

December 2020

“Public Debt as Private Liquidity: Optimal Policy”

George-Marios Angeletos, Fabrice Collard and Harris Dellas

Public Debt as Private Liquidity: Optimal Policy*

George-Marios Angeletos

MIT

Fabrice Collard

Toulouse School of Economics

Harris Dellas

University of Bern

November 22, 2020

Abstract

We study optimal policy in an economy in which public debt is used as collateral or liquidity buffer. Issuing more public debt raises welfare by easing the underlying financial friction; but this easing lowers the liquidity premium and increases the government's cost of borrowing. These considerations, which are absent in the basic Ramsey paradigm, help pin down a unique, long-run level of public debt. They require a front-loaded tax response to government-spending shocks, instead of tax smoothing. And they explain why a financial recession, more than a traditional one, makes government borrowing cheaper, optimally supporting larger fiscal stimuli.

*This paper supersedes an earlier draft, entitled "Optimal Public Debt Management and Liquidity Provision", which was concerned with the same topic but did not contain the present paper's theoretical contribution. We are grateful to Behzad Diba for his collaboration on the earlier project; to Pedro Teles and Per Krusell for illuminating discussions; and to various seminar participants for their feedback. Angeletos also thanks the University of Bern, Study Center Gerzensee, and the Swiss Finance Institute for their hospitality.

1 Introduction

Liquidity shocks and shortages of private collateral interfere with the efficient allocation of resources. A theoretical literature has emphasized that public debt issuance may ease such frictions by contributing to the supply of assets that can be used as collateral or buffer stock (Woodford, 1990; Aiyagari and McGrattan, 1998; Holmström and Tirole, 1998). In the same spirit, an empirical literature has shown that, even after controlling for default risk, the spread between government and private bonds is both substantial and sensitive to the quantity of public debt (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood and Vayanos, 2014).

What are the implications of these considerations for optimal fiscal policy? Do they lead to a well-defined optimal long-term target for public debt and, if yes, what are its determinants? How do they matter for the optimal policy response to fiscal shocks or other business cycle shocks?

To address these questions, we augment the basic Ramsey paradigm (Barro, 1979; Lucas and Stokey, 1983; Aiyagari et al., 2002) with a liquidity function for public debt as suggested by the aforementioned literature. In our model, public debt's use as collateral helps ease the reallocation of a consumption good across households, or of capital across firms. Such reallocation is necessary because of idiosyncratic taste or productivity shocks. Tractability is nevertheless preserved by use of a quasi-linear specification as in Lagos and Wright (2005). We are thus able to reduce the planner's problem in our economy to an analytically solvable, albeit non-convex, optimal control problem, in which it is *as if* public debt enters (i) the utility or the production function and (ii) the interest-rate cost of government borrowing.

The optimal policy, both in the long and the short run, is dictated by the interplay of three forces. The first is the desire to smooth taxes. The second is the desire to *ease* the financial friction so as to improve private allocations. The third, which is perhaps the most novel element of our analysis, is the desire to *preserve* the financial friction so as to suppress the interest-rate cost of public debt.

The relative importance of these three forces varies across the short and the long run. The trade off between improving the private sector's allocations and keeping the government's interest-rate costs low is the key determinant of the steady-state level of public debt. But tax smoothing naturally gains influence when considering transition to the steady state or the optimal response to shocks.

Consider first the long run. In our model, the government could eliminate the financial friction by issuing a sufficiently large amount of public debt to satiate the economy's demand for collateral. But doing so would have adverse budgetary consequences as it would eliminate the "profit" that the government enjoys by paying an interest rate on public debt that is below the underlying social discount rate. Crucially, this trade off hinges, not merely on taxation being distortionary, but also on the fact that the price of public debt varies negatively with its quantity: in the absence of this effect, the optimal quantity of public debt converges in the long run to the level that satiates the economy's demand for collateral.¹

¹This presumes, for the sake of the argument, that the satiation level is lower than the maximal sustainable level of debt.

This trade off can support a unique steady-state level of debt to which the economy converges for any initial position below satiation. This contrasts with the steady-state indeterminacy –induced by tax smoothing, in the standard model. Furthermore, there is a threshold strictly above satiation such that, for initial levels of debt in between these two, it is optimal to reduce debt below satiation. This illustrates the importance of the interest-rate suppression motive: the government may optimally create a shortage of collateral in an economy where collateral was initially abundant.

The same motive also figures prominently in the optimal response to shocks. Consider, for instance, an unanticipated, uninsured, positive shock to government spending. In the basic Ramsey paradigm, this triggers a permanent increase in taxes by an amount equal to the annualized innovation in the present discounted value of government spending. In our setting, instead, taxes increase relatively more early on in order to keep interest rates on debt low, enabling a smaller tax burden later on.

In a similar vein, consider a shock that raises the labor wedge, reduces aggregate output and tax revenue, and motivates a fiscal stimulus in the form of a temporary “payroll tax cut” to moderate the increase in the labor wedge. Suppose also that the resulting recession is associated—exogenously or endogenously—with more severe financial frictions. This has an ambiguous effect on the trade-off between liquidity provision and interest-rate suppression: while the larger friction encourages greater provision of liquidity, the higher profit that the planner can make by preserving the shortage of collateral pulls in the opposite direction. But it unambiguously increases fiscal space by reducing the interest rate on public debt, thereby supporting a larger fiscal stimulus.

This result provides a formal basis for the proposal, made by, among others, Paul Krugman and Brad DeLong, that the Great Recession called for high deficits not only because of the need to stimulate aggregate demand but also because the low interest rates made it “cheap” for the US government to borrow. This proposal has no place in the textbook Ramsey paradigm, because the price of the bonds coincides with the social discount factor. But it makes sense through the lenses of our analysis insofar as a lower interest rate is a signal of a heightened shortage of collateral and, hence, of a higher spread between the market price of public debt and the underlying social discount factor. From this perspective, a broader contribution of our paper is to draw attention to the policy implications of both the cyclicity of this spread and its endogeneity to the quantity of public debt.

Turning to the technical contribution of this paper, we highlight that, when the interest rate on public debt increases with its level, the policy problem becomes non-convex. As a result, the standard first-order approach does not apply: there can exist multiple paths that satisfy the planner’s Euler and transversality conditions, and the challenge is to find out which one of them is truly optimal.² An integral part of our contribution is to address this challenge by adapting the methods of Skiba (1978). This

²The non-convexity of the problem also explains why, in general, there can exist multiple steady states below satiation. We avoid this complication in the main analysis, and guarantee the existence of a unique such steady state, by making appropriate assumptions on the problem’s primitives. But we also explain why and how our main insights survive more broadly.

allows sharp analytical results, in contrast to Aiyagari and McGrattan (1998), who rely on numerical simulations, abstract from transitional dynamics and shocks, and, as explained in Section 4.1, mis-measure the shadow cost of government debt.

2 Micro-foundations

This section describes a micro-founded economy in which public debt serves as collateral and helps improve the allocation of resources. We characterize the equilibrium for given policy and show how the optimal policy problem nests in the class of reduced-form problems analyzed in the next section. We also illustrate the policy conflict between liquidity provision and interest-rate suppression.

2.1 Setup

There is a unit-mass of ex-ante identical households, indexed by $i \in [0, 1]$, and a representative firm. Time is discrete, indexed by $t \in \{0, 1, 2, \dots\}$, and each period is split into a “morning” and an “afternoon.” There are two edible goods. The one is the (exogenous) fruit of a tree, which becomes ripe in the morning of each period. The other is the (endogenous) output of the representative firm, which is produced in the afternoon with the labor of the households. Each good has to be consumed in the sub-period in which it is produced, or else it perishes. We refer to the first good as the “morning good” and to the second one, which is also our numeraire, as “the afternoon good.” Idiosyncratic risk takes the form of taste shocks to the utility of the morning good. The associated first-best allocation is impeded by a financial friction. This friction can be eased by saving in a riskless bond, whose supply is controlled by the government.

The representative firm. The representative firm is competitive and produces the afternoon good using labor. Aggregate output is given by $y_t = Ah_t$, where h_t denotes the labor input and A denotes the exogenous aggregate productivity (assumed to be time-invariant for the time being). It follows that, in equilibrium, the pre-tax wage is given by $w_t = A$ and all income goes to labor.

The households. Consider a household i . Let $h_{it} \in \mathbb{R}_+$ denote her period- t labor supply, and let $x_{it} \in \mathbb{R}_+$ and $c_{it} \in \mathbb{R}$ denote her period- t consumption of, respectively, the morning and the afternoon good. Her life-time utility is given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_{it}, x_{it}, h_{it}; \theta_{it}) \right]. \quad (1)$$

$\beta \in (0, 1)$ is the subjective discount factor, θ_{it} is an idiosyncratic taste shock, and U is given by

$$U(c, x, h; \theta) \equiv c + \theta u(x) - v(h), \quad (2)$$

where u is strictly increasing and strictly concave, and v is strictly increasing and strictly convex.

The taste shock is i.i.d. across households and follows a continuous Markov process, with transition density $\varphi(\theta'|\theta)$, unconditional density $\varphi(\theta)$, and support $[\underline{\theta}, \bar{\theta}]$. Its modeling role is to introduce a desire

for trade: high- θ agents would like to buy the morning good from low- θ agents. As we explain next, such trades are impeded by a financial friction—and this is where public debt enters the picture as a form of buffer stock or collateral. Finally, the linearity of U in c plays the same role as in Lagos and Wright (2005): it guarantees that the cross-sectional distribution of wealth is not a relevant state variable for the aggregate equilibrium dynamics and the planner’s problem.

Markets and frictions. In the afternoon of each period, households can buy and sell a risk-free asset, which delivers one unit of the numeraire good in the afternoon of the following period. This asset, whose price is denoted by q_t , may be issued either by the government or by other private agents: government and private bonds are perfect substitutes. In addition, households can trade short-term IOUs in each morning. These IOUs facilitate the transfer of resources within the period.

Let a_{it} denote household i ’s holdings of the risk-free asset—also, its net financial worth—in the beginning of period t . The period- t budget constraint can then be expressed as follows:

$$c_{it} + p_t x_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} + p_t \bar{e} \quad (3)$$

where p_t is the price of the morning good and \bar{e} is the fixed endowment of it.

Let $z_{it} = p_t(x_{it} - \bar{e})$ denote the value, in terms of the numeraire, of the household’s net trade of the morning good. One can interpret z_{it} as short-term (intra-period) credit lines that help cover “liquidity needs.” When $z_{it} > 0$, the household is a “borrower” in the sense that it finances its net purchase of the morning good by issuing an IOU against its afternoon labor income; and conversely, the household is a “lender” when $z_{it} < 0$.

