbrace{(1 - \delta(1 + r))\,\ell\left(mrs\,\epsilon^{A,r} + \epsilon^{A,\tau}\right)}_{\text{liquidity benefit}} \underbrace{- \frac{1 - w}{w}\left(-\epsilon^{N,\tau} - mrs\,\epsilon^{N,r}\right)}_{\text{labor supply}} \underbrace{- (\ell\, mrs - 1)}_{\text{redistribution}} = 0 \tag{25}$$

where we omitted δ as an argument of elasticities and the mrs.

Proposition 2 gives us two simple conditions to check whether (r, w) is indeed consistent with an RSS. The first is simply the government budget constraint. The second is a first-order condition. The key thought experiment underlying it is a shift towards a higher interest rate r at the RSS, and higher labor income taxation τ (or, equivalently, a lower wage w). We imagine the planner contemplating this shift, balancing the r and the w shift in a way to keep utility exactly unchanged (21). One could regard such a shift as being compensated. The three terms in (25) consist of three forces that indicate whether the planner can gain resources from undertaking such a compensated shift or not. If there are no resources to gain, the sum of the three terms nets out to zero. This has to be satisfied at any RSS. We now go over each of the three terms in turn.

The *liquidity benefit* term captures the idea that the planner can raise the present value of resources by issuing more debt when there is a gap between the social discount factor δ and the interest rate. An additional unit of debt issued brings the planner one additional resource, but requires payment of $1 + r$ in the future. In present value terms, this gives a net resource gain of $1 - \delta(1 + r)$ per unit of debt issued. The expression $\mathrm{mrs}\,\epsilon^{A,r} + \epsilon^{A,\tau}$ can be interpreted as the net increase in debt—or, equivalently, assets held by households—as a result of the compensated increase in interest rates and labor income taxes. It is worth noting that the sign of the liquidity benefit will be positive whenever the planner uses the private discount factor, $\delta = \beta$, but can be negative when the planner uses relatively high discount factors, e.g. $\delta \to 1$ as for the optimal steady state we study below.

The second term, labeled *labor supply*, captures any disincentive effects on labor supply of greater labor income taxation, via $-\epsilon^{N,\tau}$, and interest rates, via $-\mathrm{mrs}\,\epsilon^{N,r}$. One would intuit that this term is negative, indicating lower labor supply. However, as already anticipated in section 2.4, this will not always be the case.

The final term is a *redistribution* term. It captures that the combined increase in interest rates and labor income taxes has distributional consequences. The planner will, in general, have to use more than one unit of resources in terms of higher interest payments to compensate households for a one unit reduction in after-tax wages, simply because asset income earners have, on average, lower marginal utility. This is especially clear in the case $\delta = \beta$, where (22) shows that $\ell\,\mathrm{mrs}$ is simply a ratio of weighted marginal utilities.

Our next result shows that (25) can be simplified further if preferences are of the log-separable form (12).

Proposition 3. *Assume $u(c,n)$ is log-separable, as in (12). Then, condition (25) in proposition 2 is equivalent to*

$$\underbrace{(1-\delta(1+r))\,\ell\epsilon^{A,r}}_{\text{liquidity benefit}}\underbrace{-\frac{1-w}{w}\left(-\epsilon^{N,r}\right)}_{\text{labor supply}}\underbrace{+\frac{1}{mrs}}_{\text{redistribution}}=0 \tag{26}$$

Surprisingly, only interest rate elasticities appear in (26). The reason behind this result goes back to a specific kind of symmetry between r and w derivatives in a log-separable model, which is explained in detail in Auclert et al. (2023).

Diverging multiplier. The next result we study considers the case where the Lagrange multiplier λ_t diverges to infinity.

Proposition 4. *If a pair of prices (r,w) is part of a Ramsey steady state of a Ramsey plan with diverging Lagrange multipliers, $\lambda_t \to \infty$ with $\lambda_t/\lambda_{t-1} \to \eta \in [1,\delta^{-1}]$, then the government budget constraint (24) as well as two optimality conditions*

$$(1-\delta\eta(1+r))\,\ell\epsilon^{A,\tau}(\delta\eta)-\frac{1-w}{w}\left(-\epsilon^{N,\tau}(\delta\eta)\right)+1=0 \tag{27}$$

$$(1-\delta\eta(1+r))\,\ell\epsilon^{A,r}(\delta\eta)-\frac{1-w}{w}\left(-\epsilon^{N,r}(\delta\eta)\right)-\ell=0 \tag{28}$$

have to hold.

A diverging Lagrange multiplier implies that the planner eventually perceives its most important role to be to raise available resources, independent of consequences for utility. The two optimality conditions (27) and (28) capture this idea. If (27) holds, the planner cannot raise additional resources by raising labor income taxes further. If (28) the planner cannot raise additional resources by raising interest rates further. Observe that in neither of the two conditions do derivatives of the utility function enter; and the relevant discount factor is effectively no longer δ, but instead $\delta\eta$, adjusted by the growth rate of the multiplier.

It is intuitive that there are three conditions in this case, (24), (27) and (28), as there are three unknowns to pin down, namely r, w, and η.