Once the afternoon arrives, a borrower may be tempted to renege on her promise to pay back. If she does so, her lenders can confiscate a fraction $\xi \in (0, 1)$ of her labor income as well as all of her assets. For default to be averted in equilibrium, the following constraint must therefore hold:

$$z_{it} \leq \xi w_t h_{it}^{def} + a_{it} \quad (4)$$

where h_{it}^{def} denotes labor supply in the (off-equilibrium) event of default. Applying the same logic to inter-period borrowing, we get:

$$-a_{it+1} \leq \xi w_{t+1} h_{it+1}^{def}. \quad (5)$$

Condition (5) represents an upper bound on the agent’s debt, or equivalently on her net supply of the risk-free asset. Condition (4), on the other hand, shows how holdings of the risk-free asset serve as collateral in the IOU market, therefore enabling trade of the morning good.³

³By assumption, the risk-free asset is traded only in the afternoon. But because it can be posted as collateral, the equilibrium allocations remain unaffected if we let it be traded in the morning alongside, or in place of the IOUs. Accordingly, we can think of the risk-free asset interchangeably as collateral and as buffer stock.

The government. The government's budget constraint is given by

$$b_t + g = q_t b_{t+1} + \tau_t w_t h_t \quad (6)$$

where b_t is the stock of public debt inherited from period $t-1$, g is the exogenous level of government spending, $w_t h_t$ is labor income, and τ_t is the tax rate. For any given b_0 , the government chooses the sequence $\{\tau_t, b_{t+1}\}_{t=0}^{\infty}$ so as to maximize ex ante utility,

$$\mathcal{W} \equiv \mathbb{E} \left[\sum_{t=0}^{\infty} \beta^t (c_{it} + \theta_{it} u(x_{it}) - v(h_{it})) \right],$$

subject to its budget constraint and the applicable equilibrium restrictions (to be derived next).

2.2 Equilibrium

We characterize the equilibrium in three steps. First, we study the individual's problem and derive the *private* value of liquidity for a given sequence of prices and policies, $\{p_t, q_t, \tau_t, b_t\}_{t=0}^{\infty}$. Second, we solve for the equilibrium prices $\{p_t, q_t\}_{t=0}^{\infty}$ and derive the *social* value of liquidity. Finally, we show how to represent all the objects that enter the planner's problem as functions of the sequence $\{\tau_t, b_t\}$ alone.

Let us start by determining the pledgeable income of the household. This raises the question of what tax rate the household would face in the event of default. Since default is an *off*-equilibrium object, the associated tax rate can differ from the *on*-equilibrium tax rate τ_t . Without serious loss of generality, we assume that it is zero. Along with the fact that a fraction ξ of labor income is confiscated by lenders, this implies that labor supply under default solves $v'(h_{it}^{def}) = (1 - \xi)w_t$. Using $w_t = A$, we conclude that the financial constraints (4) and (5) can be restated as, respectively,

$$z_{it} \leq \phi + a_{it} \quad \text{and} \quad a_{it} \geq -\phi,$$

where $\phi \equiv \xi A (v')^{-1}((1 - \xi)A)$.⁴

We now proceed to study how the financial friction gives rise to a private value for liquidity. The optimal consumption of the morning good solves

$$\begin{aligned} & \max_x \{ \theta u(x) - p(x - \bar{e}) \} \\ & \text{subject to } p(x - \bar{e}) \leq \phi + a \end{aligned}$$

Clearly, for given p , the constraint is binding ex post when θ is high enough relative to a . By the same token, a higher a means a smaller ex ante chance that the collateral constraint will bind. This explains the precise sense in which there is a precautionary motive and the precise source of the equilibrium premium on public debt. Public debt is priced at a premium because it helps ease the collateral constraint.

⁴We henceforth treat A , ξ , and ϕ as constants and do not show the dependence of all endogenous objects on them. Also note that ϕ is strictly concave in ξ , with a maximum obtained at an interior $\bar{\xi} \in (0, 1)$. We assume $\xi < \bar{\xi}$ and think of a tighter constraint as, interchangeably, a higher ξ or a higher ϕ .

To capture this function in a convenient form, let $\hat{u}(a, \theta, p)$ denote the maximum obtained above, let

$$\tilde{u}(a, \theta, p) \equiv \beta \int \hat{u}(a, \theta', p) \varphi(\theta' | \theta) d\theta'$$

be the discounted, previous-period expectation of this maximum, and let $\tilde{c}_{it} \equiv c_{it} + z_{it}$. We can recast the household's problem as follows:

$$\max_{\{\tilde{c}_{it}, h_{it}, a_{it+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\tilde{c}_{it} - v(h_{it}) + \tilde{u}(a_{it+1}, \theta_{it}, p_{t+1}) \right) \right] \quad (7)$$

$$\text{subject to } \tilde{c}_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} \quad \text{and} \quad a_{it+1} \geq -\phi$$

It is therefore *as if* individual asset holdings entered the utility function.

For given θ and p , $\tilde{u}(a, \theta, p)$ is strictly increasing and concave in a if the next-period constraint is binding with positive probability, and constant otherwise. The following Euler condition is therefore necessary and sufficient for a_{it+1} to be optimal:

$$\tilde{u}_a(a_{it+1}, \theta_{it}, p_{t+1}) \leq q_t - \beta \equiv \pi_t, \quad (8)$$

with equality whenever $a_{it+1} > -\phi$.

Because \tilde{u}_a is equal to the expected value of the Lagrange multiplier on the morning collateral constraint, this condition means that each agent saves as much as it takes to equate the expected shadow value of collateral, or the return to liquidity, with the spread π_t , or the cost of liquidity. More succinctly, the aggregation of this condition across agents gives the aggregate demand for liquidity as a decreasing function of π_{t+1} . And because the aggregate supply of liquidity is given by b_{t+1} , we expect a higher b_{t+1} to map in equilibrium to a lower π_t , as well as to a more efficient allocation of the morning good and hence higher welfare. This intuition is incomplete because it does not take into account the endogeneity of p_{t+1} , but as the next proposition shows the essence remains the same.

Proposition 1. *There exist functions π, V and a scalar b_{bliiss} such that the following is true:*

(i) *For any policy path $\{\tau_t, b_{t+1}\}$, the equilibrium price of public debt is given by*

$$q_t = \beta + \pi(b_{t+1})$$

and welfare is given, up to a policy-invariant constant, by

$$\mathcal{W} = \sum_{t=0}^{\infty} \beta^t [c_t - v(h_t) + V(b_{t+1})]. \quad (9)$$

(ii) *For $b < b_{bliiss}$, $\pi(b) > 0$, $V(b) < V_{bliiss}$, and $V(b)$ is increasing in b ; and for $b \geq b_{bliiss}$, $\pi(b) = 0$ and $V(b) = V_{bliiss}$, where V_{bliiss} is the value of $\mathbb{E}[\theta u(x)]$ obtained at the first-best allocation of the morning good.*

In a nutshell, $V(b)$ captures how much public debt contributes to social welfare by easing the friction and improving the allocation of the morning good, whereas $\pi(b)$ measures the price that the typical

agent is willing to pay for the services provided. Finally, b_{bliss} identifies the “satiation” level of public debt, or aggregate collateral, above which the friction ceases to bind and, as a result, $V'(b) = \pi(b) = 0$.

For $b < b_{bliss}$, the private and the social value of liquidity are both positive—but they are not equal to each other, due to the pecuniary externality emerging from the dependence of the collateral constraint on the price of the morning good. This helps illustrate the potential robustness of our main lessons to pecuniary externalities a la Shleifer and Vishny (1992), Lorenzoni (2008) and Dávila (2015).⁵ At the same time, the existence of the satiation point, although neither realistic nor strictly needed, allows us to illustrate that the planner may optimally choose to *manufacture* a shortage of collateral, that is, to lead the economy below b_{bliss} even if it starts above it.

2.3 The reduced-form policy problem

The government’s problem consists of finding the sequence $\{c_t, h_t, \tau_t, q_t, b_{t+1}\}_{t=0}^{\infty}$ that maximizes (9) subject to the following four constraints:

$$q_t b_{t+1} = b_t + g - \tau_t A h_t \quad (10)$$

$$q_t = \beta + \pi(b_{t+1}) \quad (11)$$

$$v'(h_t) = (1 - \tau_t) A \quad (12)$$

$$c_t + g = A h_t \quad (13)$$

The first is the government’s budget constraint; the second is the bond pricing condition; the third is the labor supply condition; and the last one is the economy’s resource constraint.

This problem is equivalent to that of a representative-agent economy in which public debt generates a welfare flow of $V(b)$ and is priced at $q = \beta + \pi(b)$. The dependence of V and π on b epitomizes the dual role of public debt in easing the trading friction (the effect captured by V) and manipulating interest rates (the effect captured by π). How this dual role, in combination with the desire to smooth taxes, shapes the optimal policy is the subject of Section 3.

To ease that transition to that section, we further simplify the government’s problem as follows. Let $H(\tau) \equiv (v')^{-1}(\frac{1-\tau}{A})$ and $S(\tau) \equiv \tau A H(\tau)$ denote the equilibrium values of, respectively, labor supply and tax revenue, as functions of the tax rate. It is straightforward to check that S is single-peaked—i.e., there is a Laffer curve—and attains its maximum value, \bar{s} , at $\tau = \bar{\tau}$ for some $\bar{\tau} \in (0, 1)$. For any $s \leq \bar{s}$, the tax rate that raises revenue s is therefore given by $\tau = T(s) \equiv \min\{\tau : S(\tau) = s\}$. Next, let $U(s) \equiv A H(T(s)) - v(H(T(s)))$ measure the equilibrium utility from consumption and leisure as a function of s and note that this is decreasing and concave in s , reflecting the distortionary effect of taxation. The government’s problem can be re-expressed as follows:

⁵In Appendix B, we show that $\pi(b) > V'(b)$, i.e., there is a negative externality. Intuitively, when an agent decides how much to save, she does not internalize how her enhanced ability to buy the morning good will increase its price, tightening the others’ constraints. That said, our main results in Section 3 allow $\pi(b)$ and $V'(b)$ to be unequal in the opposite direction.

Proposition 2. *Let the functions V , π , and U be defined as above. The optimal policy path for taxes and public debt solves the following problem:*

$$\max_{\{s_t, b_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t [U(s_t) + V(b_{t+1})] \quad (14)$$

$$\text{subject to } q(b_{t+1})b_{t+1} = b_t + g - s_t \quad (15)$$

where $q(b) \equiv \beta + \pi(b)$.

2.4 Liquidity provision versus interest-rate suppression

To build some intuition about the key trade-off faced by the policymaker, let us momentarily consider a two-period version of the problem described in Proposition 2. Suppose further that the economy starts with zero debt and that any debt issued at $t = 1$ has to be retired at $t = 2$. Finally, abstract from optimal taxation and let some exogenous scalars λ_1 and λ_2 measure the value of tax revenue in periods $t = 1$ and $t = 2$, respectively. These are effectively the Lagrange multipliers on the respective budget constraints, except that they are treated as exogenous in the present exercise (but not in the main analysis).

Under these simplifications, the optimal debt issuance at $t = 1$ is given by

$$b^* = \arg \max_b \{ \lambda_1 q(b)b + V(b) - \beta \lambda_2 b \}. \quad (16)$$

The first term captures the benefit of relaxing the budget at $t = 1$. The second term captures the benefit of easing the financial friction at $t = 2$. The last term captures the tax burden of retiring the debt at $t = 2$.

Suppose now that $\lambda_1 = \lambda_2 = \lambda$. This amounts to imposing tax smoothing as in Barro (1979) and let us proxy within the present two-period exercise what goes on in the steady state of our infinite-horizon model. Using $q(b) = \beta - \pi(b)$, the problem can be restated as follows:

$$b^* = \arg \max_b \Omega(b, \lambda), \quad (17)$$

where

$$\Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b.$$

The first term captures the social value of the “liquidity services” of public debt, that is, the welfare gain from easing the financial friction. The second term captures the “profit” the government makes by providing these services.

This profit reminds seigniorage in monetary models. Here, it emerges because, and only because, there is a wedge between the interest rate the government has to pay on public debt and the underlying social discount rate. In the textbook paradigm (Barro, 1979; Lucas and Stokey, 1983), this wedge is zero, implying that the planner sees neither a profit nor a cost to issuing an extra unit of debt. In our setting, by contrast, the financial friction depresses the interest rate below the social discount rate. This in turn explains the precise sense in which public debt is “cheap” when rates are low.

Had the government cared only about this profit, it would have set $b = b_{\text{seig}} \equiv \arg \max_b \pi(b)b < b_{\text{bliss}}$. At the other extreme, had the government cared only about “social surplus,” as measured by $V(b)$, it would have set $b = b_{\text{bliss}}$. Here, the government strikes a balance between these two extremes, i.e., $b_{\text{seig}} < b^* < b_{\text{bliss}}$, because it values lowering financial frictions but also enjoying fiscal space. The stronger the fiscal preoccupation, as measured by λ , the closer b^* is to the “profit-maximizing” point b_{seig} .

In Appendix E, we use an example (with two types and log utility for the morning good) to obtain a closed-form solution for V and π and a sharper characterization of the determinants of b^* . We show, inter alia, that when borrowing constraints are relaxed, causing the premium π to fall, the planner may find it optimal to issue less public debt in order to moderate the fall in π . This anticipates our analysis in Section 5 of how the desire to increase fiscal space by manipulating the government’s cost of borrowing shapes the optimal policy response to shocks.

Clearly, such fiscal considerations are present only because taxation is distortionary. But it is important to recognize that distortionary taxation *alone* does not suffice to make $b^* < b_{\text{bliss}}$: had the interest rate on public debt been the same as the social discount rate, the government would have chosen to flood the economy with liquidity regardless of how large the welfare cost of taxation is. Formally, when $\pi = 0$, the terms $\lambda_1 q(b)b$ and $\beta \lambda_2 b$ in the objective of problem (16) cancel out, Ω reduces to V , and the government chooses $b = b_{\text{bliss}}$ regardless of λ .

The relevant trade off is therefore not liquidity provision versus distortionary taxation but rather liquidity provision versus interest-rate suppression: preserving the shortage of aggregate collateral makes sense only because it helps suppress the interest rate on public debt.

This insight is key for understanding the properties of optimal policy established in this paper. But while the two-period example has helped put the spotlight on the aforementioned two policy objectives (liquidity provision and interest-rate suppression), it has abstracted from the question of how these objectives interact with that of tax smoothing. It also cannot distinguish between the short and the long run, the transition between the two, and the optimal response to shocks. We address these limitations and offer a complete and precise characterization of the optimal policy in the rest of the paper.

3 Optimal Policy

In Appendix A we obtain the analogue to Proposition 2 in a model that has the financial friction distort the allocation of capital and thereby also reduce aggregate productivity, in the spirit of Kiyotaki and Moore (1997) and Holmström and Tirole (1998). While the precise channels via which public debt influences welfare and interest rates (or V and π) change, the essence of the policy problem remains the same. With this in mind, in this section we suppress the micro-foundations and focus on solving the reduced-form problem appearing in Proposition 2.

This task is complicated by the non-convexity of the problem. This complication cannot be simply

assumed away, for it is inherent to the interest-rate effects of public debt and therefore to the trade off we wish to study.⁶ To address this complication and complete the task at hand, we reformulate the policy problem in continuous time and adapt the methods of Skiba (1978).

3.1 Continuous-time formulation

Let $\bar{s} > 0$ be the maximal possible tax revenue, or the peak of the Laffer curve, let $\underline{s} \leq 0$ be an arbitrary lower bound, and let $\bar{b} \equiv \frac{\bar{s}-g}{\rho} > 0$. We henceforth consider the following continuous-time version of the problem obtained in Proposition 2.

Planner's Problem. Choose a path for (s, b) in $\mathcal{A} \equiv [\underline{s}, \bar{s}] \times [0, \bar{b}]$ so as to solve

$$\max \int_0^{+\infty} e^{-\rho t} [U(s) + V(b)] dt \quad (18)$$

$$\text{subject to } \dot{b} = R(b)b + g - s \quad \forall t \quad (19)$$

with initial condition $b(0) = b_0$, for some $b_0 \in [0, \bar{b}]$ and for $R(b) \equiv \rho - \pi(b)$.

We impose the following restrictions, which are consistent with but not limited to the micro-foundations presented in the previous section:

Main Assumptions. [A1] U , V , and π are continuously differentiable.

[A2] U is concave in s , with a maximum attained at $s = 0$.

[A3] There exists a threshold $b_{bliss} \in (0, \bar{b})$ such that $V'(b) > 0$ and $\pi(b) > 0$ for all $b < b_{bliss}$, and $V'(b) = 0$ and $\pi(b) = 0$ for all $b > b_{bliss}$.

[A4] $\pi(b) \leq \rho$ for all b .

A1 is technical. A2 and A3 mirror the properties established in our micro-founded setting. In particular, A2 means that the welfare cost of taxation is convex, while A3 captures the dual role of the financial friction on welfare and interest rates. A3 also imposes that the level of public debt that satiates the economy's demand for collateral is sustainable, a property that is not strictly needed but makes the analysis more interesting. Finally, A4 restricts the interest rate to be non-negative, an assumption that is not strictly needed but simplifies the exposition.

3.2 The Euler condition and the economics behind it

Denote the costate variable with λ and consider the Hamiltonian of the problem:

$$H(s, b, \lambda) \equiv U(s) + V(b) + \lambda [s - (\rho - \pi(b))b - g].$$

⁶Indeed, even if $V(b)$ happens to be concave, and even if $\pi(b)b$ is concave over the region $b \in (0, b_{bliss})$, non-convexity emerges from the fact that the "profit" $\pi(b)b$ switches from positive for $b < b_{bliss}$ to zero for $b > b_{bliss}$.

This can be rewritten as

$$H(s, b, \lambda) = U(s) + \lambda [s - \rho b - g] + \Omega(b, \lambda).$$

Similarly to Section 2.4, $\Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b$ measures the social value of the liquidity services of public debt plus the profit made from providing these services, and λ measures the shadow value of tax revenue. But whereas in that section we treated λ as exogenous and assumed it to be constant over time, here we let it evolve endogenously.

Given λ , the optimal s (equivalently, the optimal τ) solves $U'(s) + \lambda = 0$. Let $s(\lambda)$ denote the solution to this equation. The Euler condition can then be expressed as $\dot{\lambda} = \rho\lambda + H_b(b, \lambda, s(\lambda))$, or equivalently

$$\dot{\lambda} = \Omega_b(b, \lambda). \tag{20}$$

In a steady state, this condition reduces to $\Omega_b(b, \lambda) = 0$. This suggests that a steady state of our dynamic problem is akin to the solution of the static problem studied in Section 2.4. We will later verify that this intuition is correct, subject though to the following qualification: unless we strengthen our assumptions about π and V , there can exist multiple steady states, each one associated with a different λ and hence also with a different “static” solution $\operatorname{argmax}_b \Omega(b, \lambda)$.

Away from steady state, condition (20) equates $\dot{\lambda}$ with Ω_b . The former encapsulates the welfare cost of departing from tax smoothing; the latter captures the dual effect of the quantity of public debt on welfare and interest rates. Condition (20) therefore means that, along the transition to a steady state, the optimal policy balances not only the two objectives we emphasized before—namely liquidity provision and interest-rate suppression—but also the traditional objective of smoothing tax distortions over time.

Intuitively, when Ω_b is positive, there is value to increasing public debt, which means raising taxes tomorrow relative to today. (And the converse is true if Ω_b is negative.) If this tilt in the time profile of taxes were of no consequence for welfare, the government would move to the steady state instantaneously. The desire to smooth taxes therefore acts as some sort of an adjustment cost that slows down convergence to the steady state.

In fact, the desire to smooth taxes not only influences the rate of convergence to the relevant steady state but may also tie this steady state to initial conditions. We will see a sharp version of this point below in the form of a threshold for the initial level of debt such that the economy converges to a steady state without satiation if and only if it starts below this threshold.⁷

3.3 Steady state(s) and transitional dynamics

A caveat to some of the intuition provided above is that, like the Euler condition upon which they are based, they rely on local arguments. Such arguments are not only necessary but also sufficient for op-

⁷This discussion also underscores the intertwining of the optimal provision of liquidity and tax smoothing. As mentioned in the Introduction, this intertwining is a key high-level difference between our analysis and the literature on Friedman rule. We expand on this point in Appendix C.3.

tinality in convex optimization problems but not in our problem. We need *additional* arguments to identify the global optimum among multiple local optima that satisfy both the Euler condition and the transversality condition.

A formal treatment of this issue and a general solution are found in Appendix B. Here, we simplify the exposition by imposing the following restrictions on the problem's primitives.