Figure 2: Two conditions that determine the Ramsey steady state

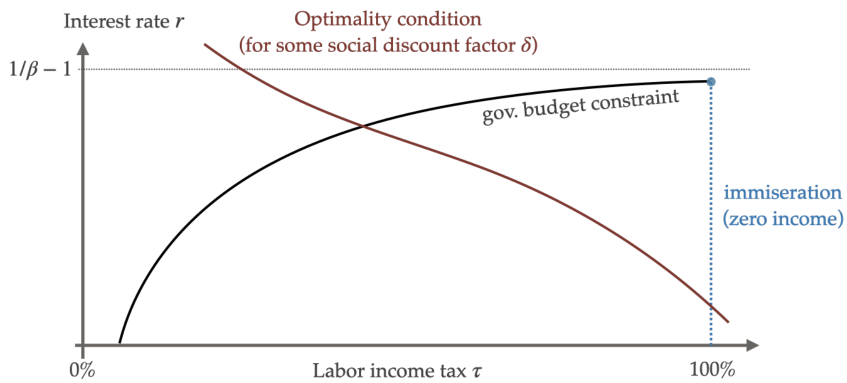

3.4 Graphical illustration of the optimality condition

Figure 2 graphically illustrates the determination of the Ramsey steady state in the case of a converging Lagrangian multiplier. We look for a pair of an interest rate r (on the y axis) and a labor income tax rate $\tau = 1 - w$ (on the x axis) such that two conditions hold: *(i)* the government budget constraint (24) (black solid line) requires that higher interest rates, and thus liquidity provision by the government, needs to be financed by greater labor income taxes; *(ii)* the optimality condition (25) (red dashed line) requires that we optimally resolve the trade-off between interest rates and taxes. The budget constraint in r, τ space typically slopes up, while the optimality condition typically slopes down, though we will see exceptions to this in section 5.

4 Immiseration in Aiyagari models

We now put our optimality conditions into action. In this section, we do so focusing on log-separable preferences (12). We discuss alternative preferences in section 5. Our main result in this section will be that if the social discount factor equals the private one, $\delta = \beta$, the two lines in figure 2 typically do not cross. Moreover, there is no rate η at which the Lagrange multiplier could diverge that gets the three conditions of proposition 4 to hold. This lets us conclude that in many standard Aiyagari models, there is no well-defined Ramsey steady state. Instead, we find suggestive evidence that the economy tends to immiseration, with labor income taxes τ approaching 100%.

4.1 Baseline calibration

As a starting point, we use the following baseline calibration of the economy. The disutility from labor $v(n)$ in (12) is isoelastic with a Frisch elasticity of 1. The Markov chain for idiosyncratic productivities e_{it} follows an AR(1) income process with annual persistence 0.90 and a cross-sectional

Figure 3: Net benefit from higher liquidity at Ramsey steady state (RSS)

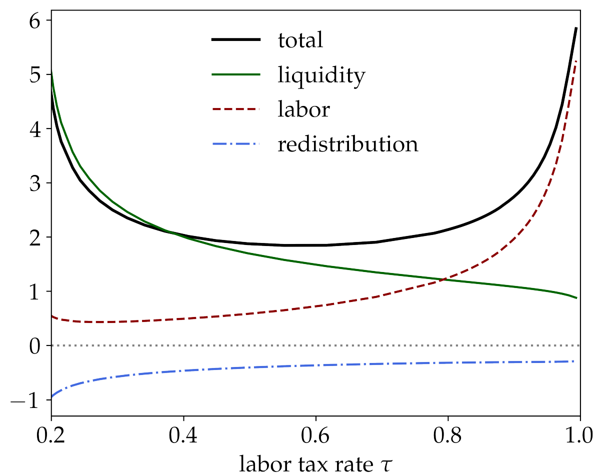

Note: This figure displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 3) and their sum under the baseline calibration (see sec. 4.1 for further details). Households have log-separable preferences and a zero-borrowing constraint, and the planner shares the households' preferences, $\delta = \beta$.

standard deviation of log e_{it} equal to 0.92. The initial steady state of the economy is one with debt B equal to 100% of GDP, government spending G equal to 20% of GDP, and an interest rate of 2%. This implies a private discount factor of $\beta = 0.897$ to satisfy market clearing; and a tax rate of $\tau = 22\%$ to satisfy the government budget constraint.

We discuss in great detail below how our results depend on these numbers. We find that for a wide range for standard calibrations, the economy either tends to immiseration, without a Ramsey steady state, or to a Ramsey steady state with very high tax rates, often north of 90%.

4.2 In search of a Ramsey steady state

We begin our search for a Ramsey steady state with an evaluation of the two conditions in proposition 2 in the case where the planner shares the households' preferences, $\delta = \beta$. We do so by varying the potential RSS labor income tax τ, then solving for the associated interest rate r that satisfies the government budget constraint (24), and finally by evaluating the optimality condition (25) for that pair of τ and r. As before, we note that w is simply equal to $1 - \tau$. Graphically, we can think of the approach as walking from the bottom left to the top right corner along the black solid line in figure 2, always evaluating the optimality condition (25), seeing whether it ever crosses zero.

Figure 3 plots the three terms on the left hand side of the optimality condition (25) (various colors) as well as the sum of all terms (black, solid). As expected, the liquidity effect is positive. Furthermore, the redistribution is negative. However, the labor supply effect is *positive*, indicating that greater labor taxes and interest rates lead to a *positive* response of the present discounted value

Figure 4: Net benefit from higher labor taxes at the optimal steady state (OSS)

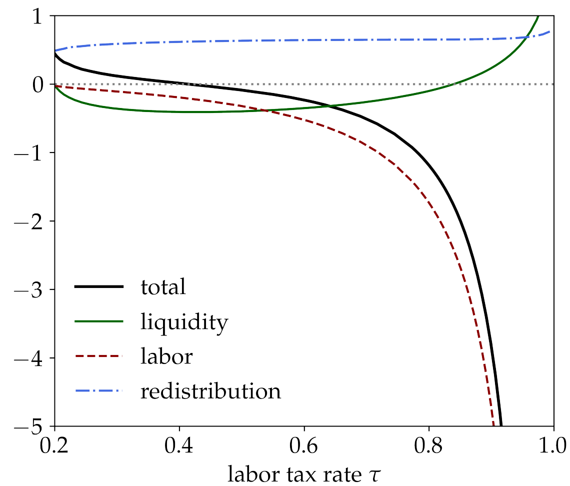

Note: This figure displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 3) and their sum under the baseline calibration (see sec. 4.1 for further details). The planner maximizes steady-state welfare, $\delta = 1$.

of labor supply—much in line with our discussion on the anticipatory labor supply effects of future labor income taxes in section 2.4.