Auxiliary Assumptions. [B1] *The ratio $V'(b)/\pi(b)$ is a constant ω .*

[B2] *The elasticity $\sigma(b) \equiv -\pi'(b)b/\pi(b)$ is increasing in $b \in (0, b_{bliss})$.*

[B3] *The level of government spending, g , is sufficiently large.*

B1 imposes that the wedge between the social and the private value of collateral is invariant to b ; this nests the special case in which the social and private value of liquidity coincide (i.e., $\pi = V'$). B2 guarantees that $\pi(b)b$ is single-peaked, achieving its maximum at some $b_{seign} \in (0, b_{bliss})$, that is, there is a Laffer curve in terms of b as well as in terms of τ . Finally, B3 guarantees that there is a sufficiently large shadow value for depressing the government's cost of borrowing. Together, these assumptions lead to the following sharp characterization of the optimal debt dynamics.⁸

Theorem 1. *There exists unique (b_{skiba}, b^*, s^*) , with $b_{seign} < b^* < b_{bliss} < b_{skiba}$ and $0 < s^* < \bar{s}$, such that:*

(i) *For $b_0 < b_{skiba}$, optimal debt and taxes converge monotonically to, respectively, b^* and s^* .*

(ii) *For $b_0 \geq b_{skiba}$, optimal debt and taxes stay constant at their initial levels.*

This result identifies b^* as the unique steady-state level of public debt below satiation, and b_{skiba} as the critical threshold for the initial level of debt below which the economy converges to this steady state. Above this threshold, the economy instead rests for ever at its initial point.⁹

A detailed proof of this result is provided in Appendix B. Here, we sketch out the main ideas with the help of the phase diagram in Figure 1.


To start with, consider the $\dot{b} = 0$ locus. This corresponds to balanced budget, or the value of λ (equivalently, the rate of taxation) that solves $s(\lambda) = g + (\rho - \pi(b))b$. It is illustrated by the curve labeled " $\dot{b} = 0$ " in the figure. This curve is upward sloping because a higher level of debt requires a higher rate of taxation (equivalently, a higher λ) for the budget to be balanced.¹⁰

⁸The micro-foundation of B1 and B2 is an open question. But since our main lessons survive without them, in the sense described later, we do not find it necessary to search for such micro-foundations. Also note that we do not restrict V to be concave because it would not have helped eliminate the non-convexity of the problem.

⁹This threshold is an example of the "Skiba points" that emerge in non-convex, optimal control problems. These points separate the state space in different sub-regions, each one corresponding to the basin of attraction of a different steady state. A peculiarity of the particular problem studied here is that every point above b_{skiba} is itself a steady state, mirroring the continuum of steady states in Barro (1979).

¹⁰To be precise, this monotonicity holds if and only if the cost of debt, $Rb = (\rho - \pi(b))b$, is increasing in b , which is necessarily the case in regions M and H of Figure 1 but may fail in region L. However, the result presented here are not unsettled by this possibility, because it is always optimal to move the economy outside region L and into region M.

Figure 1: Phase Diagram and the Optimal Path.



Next, consider the $\dot{\lambda} = 0$ locus. There are three scenarios to consider here, corresponding to the regions L, M and H in the figure.

In region L, which is defined by $b < b_{\text{seign}}$, increasing b raises both $\pi(b)b$ and $V(b)$, so there is no trade off between “liquidity provision” and “interest rate manipulation.” It follows that, for any $\lambda \geq 0$, the marginal value of raising debt is positive, $\Omega_b(b, \lambda) > 0$, and therefore also $\dot{\lambda} > 0$. That is, the locus $\dot{\lambda} = 0$ does not exist in region L.¹¹ By direct implication, there is also no steady state within this region.

In region M, which is defined by $b \in (b_{\text{seign}}, b_{\text{bliss}})$, increasing b raises $V(b)$ at the expense of reducing $\pi(b)b$, so the aforementioned trade off is now active. Which of the two sides of the trade off, liquidity provision or interest-rate suppression, dominates depends on how large the shadow value of government resources, λ , is. Holding b constant, a large enough λ tilts the balance in favor of interest-rate suppression and maps to $\dot{\lambda} = \Omega_b(b, \lambda) < 0$. Conversely, $\dot{\lambda} = \Omega_b(b, \lambda) > 0$ for λ small enough. By the same token, for any $b \in (b_{\text{seign}}, b_{\text{bliss}})$, there exists a critical value of $\lambda = \gamma(b)$ such that $\Omega_b(b, \gamma(b)) = 0$. This gives the curve labeled “ $\dot{\lambda} = 0$ ” in Figure 1.

This curve is decreasing, reflecting the idea that a higher λ shifts the balance in favor of interest-rate suppression (i.e., its brings $\text{argmax} \Omega$ closer to $\text{argmax} \pi(b)b$). The balanced-budget line, on the other hand, is increasing, reflecting the higher tax distortion implied by higher level of debt. It follows that the two lines intersect at a unique point $(b, \lambda) = (b^*, \lambda^*)$, which identifies the unique steady state within regions L and M. To the left of this point, debt and taxes increase over time; and to the right of it, debt and taxes fall over time.

Finally, consider region H, which corresponds to levels of debt above satiation. In this region, we have that both $V(b) = V_{\text{bliss}}$ and $\pi(b)b = 0$, so $\Omega_b(b, \lambda) = 0$ and $\dot{\lambda} = 0$ for all λ . That is, the locus of $\dot{\lambda} = 0$ is

¹¹This locus does not exist insofar as non-negative lump-sum transfers are allowed, for this restricts $\lambda \geq 0$. Otherwise, the $\dot{\lambda} = 0$ locus exists in the negative territory of the L region.

now the *entirety* of region H.

The last property may look peculiar but it actually epitomizes the optimality of tax smoothing in the textbook paradigm (Barro, 1979; Lucas and Stokey, 1983). In that model, both the liquidity-provision and the interest-rate concerns are absent, so $\dot{\lambda} = \Omega_b = 0$ over the entire phase diagram. Here, the same applies to the portion of the phase diagram above satiation.

This also explains why, in region H, there exist a continuum of *apparently* optimal steady states, corresponding to the segment of the “balanced budget” line inside that region. That is, for any $b_0 > b_{bliss}$, the policy plan that keeps debt and taxes constant for ever satisfies not only the budget constraint and the Euler condition, but also the transversality condition. However, for $b_0 \in (b_{skiba}, b_{bliss})$, this plan is actually dominated by an alternative plan, illustrated in Figure 1 by the segment of the saddle path that starts inside region H and enters into region M.

Along this plan, debt falls gradually, crossing b_{bliss} within finite time and converging asymptotically to b^* . Compared to the Barro-like plan of staying at $b = b_0$ for ever, this plan necessitates a departure from tax smoothing (higher taxes early, lower taxes later), which is costly. But it allows the government to extract a profit in terms of interest-rate suppression, once debt has fallen below b_{bliss} . Provided that this happens fast enough, which is the case if $b_0 \in (b_{bliss}, b_{skiba})$, the sacrifice in terms of tax smoothing is justified. The converse is true if $b_0 > b_{skiba}$.

4 Discussion

In this section we comment on the nature of optimal-long run quantity of public debt, the role played by the desire to smooth taxes, the robustness of our insights to the possibility that public debt crowds out (or, in) capital, and the importance of the assumption that debt is non-neutral.

4.1 The optimal long-run quantity of public debt

By the fact that (b^*, λ^*) is a steady state below satiation, $\Omega_b(b^*, \lambda^*) = 0$. By the Auxiliary Assumptions, $\Omega(b, \lambda)$ is concave in b over $[0, b_{bliss}]$ for any $\lambda \geq 0$. The following is then immediate:

Proposition 3. *Consider the steady state to which the economy converges whenever it starts with $b_0 < b_{skiba}$. In this steady state, the level of debt satisfies*

$$b^* = \arg \max_{b \in [0, b_{bliss}]} \{V(b) + \lambda^* \pi(b)b\} \quad (21)$$

where $\lambda^* = U'(s^*)$ and $s^* + \pi(b^*)b^* = g + \rho b^*$.

This verifies that the discussion in Section 2.4 provides the right intuition about the optimal steady-state level of debt, subject to two caveats: that we take into account the fixed-point relation between the debt level b^* and the weight λ^* that appears inside Ω ; and that this steady state is applicable only for some initial conditions, not all.

We have thus established the existence of a well-defined, stable, long-run target debt level, b^* , that falls short of the level that satiates the economy. And also that this level balances the value of easing the financial friction with the need to suppress interest rates. The former property contrasts with the textbook policy paradigm (Lucas and Stokey, 1983; Barro, 1979), where tax smoothing dictates that the long-run level of debt moves one-to-one with its initial level. The same property also allows us to study the optimal policy response to shocks (the topic of Section 5 below) using the transitional dynamics in the neighborhood of this steady state.

The second property underscores the value of preserving the financial friction so as to keep government borrowing cheap. This value is missing from the analyses of Woodford (1990) and Holmström and Tirole (1998), because they do not consider distortionary taxation. And although it is present in the environment of Aiyagari and McGrattan (1998), it is not properly accounted for, because their solution strategy—maximizing steady-state welfare subject to the steady-state budget—incorrectly treats the entire interest rate payments on public debt as a fiscal cost. By contrast, the correct planning problem ought to recognize that the component ρb of these interest rate payments is *not* a cost and, instead, debt issuance is a profit-generating business to the tune of $\pi(b)b$.¹²

4.2 Tax smoothing and the tripartite trade off

When $b_0 > b_{bliss}$, the economy starts from a position of collateral abundance and has no role for public debt. While we find this possibility of little practical interest, allowing for it in the model helps illustrate two broader points. First, that the planner may intentionally *manufacture* a shortage of collateral, for the sake of suppressing the interest rate on public debt. And second, that the desire to smooth taxes not only shapes the rate of convergence to the applicable long-run target for public debt but may also justify convergence to different long-run positions from different starting positions.¹³

Together, these points underscore how the optimal policy balances *three* objectives at once: the value of easing the financial friction; the need to contain interest-rate costs; and the desire to smooth taxes. It is this *tripartite* trade off that ultimately shapes *all* the properties of the optimal policy—including the steady state(s), the transitional dynamics, and the response to shocks.

4.3 Crowding out (or in) capital

Consider the following extension of our baseline model. In each afternoon, households have access to a technology of transforming the current consumption good into capital, which in turn can be used not

¹²See Appendix C.4 for a precise explanation of what the Aiyagari and McGrattan (1998) solution strategy is, of why it may miss to detect the existence of multiple steady states, and of why it mis-characterizes the steady state even when it is unique.