The positive labor supply effect prevents the sum of all three effects (black solid) from ever crossing zero. There is no Ramsey steady state with a converging Lagrange multiplier in this economy.

4.3 Optimal steady state and other social discount factors

Figure 4 repeats the same calculation for the case of a social discount factor equal to 1, $\delta = 1$. This can be interpreted as a planner that is maximizing steady state welfare only, ignoring the transition it takes to reach the steady state. As it turns out, by assuming planner and household preferences are misaligned, the three terms in (25) become harder to interpret. For example, the first term capturing the liquidity benefit is now negative, as the planner no longer uses a discount factor below $1/(1+r)$, like β, and instead ceases to discount altogether. Additionally, the redistribution term is now positive as greater steady state interest rates no longer simply benefit today's asset rich households, but *any* household, as they all hold some assets eventually.

The labor supply term, however, is still capturing disincentive effects for greater interest rates and labor income taxation. As figure 4 shows, these now go in the intuitive direction and turn negative eventually. This ensures a unique intersection of our optimality condition with zero, and hence ensures a unique optimal steady state (OSS) in this economy.

Similar to our procedures for $\delta = \beta$ and $\delta = 1$, we can also search for Ramsey steady state for any intermediate social discount factor δ. Figure 5 follows this approach and computes various steady state outcomes, such as consumption, wages, or liquidity, whenever a candidate Ramsey steady state for a given δ was found. The figure shows that, such a steady state can be found for δ's very close to β.[4] Interestingly, as δ is moved from 1 all the way to the left, towards β, many quantities and prices approach a corner. For example, consumption converges to zero, GDP converges to government spending G, and labor taxes converge to 100%. This is a first indication as to what might happen in lieu of convergence to a Ramsey steady state when $\delta = \beta$: the economy might tend to long-run immiseration.

4.4 Immiseration

Our results in figure 5 suggest that the Ramsey plan of a planner that shares the household preferences, i.e. $\delta = \beta$, tends to immiseration: tax rates τ_t rise to 100%, after-tax wages w_t fall to zero. None of our conditions in section 3.3 apply to this situation as there is no well-defined long-run steady state equilibrium here. We next provide a necessary condition that specifically captures immiseration dynamics in the case of log-separable preferences (12).

Proposition 5. *Assume $u(c,n)$ is of the log-separable form, (12). Let $\{r_t, w_t\}$ be an optimal Ramsey plan such that:*

- *w_t falls to zero with asymptotic decay factor $\gamma \in [\delta, 1)$, that is, $\lim_{t\to\infty} w_t/\gamma^t$ exists and is positive.*
- *r_t converges to some constant $r < \gamma/\beta - 1$.*
- *λ_t diverges at with asymptotic factor $\eta \in (1, \delta^{-1}]$, that is, $\lim_{t\to\infty} \lambda_t/\eta^t$ exists and is positive.*

Then, the following two conditions have to hold, evaluated at a de-trended steady state with interest rate $1 + \hat{r} = (1 + r)/\gamma$, same discount factor β, and wage $\hat{w} = 1$:

1. *an immiseration-adjusted government budget constraint*

$$G = \mathcal{N}^{ss}(\hat{r}, \hat{w}) \tag{29}$$

2. *an immiseration-adjusted optimality condition*

$$\epsilon^{N,\tau}(\delta\eta) = \epsilon^{N,r}(\delta\eta) = 0 \tag{30}$$

Proposition 5 is, once more, a necessary condition that has to hold if the economy tends to immiseration at some exponential rate. It is based on the idea that a household side whose after-tax wages are shrinking at an exponential rate, can be "de-trended". With log-separable

[4] There is a small interval of positive measure $[\beta, \tilde{\beta})$ just above β in which social discount factors do not give rise to a candidate Ramsey steady state.

Figure 5: Ramsey steady state as function of social discount factor

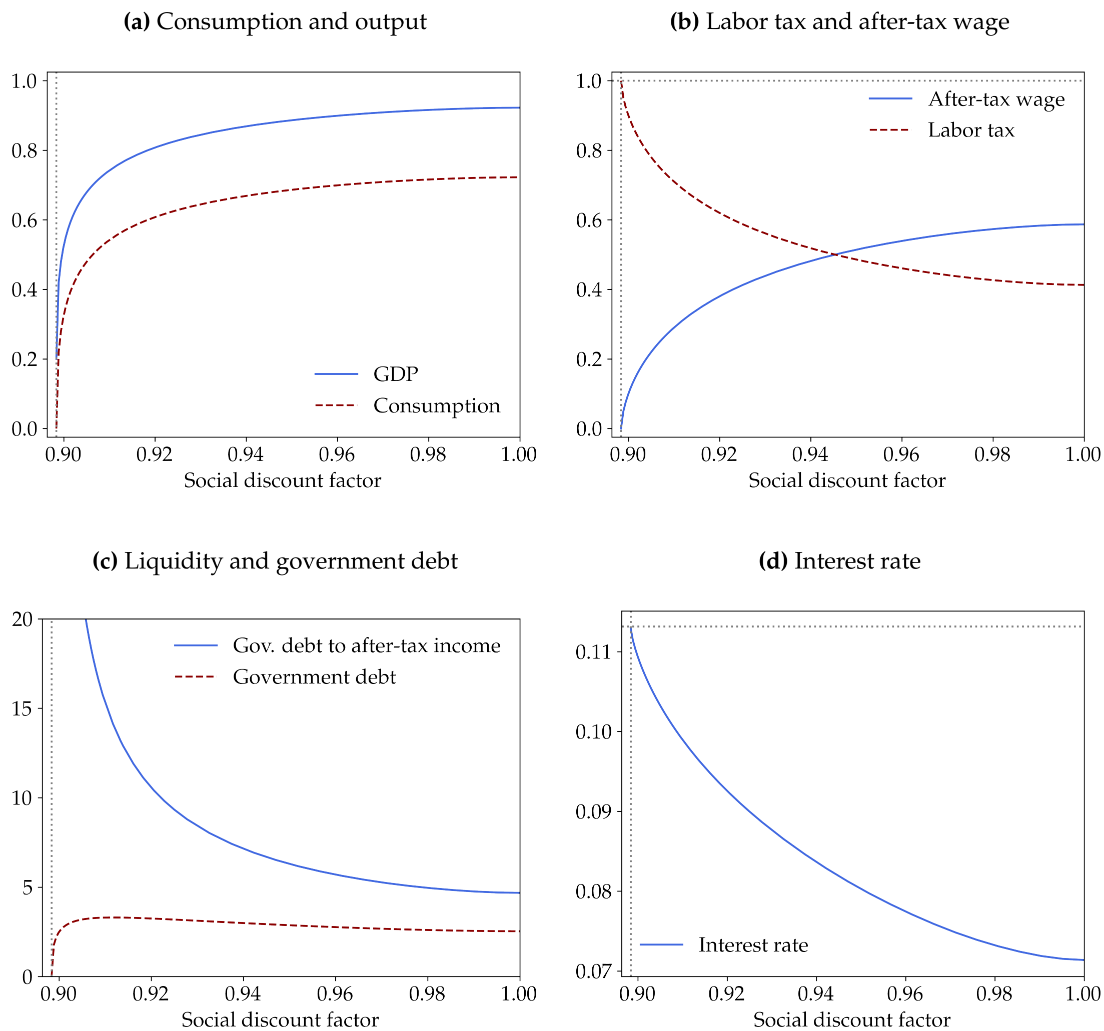

Note: Each panel in this figure displays the interior Ramsey steady state values of variables attained as the social discount factor δ is varied from $\delta = \beta$ (grey-dotted vertical line) to $\delta = 1$ under the baseline calibration (see sec. 4.1 for further details).

Figure 6: Looking for immiseration

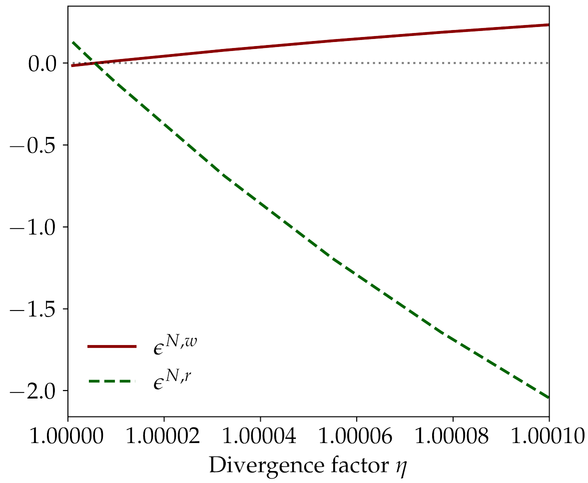

Note: This figure displays labor elasticities with respect to post-tax wages and interest rates for the optimal Ramsey plan whose Lagrange multipliers diverge at exponential rate η. Parameters are given by the baseline calibration (see sec. 4.1 for further details).

preferences, (12), the de-trended household side is one whose labor supply is given by $\mathcal{N}^{ss}\left(\hat{r}, \hat{w}\right)$, where $1+\hat{r} = \left(1+r\right)/\gamma$ is a "de-trended interest rate" and the wage $\hat{w}$ is normalized to 1; and whose asset demand is given by $w_t \cdot \mathcal{A}^{ss}(\hat{r}, \hat{w})$. With these relationships, the government budget constraint (24) collapses to (29).

The optimality condition (30) can, at least heuristically, be derived from our optimality conditions with diverging multipliers (27) and (28). For example, the only way (27) can continue to hold as $w \to 0$ is if $\epsilon^{N,\tau}(\delta\eta) = 0$. Similarly, $\epsilon^{N,r}\left(\delta\eta\right) = 0$ follows from (28) in this limit. As it turns out, with log-separable preferences, $\epsilon^{N,\tau}\left(\delta\eta\right) = 0$ precisely if and only if $\epsilon^{N,r}\left(\delta\eta\right) = 0$, for arbitrary $\delta\eta$. Thus, (30) should be thought of as a single condition, not two independent ones.

It is important to acknowledge that proposition 5 does not allow us to back out the rate at which the economy tends to immiseration. Instead, (29) and (30) are two conditions that allow to determine the de-trended interest rate $\hat{r}$ as well as the factor η with which the Lagrange multiplier diverges. To show how this works in our baseline calibration, we solve (29) for our baseline economy, finding $\hat{r} = 11.47\%$. Figure 6 plots $\epsilon^{N,\tau}\left(\beta\eta\right)$ and $\epsilon^{N,r}\left(\beta\eta\right)$ varying η. As the figure clearly shows, an η just above 1 is sufficient to turn $\epsilon^{N,\tau}\left(\beta\eta\right)$ and $\epsilon^{N,r}\left(\beta\eta\right)$ to zero.