¹³This same logic also explains why, without the Auxiliary Assumptions introduced above, it is possible to have multiple steady states below satiation. Intuitively, the “adjustment cost” of a long-lasting departure from tax smoothing can justify remaining at one steady state when another, seemingly superior, steady state exists but is sufficiently far away (in terms of initial conditions). See Appendix B for a detailed treatment.

only for the production of the consumption good next afternoon but also as collateral next morning. This brings exactly two changes in the model. First, it changes the borrowing constraints (4)-(5) replacing a_{it} with $a_{it} + k_{it}$, where k_{it} denotes the amount of capital held in the beginning of period t . And second, it modifies the budget constraint (3) as follows:

$$c_{it} + p_t x_{it} + q_t a_{it+1} + k_{it+1} = a_{it} + f(k_{it}) + (1 - \tau_t) w_t h_{it} + p_t \bar{e}.$$

where f is a production function with $f' > 0 > f''$, $f(0) = 0$, and $f'(0) \geq 1/\beta$.

Because capital and the risk-free bond are equally good forms of collateral, the following arbitrage condition has to hold for all i and t :

$$f'(k_{i,t+1}) = \frac{1}{q_t} \equiv \frac{1}{\beta + \pi_t}.$$

It follows that all households choose the same amount of capital, and this amount is positively related to π_t . Intuitively, when the financial friction is more binding, there is an incentive to hold both more capital and bonds. By the same token, when the government issues more debt, it lowers π_t and crowds out capital, similarly to Aiyagari and McGrattan (1998). This however does not change the essence of the policy problem.

Proposition 4. *In the extension described above, public debt crowds out capital. Still, the reduced-form representation of the policy problem given in Proposition 1 continues to hold.*

The precise micro-foundations of V and π are now different, but neither their qualitative properties nor the implications for policy are affected. The fact that the issuance of public debt may crowd out capital is *not* an additional, separate element of the costs and benefits of debt issuance. It is merely a symptom of the role of public debt in easing the underlying financial friction.

Furthermore, this symptom can be turned upside down by letting the financial friction impact the production side of the economy, as in Kiyotaki and Moore (1997), Holmström and Tirole (1998). We offer an example of such a model in Appendix A. There, public debt can crowd *in* capital accumulation by easing the friction among firms, improving the cross-sectional allocation of capital and labor, and raising aggregate TFP. Furthermore, this crowding-in can be strong enough to offset the crowding-out of the higher taxes associated with higher levels of debt. Still, the basic trade-off we have emphasized—between the benefits of easing the friction and the desire to suppress interest rates—remains present.

These points underscore the likely robustness of our policy lessons to different micro-foundations of the role of public debt as buffer stock or collateral. A similar point applies to the substitutability of private assets and government bonds in this function. Our baseline model allowed for perfect substitutability. Assuming imperfect substitutability changes the magnitudes of V and π , but does not change the essence of the problem and does not upset the results.

4.4 On Ricardian Equivalence

What is essential for our results is the non-neutrality of public debt. To see this more clearly, modify our baseline model so that the private sector's pledgeable income moves one-to-one with future tax obligations. This preserves the financial friction but renders public debt neutral: any increase in aggregate collateral in the form of additional public debt issuance is perfectly offset by a commensurate reduction in pledgeable income. The same point applies to Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998): if borrowing constraints adjusted to future tax obligations, public debt would be neutral in those papers as well.

This, however, does not mean that the economy reduces to that in Barro (1979). Although π is now invariant to b , it is still positive (insofar as the financial friction binds). That is, the interest rate is still depressed, although insensitive to the level of public debt. Accordingly, the Euler condition gives $\dot{\lambda} = \pi\lambda > 0$, where $\pi > 0$ is fixed, and the following result obtains.

Proposition 5. *Suppose that the friction is present but public debt is neutral. Then, optimal taxes and debt exhibit a positive drift. In the long run, debt converges to \bar{b} , the highest sustainable level.*

This result is a “sanity test,” which further clarifies how our main lessons depend on the causal effect of public debt on liquidity premia and interest rates—a causal effect that is corroborated by the evidence in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood and Vayanos (2014).

5 Optimal Response to Shocks


We now study how the tripartite trade off between liquidity provision, interest-rate suppression and tax smoothing shapes the optimal response to shocks. Thanks to the (local) determinacy and stability of the steady state, this can be understood by studying the comparative dynamics in the phase diagram introduced in Section 3. For further illustration, we also use the numerical, non-linear solution of a stochastic extension of the original, discrete-time model from Section 2.¹⁴

5.1 Wars

Consider the comparative dynamics associated with an unexpected, once and for all, increase in g . They are illustrated in Figure 2. Prior to the change, the economy is assumed to be resting at the steady-state point b_{old}^* . The increase in g causes the $\dot{b} = 0$ locus to shift upwards, reflecting the increase in the taxes required for balanced budget. By contrast, the $\dot{\lambda} = 0$ locus does not move, because g does not enter the planner's Euler condition.

¹⁴Throughout this section, we let public debt be risk-free, as in Barro (1979) and Aiyagari et al. (2002). The opposite scenario, which allows public debt to be fully state-contingent, is considered in Appendix C.2. As in Lucas and Stokey (1983), this scenario allows the government to insure its budget against shocks; but now the optimal state-contingency balances such insurance with the objectives of providing liquidity and suppressing interest rates.

Figure 2: Permanent Increase in Government Spending



As a result, the steady-state level of debt drops from b_{old}^* to b_{new}^* and the optimal dynamic response is as follows: on impact, λ and the associated tax rate jump up from their old steady-state values to values that set the economy on a new saddle path; thereafter, debt and tax monotonically decrease towards the new steady state. Initially, taxes increase by more than the increase in g in order to allow debt to decrease.. But as both the level of debt and the interest rate on it fall, the government can eventually afford a lower increase in taxes than the increase in g .

Proposition 6. *An unanticipated permanent increase in g calls for an increase in taxes by more than one-to-one in the short run and by less than one-to-one in the long run.*


Compare this result to Barro (1979) or Aiyagari et al. (2002). There, the optimal response to a fiscal shock gives prominence to tax smoothing. Here, the optimal response deviates from tax smoothing in order to squeeze liquidity and allow the government to enjoy a profit by means of lower interest rates.

The same logic applies to transitory fiscal shocks, what is often referred to in the literature as “wars”. We illustrate this in Figure 3, using a stochastic example in which government spending follows a symmetric two-state Markov process, with the probability of staying in the same state equal to 0.9. The black lines correspond to our baseline model, the orange lines to its Barro/AMSS counterpart.

In Barro/AMSS, the war leaves a permanent mark on the level of debt and the rate of taxation, reflecting the unit-root property of that benchmark. Furthermore, the size of the tax response is simply the change in the annuity value of government spending. In our setting, by contrast, the debt level eventually reverts to its initial position, reflecting the determinacy of the long-run target level of debt. Finally, the accumulation of debt during the war is less pronounced than that in Barro/AMSS, because doing so permits the planner to moderate the increase in interest rates, which would have further tightened the budget.¹⁵ By the same token, the planner raises larger taxes during the war, but also enjoys lower taxes in the aftermath of the war.

¹⁵If the war is sufficiently persistent, this mechanism becomes so strong that the level of debt actually falls, as in the example with the permanent change discussed earlier.

Figure 3: Optimal Response to a War



5.2 Flight to Quality

Consider a shock that tightens the financial friction and raises the demand for public debt, without however affecting aggregate output, tax revenue, and the wedge between the private and social value of liquidity. Formally, let $V(b) = \theta \tilde{V}(b)$ and $\pi(b) = \theta \tilde{\pi}(b)$, for some fixed functions \tilde{V} and $\tilde{\pi}$, and consider an increase in θ . We think of this situation as a “flight to quality.”

Because this raises the social value of liquidity and the profit from interest-rate suppression in proportion to each other, it leaves the $\dot{\lambda} = 0$ locus unaffected. It follows that, *if* the $\dot{b} = 0$ locus had also been unaffected, the optimal policy response would have been to stay put. But $\dot{b} = 0$ actually shifts down, because the shock reduces the interest-rate costs on public debt, thus also reducing the taxes needed for balanced budget.

In a nutshell, a positive θ shock acts similarly to a negative g shock: the private sector’s flight to quality brings a bonanza for the government.

Proposition 7. *A unexpected permanent increase in θ causes an increase in the long-run level of public debt, and a front-loaded reduction in taxes.*

Of course, reality is more complicated than the scenario just described. A financial shock is likely to have additional, and possibly countervailing, effects on the government budget, such as shrinking the tax basis or necessitating a fiscal stimulus. Still, our insights provide a rationale for why financial shocks may justify larger deficits than other shocks. We further expand on this below.

Finally, our analysis qualifies the familiar intuition that an increase in the demand for liquidity calls for an increase in the government’s provision of it. The conventional intuition, though, fails to take into account how such a shock may also raise the marginal return to interest-rate suppression. This explains why, in our context, the optimal supply of liquidity would have *not* increased had it not been for the aforementioned “bonanza” effect and the resulting drop in λ .¹⁶


¹⁶Only a pure externality shock, which raises V without affecting π , is fully consistent with the aforementioned intuition.

5.3 Traditional vs Financial Recessions

Let us capture a “recession” as an exogenous shock to the labor wedge. This naturally leads to lower aggregate output and tax revenue, and an increase in the deficit.¹⁷ Next, let us distinguish between “traditional” and “financial” recessions as follows: the former leaves the functions π and V unaffected, the latter raises them by tightening the underlying financial constraints.

Figure 4 illustrates the optimal policy response to two such recessions of comparable size, in the sense that the exogenous shock to the labor wedge is the same in both cases. The difference is whether the shock comes together with an increase in π (black lines) or not (orange lines). The figure indicates that it is optimal to run a larger deficit in the former case. And yet, the higher deficits do not translate into faster debt accumulation. This is because the government is able to roll over its original debt at lower interest rates, as well as to pay less interest on newly issued debt. For the same reason, the government is also able to afford a larger optimal stimulus in the form of a larger “payroll tax cut.”¹⁸

Figure 4: IRFs to a Financial vs Traditional Recession



This provides a formal basis for the argument made by Paul Krugman, Brad DeLong and others that the reduction in the government’s cost of borrowing during a financial crisis calls for (makes it optimal) to run larger deficits. But it is important to emphasize the part of the statement that says “during a financial crisis”: what is key is not the variation in the observed interest rate *per se*, but rather the extent to which this represents variation in the wedge between that rate and the counterfactual rate that would have obtained in the absence of a financial friction. Were ρ to drop along side the interest rate, leaving the wedge, π , the same or smaller, public debt would *not* be cheaper.

Such a shock shifts the $\dot{\lambda} = 0$ curve to the right without shifting the $\dot{b} = 0$ curve. It therefore raises b^* , but for a different reason than that associated with a θ shock: the aforementioned “bonanza” effect is gone.