4.5 Role of government spending

Our immiseration result raises many questions. First among them is the role our baseline calibration plays. As we mentioned in section 4.1, we find immiseration or very high limiting tax rates τ across a wide variety of calibration choices, such as the level of initial debt, or the initial interest rate,

Figure 7: Role of government spending G

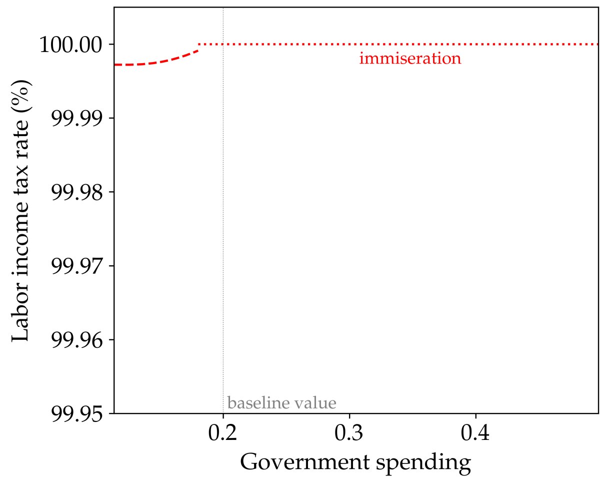

moves the needle away from immiseration. In the following four subsections, we investigate the roles of four different modeling assumptions: the amount of government spending G, the income process, the Frisch elasticity, and the role of borrowing constraints.

We begin with government spending G. Figure 7 plots the labor income taxes and consumption in the Ramsey steady state as we vary government spending G, keeping all other model parameters unchanged. As can be seen, for G below some threshold value around 0.17, a Ramsey steady state exists, albeit one associated with very high labor income taxes and small consumption. As we raise government spending beyond 0.17, no Ramsey steady state exists. In this region, we our results suggests that immiseration is optimal, with taxes approaching 100% and consumption falling to zero.

4.6 Role of income inequality

We have explored many different kinds of income processes used in the literature, ranging from standard AR(1) specifications with various parameterizations all the way to processes with extreme values (as in Castañeda et al. 2003) and leptokurtic shocks (as in Kaplan et al. 2018). All give rise to qualitatively very similar conclusions to those shown above. In the following, we explore a specification that gives rise to extreme levels of income inequality, just to see what can happen in that case. This is especially interesting given that there was a negative redistribution effect from higher interest rates and labor income taxes in figure 3. Maybe that redistribution effect will be sufficiently big to pull the total net benefit from greater liquidity provision down to zero, producing a well-defined Ramsey steady state?

For this subsection, we modify our AR(1) income process in two ways. First, we raise the standard deviation of the innovations to $\log e_{it}$ such that $\log e_{it}$ has a standard deviation of 1.5, far bigger than even the high standard deviation of pre-tax income in the U.S. economy, which is around 0.90. Second, we add an additional state to the Markov chain, which we call "poverty state". Households transition into the poverty state with a probability p and leave the poverty state with probability q. The mass of households in the poverty state is given by $\mu = p/(p+q)$. To be as extreme as possible, we assume that the poverty state is essentially permanent, $p, q \to 0$, but fix the ratio q/p in order to target a specific fraction μ in the poverty state. We also make the extreme assumption that households in the poverty state essentially have a labor productivity of zero, $e_{it} \to 0$.

We show in appendix A.17 that the inclusion of such a poverty state modifies the optimality condition (25) by scaling up $\mathrm{mrs}(\delta)$ by a factor of $\frac{1}{1-\mu}$. This is intuitive, as households in the poverty state see no reason to accumulate assets and thus rely entirely on labor income. They are especially affected by labor income tax increases and do not benefit at all from higher interest rates. Thus, the relevant optimality condition is

$$\underbrace{(1-\delta(1+r))\,\ell\left(\frac{\mathrm{mrs}}{1-\mu}\,\epsilon^{A,r}+\epsilon^{A,\tau}\right)}_{\text{liquidity benefit}}\underbrace{-\frac{1-w}{w}\left(-\epsilon^{N,\tau}-\frac{\mathrm{mrs}}{1-\mu}\,\epsilon^{N,r}\right)}_{\text{labor supply}}\underbrace{-\left(\ell\,\frac{\mathrm{mrs}}{1-\mu}-1\right)}_{\text{redistribution}}=0 \qquad (31)$$

where all elasticities and the mrs are computed in the economy without poverty state, i.e. assuming $\mu = 0$. Among other effects, this increases the magnitude of the negative redistribution term.

Figure 8(a) plots all three terms of (31) with our extreme calibration. We see that, indeed, the redistribution term is sufficiently large to pull the sum of all three terms down to zero. Two candidate Ramsey steady states emerge. Both have relatively high labor taxation—one at around 70%, the other at around 95%. Though we cannot check second-order conditions, we speculate that the lower tax steady state is a local maximum, while the higher tax steady state is a local minimum.

Importantly, however, to the right of the high tax steady state, the left hand side of (31)—the net benefit from further liquidity creation—is positive again. This suggests that, despite having found two candidate Ramsey steady states, immiseration may still be possible. Figure 8(b) follows our approach from figure 5 and computes Ramsey steady states varying the social discount factor δ. We see that, zooming in close to β, there are, in fact three candidate Ramsey steady states: two that converge to those found in figure 8(a), and one that leads to immiseration. This corroborates our suspicion that even extreme inequality does not affect the viability of immiseration as long-run outcome of the Ramsey plan.

We can also apply the necessary conditions for immiseration from proposition 5. As the economy still has log-separable preferences, and (30) is independent of the mrs, the immiseration conditions are satisfied for a given pair of r and γ irrespective of the size of the poverty state μ.