¹⁷Formally, we modify the model of Section 2 by letting the equilibrium condition for labor be $v'(n_t) = (1 - \tau_t)(1 + \omega_t)$, where ω_t is an exogenous shock. We then capture a recession as a transitory negative shock to ω_t . As usual, this proxies a negative demand shock in the New Keynesian model. A supply (productivity) shock has similar effects on output and tax revenue, but negates the need for a fiscal stimulus.

¹⁸Clearly, the same is true for government spending if we endogenize g and let the recession raise its marginal value.

6 Conclusion

We have studied optimal policy in a setting where public debt management helps not only smooth taxes (as in Barro, 1979; Lucas and Stokey, 1983) but also regulate the amount of collateral or liquidity (as in Woodford, 1990; Aiyagari and McGrattan, 1998; Holmström and Tirole, 1998). Issuing more public debt raises welfare by easing the underlying financial friction. But it also tightens the government budget by raising interest rates relative to the social discount rate.

This trade off creates the possibility that the government could optimally restrict the amount of liquidity in the market in order to keep the cost of debt finance low. It necessitates a departure from tax smoothing in the short run, so as to help attain an appropriate long-run level of debt. And it modifies the optimal response to shocks. In particular, it becomes optimal to run smaller deficits during wars, so as to contain the increase in interest rates; and larger deficits during financial crises, because such episodes are associated with cheap borrowing opportunities.

An obvious direction for future research is the quantification of the effects documented in this paper. The success of any such attempt would depend crucially on how sensitive the spread between the market price of public debt and the appropriate social discount factor is both to the state of the economy and to the quantity of public debt. The evidence in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood and Vayanos (2014) is suggestive of high sensitivity in both dimensions. That work, however, has focused on a different spread, that between government and high-grade corporate bonds, which is likely to be only imperfectly correlated with the spread that, at least under the prism of our analysis, is most relevant for optimal fiscal policy. We hope that these observations will guide future empirical and quantitative work on the topic.

Another interesting direction for future research is the interaction of fiscal and monetary policy in a sticky-price extension of our setting. We have in mind the following two issues in particular. During normal times, monetary policy could help create fiscal space by manipulating the real interest rate on public debt. And during a liquidity trap, public debt issuance could ease the zero lower bound constraint on monetary policy by providing liquidity and raising the underlying natural rate of interest. We hope that the tractability of the framework introduced in this paper will facilitate the exploration of these and other questions.

References

- Aiyagari, S. Rao, Albert Marcet, Thomas Sargent, and Jesus Seppälä.** (2002) "Optimal Taxation without State-Contingent Debt." *Journal of Political Economy*, 110(6): 1220–1254.
- Aiyagari, S. Rao, and Ellen McGrattan.** (1998) "The Optimum Quantity of Debt." *Journal of Monetary Economics*, 42(3): 447–469.

- Barro, Robert J.** (1979) "On the Determination of the Public Debt." *Journal of Political Economy*, 87(5): 940–971.
- Brock, William A., and W.D. Dechert.** (1983) "The Generalized Maximum Principle." University of Wisconsin mimeo.
- Dávila, Eduardo.** (2015) "Dissecting Fire Sales Externalities." New York University, Stern School of Business mimeo.
- Dechert, W.D., and K. Nishimura.** (1981) "A Note on Optimal Growth with a Convex-Concave Production Function." State University of New York at Buffalo Department of Economics Discussion Papers.
- Greenwood, Robin, and Dimitri Vayanos.** (2014) "Bond Supply and Excess Bond Returns." *The Review of Financial Studies*, 27(3): 663–713.
- Holmström, Bengt, and Jean Tirole.** (1998) "Private and Public Supply of Liquidity." *Journal of Political Economy*, 106(1): 1–40.
- Kiyotaki, Nobuhiro, and John Moore.** (1997) "Credit Cycles." *Journal of Political Economy*, 105(2): 211–248.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** (2012) "The Aggregate Demand for Treasury Debt." *Journal of Political Economy*, 120(2): 233–267.
- Lagos, Ricardo, and Randall Wright.** (2005) "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy*, 113(3): 463–484.
- Lorenzoni, Guido.** (2008) "Inefficient Credit Booms." *Review of Economic Studies*, 75(3): 809–833.
- Lucas, Robert E., Jr., and Nancy L. Stokey.** (1983) "Optimal Fiscal and Monetary Policy in an Economy without Capital." *Journal of Monetary Economics*, 12(1): 55–93.
- Shleifer, A., and R.W. Vishny.** (1992) "Liquidation Values and Debt Capacity: A Market Equilibrium Approach." *Journal of Finance*, 47(4): 1343–1366.
- Skiba, A. K.** (1978) "Optimal Growth with a Convex-Concave Production Function." *Econometrica*, 46(3): 527–539.
- Woodford, Michael.** (1990) "Public Debt as Private Liquidity." *American Economic Review*, 80(2): 382–388.

Online Appendices

A Variant Micro-foundations with Capital

In this appendix we present a variant model in which the financial friction impedes the allocation of capital across entrepreneurs, as opposed to the allocation of a good across consumers. This variant offers, not only an illustration of the broader applicability of the policy insights we developed in the main text, but also a bridge to the literature that emphasizes the role of collateral in the production side of the economy, as in Kiyotaki and Moore (1997) and Holmström and Tirole (1998).

There is only one good, which can be either consumed or converted into capital. There are no taste shocks and per-period utility is given by $c_{it} - v(h_{it})$, where c_{it} denotes consumption and h_{it} denotes labor supply. Each household comprises a “worker”, who supplies h_{it} in a competitive labor market, and an “entrepreneur”, who runs a private firm. The latter’s output is given by $y_{it} = \theta_{it} f(k_{it}, n_{it})$, where k_{it} is the firm’s capital input, n_{it} is the firm’s employment, and θ_{it} is an idiosyncratic productivity shock. $f(\cdot, \cdot)$ is strictly increasing and strictly concave.

Let κ_{it} denote the amount of capital owned by household i in the morning of period t . It is given by $\kappa_{it} = (1 - \delta)\kappa_{it-1} + l_{it-1}$, where δ denotes depreciation and l_{it-1} denotes last period’s saving. The firm’s input k_{it} can differ from κ_{it} insofar as entrepreneurs can rent capital from one another. Such trades are beneficial because κ_{it} is fixed prior to the realization of the current shocks, whereas k_{it} and n_{it} adjust ex post. In short, there are gains from reallocating capital.

Importantly, this reallocation is impeded by a financial friction. Let p_t denote the rental rate of capital. To use $k_{it} > \kappa_{it}$, the entrepreneur must borrow $z_{it} = p_t(k_{it} - \kappa_{it})$ in a short-term IOU market. As in the baseline model, he can do so by pledging ϕ and/or by posting his financial assets, a_{it} , as collateral. Moreover, he can use a fraction of the invested capital and/or the firm’s output as additional collateral. That is, the relevant constraint is

$$z_{it} \leq \phi + a_{it} + \xi_k k_{it} + \xi_y y_{it}$$

where $\xi_k, \xi_y \in (0, 1)$ are the fractions of invested capital and of anticipated income that can serve as collateral. Finally, the agent can also borrow in the afternoon, if he wishes so, but only subject to the constraint $a_{it+1} \leq \phi + \kappa_{it+1}$; that is, his net worth cannot fall below ϕ .

Relative to the baseline model, the model described above allows the quantity of public debt to enter the economy’s aggregate production function. In particular, by improving the allocation of capital, more aggregate collateral in the form of more public debt can map to higher aggregate TFP. Furthermore, public debt can have an ambiguous effect on capital accumulation. On the one hand, more public debt can crowd *in* capital via the aforementioned channel, namely by raising aggregate TFP and thereby the mean return to investment. On the other hand, more public debt can crowd *out* capital by offering a substitute form of collateral or buffer stock, as in Aiyagari and McGrattan (1998).

Notwithstanding these differences, the nature of the policy problem remains essentially the same. In particular, it can be shown that the following variant of Proposition 2 holds.

Proposition 8. *There exist functions $W, Q,$ and S such that the optimal policy path $\{\tau_t, b_{t+1}\}_{t=0}^{\infty}$ solves the following problem:*

$$\max_{\{\tau_t, b_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t) \quad (22)$$

$$s.t. \quad Q(\tau_{t+1}, b_{t+1})b_{t+1} = b_t + g - S(\tau_t, b_t) \quad (23)$$

Proof. See Appendix D. □

To relate this proposition to Proposition 2, note that $W, Q,$ and S capture, respectively, the per-period welfare flow, the market price of public debt and the tax revenue.¹⁹ As we move from the baseline model to the new model, the micro-foundations that underlie these objects change, and so do their functional forms. For instance, the two distortions now have non-separable effects on welfare, interest rates, and the tax base. Yet, the strategy for obtaining the desired representation remains the same: the key step is to define W as the welfare flow that obtains when the planner takes as given (τ_t, b_t) and optimizes over the set of the cross-sectional allocations of labor, capital, and asset holdings and the aggregate supplies of capital and labor; Q and S are then defined by, respectively, the interest rate that supports the best implementable allocation and the primary surplus induced by it. Importantly, the only reason why W, Q and S depend on b is that the latter controls the financial friction. The representation obtained therefore encapsulates, once again, the dual role of the financial distortion on welfare and the government budget. What is new relative to the baseline model is that the financial friction affects the budget, not only via interest rates, but also via the tax base: by interfering with the allocation of capital, it affects wages, income, and tax revenue for any given tax rate. However, neither this feature nor the details of the underlying micro-foundations need alter the properties of optimal policy.

In particular, consider the following continuous-time policy problem which is motivated by the preceding micro-foundations and which also nests the policy problem we studied before:

$$\max \int_0^{+\infty} e^{-\rho t} W(\tau, b) dt \quad (24)$$

$$s.t. \quad \dot{b} = [\rho - \pi(\tau, b)]b + g - S(\tau, b) \quad \forall t \quad (25)$$

$$b(0) = b_0 \quad (26)$$

Suppose that the functions W, S, π are continuously differentiable in both τ and b . Suppose further that there exists a function b_{bliss} such that $\rho > \pi(\tau, b) > 0$ and $W_b(b, \tau) > 0$ if $b < b_{bliss}(\tau)$, whereas $\pi(\tau, b) =$

¹⁹In the baseline model, we could express the policy problem in terms of s rather than τ because there was a one-to-one mapping between them. In the current model, this is no more true and, accordingly, we keep the tax rate as a control variable.