Figure 8: Ramsey steady states with extreme income inequality

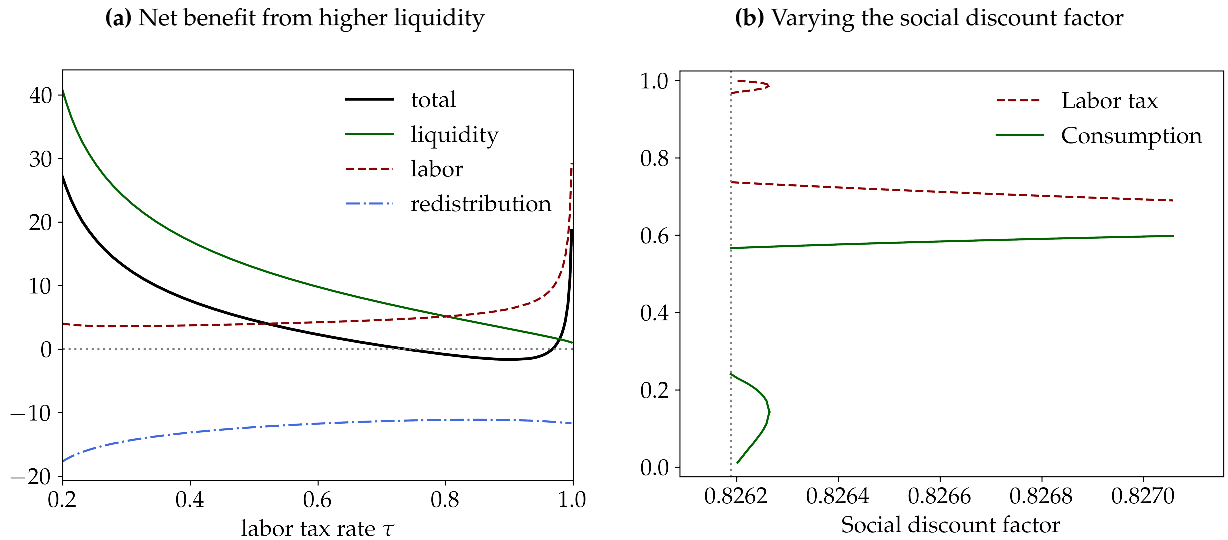

Note: Panel a displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 2) and their sum for a modified calibration where 90 percent of households are hand-to-mouth and the remaining households face heightened idiosyncratic income risk, with a cross-sectional standard deviation of $\log e_{it}$ of 1.5. Panel b displays the interior Ramsey steady state values of labor taxes and aggregate consumption attained as the social discount factor δ is varied from $\delta = \beta$ (grey-dotted vertical line) to $\delta = 1$.

4.7 Role of the Frisch elasticity

We move on to the Frisch elasticity. Our baseline calibration assumed a Frisch elasticity of 1. Varying the Frisch elasticity in a reasonable range doesn't affect our results qualitatively. Our results do change, however, when the Frisch elasticity is assumed to be very close to zero.

Figure 9(a) plots the net benefit from higher liquidity for an economy with a Frisch elasticity of 0.05. The labor supply margin is now nearly absent for all reasonable levels of labor taxes. For taxes just below 100%, however, we find that the labor supply margin turns negative, implying that in this case, a Ramsey steady state exists, albeit one very close to immiseration, with tax rates around 99%. Panel (b) shows that consumption is relatively steady with a low Frisch elasticity. In the limit of a zero Frisch elasticity, consumption would be exactly constant at the then constant labor supply N minus government spending G. Still, we suspect that in this limit, Ramsey steady state tax rates are close to 100%.

Figure 9: Ramsey steady states with relatively inelastic labor supply

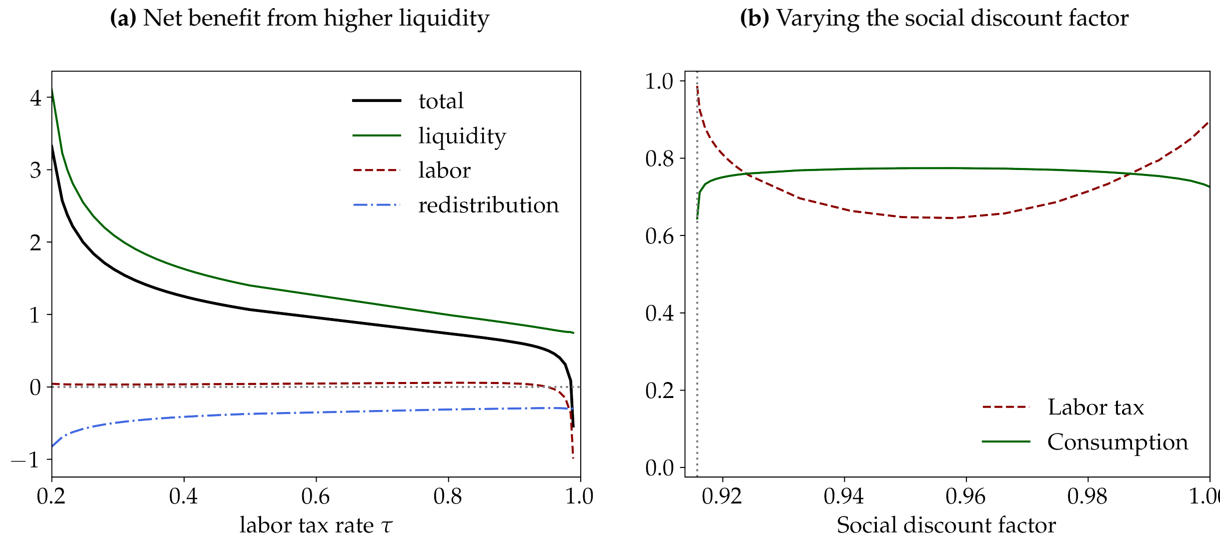

Note: Panel a displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 2) and their sum for a modified calibration with a Frisch elasticity of 0.05. Panel b displays the interior Ramsey steady state values of labor taxes and aggregate consumption attained as the social discount factor δ is varied from $\delta = \beta$ (grey-dotted vertical line) to $\delta = 1$.

4.8 Role of the borrowing constraint

For our final subsection, we relax the assumption of a borrowing constraint of zero in the household problem (3). Instead, we assume that the borrowing constraint is given by

$$\underline{a}_t = -\frac{\underline{e}\, w_t\, \bar{n}}{\beta^{-1} - 1}$$

We choose this form as it naturally scales with the average post-tax wage w. $\underline{e}$ is the lowest productivity in the Markov chain, $\bar{n}$ is chosen to be equal to 1, consistent with the optimal labor supply of the borrowing-constrained, lowest-productivity household under the baseline calibration. Figure 10 shows that even with a relaxed borrowing constraint, the net benefit from greater liquidity is still positive above a labor tax of around 25%. It is not surprising that little is change as we could already see in figure 5 that household liquidity ℓ becomes very large as $\delta \searrow \beta$, in which case most households will be far away from their borrowing constraints anyway.

5 Alternative preferences

We now start considering classes of preferences other than the log-separable ones used in the previous section. We consider three kind: general balanced growth preferences; general additively

Figure 10: Net benefit from higher liquidity with relaxed borrowing constraint

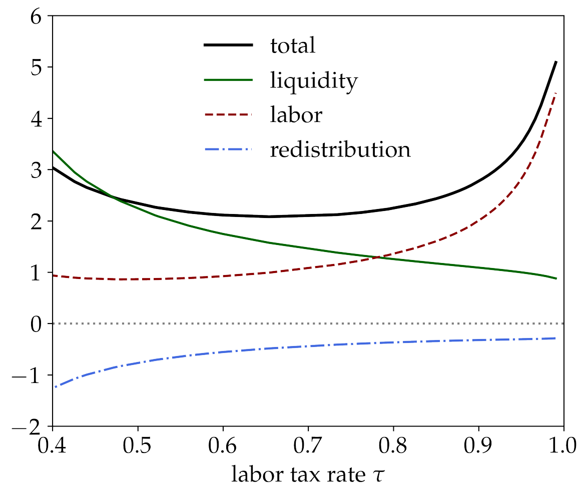

Note: This figure displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 2) and their sum for a modified calibration where households have a non-zero borrowing constraint, given by $\underline{a}_t = -\frac{\underline{e}\, w_t \bar{n}}{\beta^{-1}-1}$ where $\underline{e} \approx 0.07$ and $\overline{n} = 1$.

separable preferences; and Greenwood et al. (1988) preferences without wealth effects on labor supply.

5.1 Balanced growth preferences a la King et al. (1988)

We begin by relaxing the assumption of a unitary elasticity of intertemporal substitution (EIS) to move from log-separable preferences to more general balanced-growth compatible preferences. Figure 11(a) computes the net liquidity benefit as we vary labor taxes for an economy with $\sigma^{-1} = 1.05 > 1$, showing that it is consistently above 0.[5] Figure 11(b) shows that, just like before, as $\delta \searrow \beta$, we find candidate Ramsey steady states that converge to an allocation with zero consumption and labor tax rates of 100%.

Figure 12 focuses on the case of an EIS of 0.5. Here, we do find a Ramsey steady state, as the contribution of the labor term eventually turns negative. However, the associated tax rate is very high, above 90%, and consumption is only 30% of the calibrated steady state output.

[5] For greater values of the EIS, we have not been able to find a unique stationary household wealth distribution at the initial steady state. We suspect that this is a consequence of the substitutability between hours and consumption induced by $\sigma^{-1} > 1$.

Figure 11: RSS with balanced growth preferences and EIS $\sigma^{-1} > 1$

(a) Net liquidity benefit **(b)** Varying the social discount factor

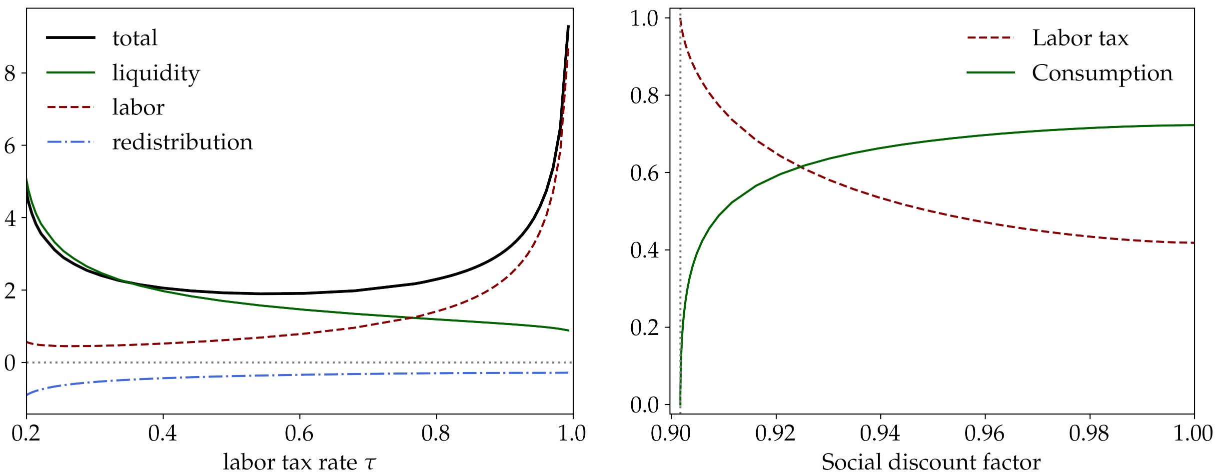

Note: Panel a displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 2) and their sum for a modified calibration with balanced growth preferences and EIS $\sigma^{-1} = 1.05$. Panel b displays the interior Ramsey steady state values of labor taxes and aggregate consumption attained as the social discount factor δ is varied from $\delta = \beta$ (grey-dotted vertical line) to $\delta = 1$.