$W_b(b, \tau) = 0$ if $b \geq b_{bliss}(\tau)$; this allows for the possibility that the “satiation point” beyond which the friction ceases to bind may depend on the tax rate. Similarly, let $b_{seign}(\tau) \equiv \arg \max \{\pi(\tau, b)b + S(\tau, b)\}$; this is the analogue to the level of debt that maximized seigniorage in our baseline model, except that now we accommodate the possibility that the quantity of aggregate collateral affects the government budget, not only via the interest rate on public debt, but also via aggregate output and tax revenue. Adjusting the notion of “liquidity plus seigniorage” accordingly gives

$$\Omega(b, \lambda) \equiv \max_{\tau} \{W(\tau, b) + \lambda[\pi(\tau, b)b + S(\tau, b)]\}.$$

We can express the planner’s Euler condition as


$$\dot{\lambda} = \Gamma(b, \lambda) \equiv \Omega_b(b, \lambda),$$

which has exactly the same interpretation as its counterpart in our baseline model. Similarly, we can express the budget constraint as

$$\dot{b} = \Psi(b, \lambda),$$

where $\Psi(b, \lambda) \equiv [\rho - \pi(T(\lambda), b)]b - S(T(\lambda), b)$ and $T(\lambda) = \arg \max_{\tau} \{W(\tau, b) + \lambda[\pi(\tau, b)b + S(\tau, b)]\}$. We therefore obtain essentially the same ODE system as in our baseline model; the underlying micro-foundations and some details are different but the essence remains the same.

Figure 5: Entrepreneurial Model




We illustrate this in Figure 5. For this example, we assume that $[\pi(\tau, b)b + S(\tau, b)]$, is single-peaked in b . This guarantees that the phase diagram can be split in three regions, similar to regions L, M and H in Figure 1. The boundaries of these regions are now curved, rather than vertical, reflecting the fact that b_{seign} and b_{bliss} are allowed to vary with the rate of taxation and thereby with λ . Other than this difference, however, the analysis of the phase diagram remains intact: there is a unique steady state in

which the financial friction does not bind, and the economy converges to it for all initial $b_0 < b_{skiba}$, for some b_{skiba} .

Although we will not provide a complete characterization of the more general class of policy problems using this model, we hope to have conveyed the message that our insights are robust to different micro-foundations of the financial friction and of the liquidity-enhancing role of public debt.

We close this appendix by illustrating how the present model allows for public debt to crowd *in* capital, in contrast to Aiyagari and McGrattan (1998). This is done in Figure 6, for a particular parameterization of the model.

Figure 6: Crowding in or out



The left panel of Figure 6 considers the policy rule for aggregate capital. In particular, we consider two economies: one with a relatively low level of government spending ($g = 17\%$ of steady-state output); and another with a relatively high level of government spending ($g = 27\%$ of steady-state output) corresponding higher taxes in steady state. For each of these economies, we then show how the optimal amount of capital varies with the level of public debt, holding constant the tax rate at the respective steady-state level.²⁰ As can be seen from this panel, public debt crowds *in* capital. This is unlike Aiyagari and McGrattan (1998), because here public debt helps improve production efficiency and thereby raise the return to capital, which in turn encourages capital accumulation.

The right panel of Figure 6 shifts attention to the aggregate capital dynamics along the transition to steady state, starting from an initial level of debt below steady state. Along this transition, the increase in public debt crowds in capital by easing the underlying financial friction. But taxes increase in tandem with public debt, and this contributes in the opposite direction, by discouraging labor supply. It follows that capital could either increase or decrease along the transition to the steady state. But it is interesting to note that, as illustrated by the low- g scenario in the figure, it is *possible* that the crowding-in effect of public debt can dominate the crowding-out effect of taxes.

²⁰Both public debt and private capital are normalized by the steady-state level of output in the respective economy

B Characterization of Optimal Plan

In this Appendix we offer a complete, self-contained, characterization of the solution to problem (18)-(19). In particular:

- We show how to adapt the methods of Skiba (1978) to our setting so as to identify the truly optimal path among the many that satisfy the Euler and transversality conditions
- We fill in the details of the benchmark case considered in the main text. We show how Assumption B guarantees the existence of a unique steady state below satiation and prove Theorem 1.
- We show how, away from the aforementioned benchmark, it is possible to have multiple steady states below satiation, as well as no such steady state.
- We finally explain the precise sense in which the lessons obtained in the main text remain robust to the richer cases allowed here.

Also note that some of the results from this appendix are used in the proofs found in Appendix D.

B.1 The ODE system

As shown in the main text, the Hamiltonian of the planner's problem can be written as follows:

$$H(s, b, \lambda) = U(s) + \lambda [s - \rho b - g] + \Omega(b, \lambda),$$

where $\Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b$ measures the social value of the liquidity services of public debt plus the profit made from providing these services, and λ measures the shadow value of tax revenue. Throughout this Appendix, we are ruling out both lump-sum taxes and lump-sum transfers. This allows the possibility that $\lambda < 0$, or equivalently $s < 0$ and $\tau < 0$, which means the planner may be using a distortionary subsidy in order to accumulate debt fast enough.²¹

We now study the ODE system for b and λ implied by the budget constraint and the planner's Euler condition.

Consider first the budget constraint. This can be expressed as follows:

$$\dot{b} = \Psi(b, \lambda) \equiv g + (\rho - \pi(b))b - s(\lambda), \quad (27)$$

where $s(\lambda)$ denotes the optimal tax revenue. It is straightforward to check that $s(\lambda)$ is increasing in λ as the economy lies on the increasing branch of the Laffer curve and therefore that $\Psi(b, \lambda)$ is decreasing

²¹Had we allowed the planner to use lump-sum transfers, this possibility would not have emerged: the optimal policy would have achieved the same goal with a non-distortionary lump-sum transfer. This curtails the negative territory of the phase diagram (i.e., it restricts $\lambda \geq 0$) but does not otherwise affect the optimal dynamics.

in λ : a higher λ means higher taxes today, which in turn means lower debt tomorrow.²² By the Implicit Function Theorem, there exists a function $\psi : [\underline{b}, \bar{b}] \rightarrow \mathbb{R}_+$ such that $\Psi(b, \psi(b)) = 0$ for all b ; equivalently,

$$\dot{b} = 0 \text{ if and only if } \lambda = \psi(b).$$

The mapping $\psi(b)$ identifies the value of λ , or equivalently the tax rate, that balances the budget when the level of debt is b . Note that $\dot{b} < 0$ when $\lambda > \psi(b)$, that is, debt falls if taxes exceed the aforementioned level, and symmetrically $\dot{b} > 0$ if $\lambda < \psi(b)$. Finally, note that the function ψ satisfies the following properties.

Lemma 1. *ψ is continuous and strictly increasing, with $\psi(\underline{b}) = 0$ and $\lim_{b \rightarrow \bar{b}} \psi(b) = +\infty$.*

Proof. $\psi(b)$ is strictly increasing in b because higher debt requires higher taxes to balance the budget; $\psi(b)$ starts at zero when $b = \underline{b}$ because taxes are zero when the government has a large enough asset position to fully finance its spending using interest income received on its assets; and $\psi(b)$ diverges to $+\infty$ as b approaches \bar{b} because the shadow cost of taxation explodes as debt approaches the maximal sustainable level and, equivalently, the tax rate approaches the peak of the Laffer curve. \square

Consider next the Euler condition. As explained in the main text, this can be written as

$$\dot{\lambda} = \Omega_b(b, \lambda),$$

where $\Omega(b, \lambda) \equiv V(b) + \lambda\pi(b)b$. Equivalently,

$$\dot{\lambda} = v(b) - \lambda\pi(b)(\sigma(b) - 1). \quad (28)$$

where $v(b) \equiv V'(b)$ is the *social* marginal value of liquidity, $\pi(b)$ is the corresponding *private* value, or the liquidity premium, and

$$\sigma(b) \equiv -\frac{\pi'(b)b}{\pi(b)} \geq 0$$

is the elasticity of the liquidity premium with respect to the quantity of public debt.

As a reference point, consider momentarily the case in which public debt has no liquidity value, so that $v(b) = \pi(b) = 0$ for all b . Condition (28) then reduces to $\dot{\lambda} = 0$, which represents Barro's celebrated tax-smoothing result: when debt is priced at the social discount rate, λ is constant over time, and hence the optimal tax is also constant. Relative to this reference point, we see that whenever the right-hand-side of (28) is non-zero, optimality requires a non-zero drift in λ , that is, a deviation from tax smoothing.

Let $\Delta \equiv \{b \in [\underline{b}, b_{bliss}) : \sigma(b) \neq 1\}$ and define the function $\gamma : \Delta \rightarrow \mathbb{R}$ as follows:

$$\gamma(b) \equiv \frac{v(b)}{\pi(b)(\sigma(b) - 1)}.$$

²²Note also that $\Psi(b, \lambda)$ has a kink at $\lambda = 0$, because the corner solution $\tau = 0$ binds as λ crosses zero from below. Relaxing the lower bound on τ and/or introducing lump sum transfers would help speed up the accumulation of debt in situations in which $\lambda < 0$, but would not otherwise affect the results.

We can then restate the Euler condition (28) as follows:

$$\dot{\lambda} = \begin{cases} v(b) \left[1 - \frac{\lambda}{\gamma(b)} \right] & \text{if } b \in \Delta \\ 0 & \text{if } b \notin \Delta \end{cases} \quad (29)$$

By implication,

$$\dot{\lambda} = 0 \text{ if and only if } \begin{cases} \text{either } & b \in \Delta \text{ and } \lambda = \gamma(b) \\ \text{or } & b \notin \Delta \text{ and } \lambda \in \mathbb{R} \end{cases}$$

It follows that the graph of γ identifies the $\dot{\lambda} = 0$ locus over the region to the left of the satiation point (that is, for $b < b_{bliss}$). To the right of this point, we instead have $\dot{\lambda} = 0$ regardless of (λ, b) .

The graph of γ can be quite complicated, in part because there may exist multiple “holes” in the domain Δ , that is, multiple points at which $\sigma(b) = 1$. To interpret these points, note that

$$\frac{d[\pi(b)b]}{db} = \pi'(b)b + \pi(b) = -(\sigma(b) - 1)\pi(b). \quad (30)$$

It follows that the points at which $\sigma(b) = 1$ correspond to the critical points of the function $\pi(b)b$, which, as explained before, represents the rent, or the profit, that the government can make by falling short of satiating the economy’s demand for liquidity. With abuse of language, we henceforth refer to this rent as “seigniorage”. Next, note that $\pi(b)b$ is continuous over the closed interval $[0, b_{bliss}]$, it is zero at the boundaries of the interval, and is strictly positive in the interior of the interval. It follows that seigniorage attains a global maximum in the interior of that interval. In general, $\pi(b)b$ may admit an arbitrary number of local maxima and minima in addition to its global maximum. By the same token, σ may cross 1 multiple times. Note, however, that the derivative of $\pi(b)b$ crosses zero from above at any point that attains the global maximum, which in turn means that $\sigma(b)$ is necessarily increasing in an area around such a point.