Figure 12: RSS with balanced growth preferences and EIS < 1

(a) Net liquidity benefit **(b)** Varying the social discount factor

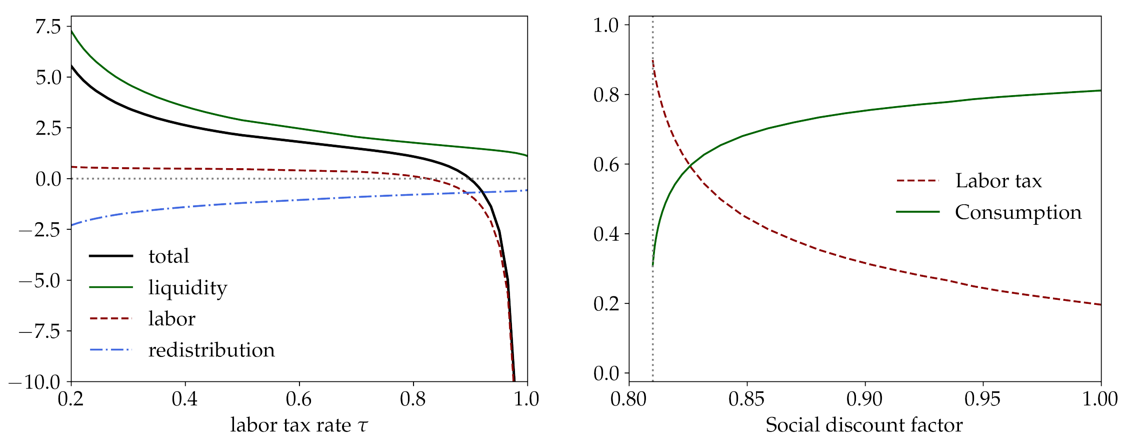

Note: Panel a displays the three terms of the interior Ramsey steady state optimality condition (stated in Proposition 2) and their sum for a modified calibration with balanced growth preferences and EIS $\sigma^{-1} = 0.5$. Panels b displays the interior Ramsey steady state values of labor taxes and aggregate consumption attained as the social discount factor δ is varied from $\delta = \beta$ (grey-dotted vertical line) to $\delta = 1$.

Figure 13: Two RSS conditions with additively separable preferences

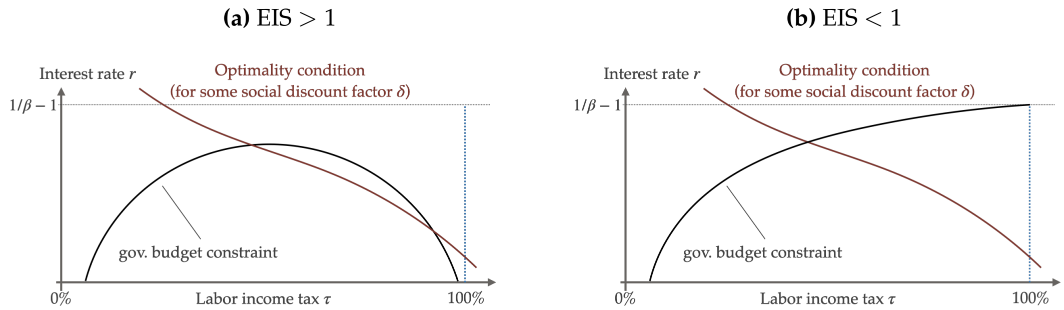

5.2 Additively separable preferences

Next, we consider additively separable preferences,

$$u(c,n) = \frac{c^{1-\sigma}-1}{1-\sigma} - \phi\frac{n^{1+\nu}}{1+\nu} \tag{32}$$

where $\phi > 0$ is a constant and ν is the inverse of the Frisch elasticity of labor supply. For $\sigma \neq 1$, these preferences are no longer compatible with balanced growth. This has a significant effect on the possible long-run behavior of Ramsey plans. To see why, consider the two diagrams in figure 13.

Panel (a) focuses on the case where the EIS $1/\sigma$ lies above 1. This is known as the case where the substitution effect from wage changes is stronger than the income effect. This leads the steady state government budget constraint look very different from the one in figure 2. As the labor tax gets closer to 100%, the substitution effect becomes sufficiently strong that households increasingly work less. This reduces tax revenue and hinders liquidity provision, lowering the equilibrium interest rate again. We obtain a hump-shaped profile.

Panel (b) focuses on the case where the EIS $1/\sigma$ lies below 1. Here, we have the opposite. As labor taxes rise, households feel increasingly poor, and start working harder, generating additional tax revenue that allows the government to raise liquidity provision and raise the equilibrium interest rate r.

EIS $1/\sigma > 1$. We begin numerically searching for candidate Ramsey steady states in the case where the EIS $1/\sigma$ lies above 1. Other than the different preferences, the calibration is the same as the one introduced in section 4.1. Due to the hump-shaped profile in figure 13(a), there can be anywhere between zero and two intersections of the optimality curve with the government budget constraint. Zero intersections are more likely when the budget constraint line is lower, as is for example the case with higher levels of government consumption G.

Panel (a) of figure 14 illustrates how there are two candidate Ramsey steady states when