B.2 The case studied in the main text

We now focus on a slightly more general case than the one studied in the main text—more specifically we dispense with Auxiliary Assumption **B3** and only maintain the following two assumptions

- B1.** the ratio v/π is constant;
- B2.** the elasticity σ is increasing in $b \in (0, b_{bliss})$.

The first assumption imposes that the wedge between the social and the private value of collateral is invariant to b , the second guarantees that $\pi(b)b$ is single-peaked and also extends the aforementioned local monotonicity of σ to its entire domain. In the sequel, we will refer to the peak in $\pi(b)b$ as b_{seign} . This peak satisfies $\pi(b_{seign}) + \pi'(b_{seign})b_{seign} = \pi(b_{seign})(1 - \sigma(b_{seign})) = 0$. An implication of **B2** is then that $\sigma(b) < 1$ for $b < b_{seign}$ and $\sigma(b) > 1$ for $b > b_{seign}$. Dispensing from **B3** will allow use to obtain a more general characterization of the cases implied by **B1** and **B2**.

Together, these assumptions lead to following characterization of the optimal debt dynamics.

Proposition 9. *Let Assumptions **B1** and **B2** hold. There exists a unique $b^* \in (\underline{b}, b_{bliss}]$ such that, for any initial point $b_0 < b_{bliss}$, the optimal level of public debt converges monotonically to b^* . Furthermore, $b^* < b_{bliss}$ if $g > \hat{g}$ and $b^* = b_{bliss}$ if $g < \hat{g}$, for some \hat{g} .*

This result identifies b^* as the steady state to which the economy converges from *any* initial point $b_0 < b_{bliss}$. It also relates b^* to the satiation point b_{bliss} . In particular, it shows that b^* is strictly lower than b_{bliss} if and only if g is high enough. Theorem 1 in the main text then follows directly from noting that Property **B3** in the main text is the same as $g > \hat{g}$ here. The rest of the section is dedicated to proving Proposition 9 in multiple steps, developing additional insights on the way. We start by noting that Property **B1** and **B2** imply the following structure for the function γ , which is instrumental for the subsequent analysis.

Lemma 2. *Let Assumptions **B1** and **B2** hold. The domain of γ is $\Delta = [\underline{b}, b_{seign}) \cup (b_{seign}, b_{bliss})$, where $b_{seign} \equiv \operatorname{argmax}_b \pi(b)b$. For $b \in [\underline{b}, b_{seign})$, γ is negatively valued and decreasing. For $b \in (b_{seign}, b_{bliss})$, γ is positively valued and decreasing. Finally, $\gamma(b) \rightarrow -\infty$ as $b \rightarrow b_{seign}$ from below and $\gamma(b) \rightarrow +\infty$ as $b \rightarrow b_{seign}$ from above.*

Proof. Recall that $b_{seign} = \operatorname{argmax}_b \pi(b)b$, so that b_{seign} solves $\pi(b)(1 - \sigma(b)) = 0$. Note that, as aforementioned, for b_{seign} to be a maximum, the following has to hold: $\pi(b)(1 - \sigma(b)) \geq 0$ for $b \leq b_{seign}$. From the definition of γ and the assumption $V'(b) \propto \pi(b)$, we have

$$\gamma(b) \propto \frac{1}{\pi(b)(\sigma(b) - 1)} \leq 0 \text{ for } b \leq b_{seign}$$

The latter result together with the definition of b_{seign} implies that $\lim_{b \uparrow b_{seign}} \gamma(b) = -\infty$ and $\lim_{b \downarrow b_{seign}} \gamma(b) = \infty$. Finally, as b increases above b_{seign} , $\pi(b)(1 - \sigma(b)) < 0$ and $\gamma(b) < \infty$. Together with the monotonicity of $\sigma(b)$, this implies that $\gamma(b)$ is decreasing over the domain $[\underline{b}, b_{bliss})$. \square

Recall that the graph of γ identifies the $\dot{\lambda} = 0$ locus in the region to the left of the satiation point, whereas the $\dot{b} = 0$ locus is given by the graph of ψ . By Lemma 1, ψ is positively valued and strictly increasing. Together with Lemma 2, this means that γ and ψ can intersect at most once. In particular, letting $\gamma_{bliss} \equiv \lim_{b \uparrow b_{bliss}} \gamma(b)$ and $\psi_{bliss} \equiv \psi(b_{bliss})$,²³ we have the following property.

Lemma 3. *Let Assumptions **B1** and **B2** hold. If $\gamma_{bliss} > \psi_{bliss}$, then γ and ψ never intersect. If instead $\gamma_{bliss} < \psi_{bliss}$, then γ and ψ intersect exactly once, and this intersection occurs at $b = b^*$, for some $b^* \in (b_{seign}, b_{bliss})$.*


²³Recall that γ is defined to the left of the satiation point but not *at* it, which explains why we write $\gamma_{bliss} \equiv \lim_{b \uparrow b_{bliss}} \gamma(b)$ rather than $\gamma_{bliss} \equiv \gamma(b_{bliss})$. Also, the existence of the limit follows from the property that, in the neighborhood of b_{bliss} , γ is decreasing and bounded from below by 0. Finally, note that this last property is true in general, not just in the special case under consideration.

Proof. From Lemma 2, we know that $\psi(b)$ and $\gamma(b)$ can only intersect in (b_{seign}, b_{bliiss}) . Given (i) the monotonicity of $\sigma(b)$ and hence $\gamma(b)$, (ii) the fact that $\psi(b)$ is increasing and (iii) $\lim_{b \downarrow b_{seign}} \gamma(b) = \infty$, $\gamma(b)$ and $\psi(b)$ can intersect at most once. If $\gamma_{bliiss} > \psi_{bliiss}$, (i) and (iii) imply that $\gamma(b)$ lies above $\psi(b)$ everywhere in $(b_{seign}, b_{bliiss}]$ and therefore they never intersect. In $\gamma_{bliiss} < \psi_{bliiss}$, (i)–(iii) imply that they intersect only once. \square

The two scenarios are illustrated in, respectively, Figures 7 and 8. The latter is the same as Figure 1 in the main text, reproduced here to ease the exposition.

Let us first consider Figure 7. The phase diagram is split in three regions: the region L, for $b < b_{seign}$; the region M, for $b \in (b_{seign}, b_{bliiss})$; and the region H, for $b > b_{bliiss}$. The dynamics of b are qualitatively similar across all three regions: $\dot{b} > 0$ below the graph of ψ and $\dot{b} < 0$ above it. By contrast, the dynamics of λ differ qualitatively across the three regions. In region L, γ is negatively valued; $\dot{\lambda} > 0$ above the graph of γ ; and $\dot{\lambda} < 0$ below it. In region M, the reverse is true: γ is positively valued; $\dot{\lambda} < 0$ above the graph of γ ; and $\dot{\lambda} > 0$ below it. Finally, in region H, γ is undefined and $\dot{\lambda} = 0$ throughout. These properties also hold true in Figure 8. What distinguishes the two figures is whether γ and ψ admit an intersection within region M. In Figure 7, they do not. This is because we have imposed $\gamma_{bliiss} > \psi_{bliiss}$, which together with the monotonicity of γ and ψ guarantees that γ lies above ψ throughout region M.

Figure 7: Benchmark, with $\psi_{bliiss} < \gamma_{bliiss}$.



What do these properties imply for the solution to the planner's problem? Since γ and ψ never intersect for the case depicted in Figure 7, the ODE system (27)–(29) admits no steady state to the left of the satiation point b_{bliiss} (regions L and M). By contrast, there is a continuum of such steady-state points to the right of the satiation point (region H): any point along the segment of ψ that lies to the right of b_{bliiss} trivially satisfies both $\dot{\lambda} = 0$ and $\dot{b} = 0$. Whether the planner finds it optimal to rest at such a point or move away from it—*i.e.*, whether these points correspond to a steady state of the optimal dynamics as

opposed to merely a fixed point of the ODE system—remains to be seen. For now, let us note that the lowest of these fixed points is associated with $b = b_{bliss}$ and $\lambda = \psi_{bliss} \equiv \psi(b_{bliss})$; the latter corresponds to the level of taxes that balances the budget when the economy rests at the satiation point.

For any $b_0 < b_{bliss}$, there exists a unique value of the costate, $\lambda_0 < \psi(b_0)$, such as the following is true: if the economy starts from (b_0, λ_0) and thereafter follows the dynamics dictated by (27)-(29), then, and only then, the economy converges asymptotically to $(b_{bliss}, \lambda_{bliss})$. In other words, there is a unique path that satisfies the planner's Euler condition and the budget constraint at all dates, and that eventually leads to satiation. This path is indicated with blue color in the figure.²⁴

The aforementioned path trivially satisfies the transversality condition, and is therefore a candidate for optimality. By contrast, any path that starts with $\lambda(0) > \lambda_0$ (higher taxes) and that follows the ODEs causes the level of debt to reach the lower bound \underline{b} in finite time; at this point, λ would have to jump down, violating the Euler condition, which means that this path cannot be optimal. Similarly, any path that starts with $\lambda(0) < \lambda_0$ (lower taxes) causes the level of debt to increase past the satiation point b_{bliss} and to reach the upper limit \bar{b} in finite time; at this point, λ would diverge to infinity and the transversality condition would be violated, which means that neither this path can be optimal.

Consequently, for any $b_0 < b_{bliss}$, the path that leads to satiation is the optimal path, and Proposition 9 applies with $b^* = b_{bliss}$. For any $b_0 \geq b_{bliss}$, the only candidate for optimality is the steady-state point associated with smoothing taxes and “staying put” at the initial level of debt: $(b, \lambda) = (b_0, \lambda_0)$ for all t , with $\lambda_0 = \psi(b_0)$.

Proposition 10. *Let Assumptions B1 and B2 hold and suppose $\psi_{bliss} < \gamma_{bliss}$. If $b_0 < b_{bliss}$, debt converges monotonically to b_{bliss} and taxes exhibit a positive drift along the transition. If instead $b_0 \geq b_{bliss}$, debt stays constant at b_0 for ever, and tax smoothing applies.*

Proof. Let us first consider $b_0 \geq b_{bliss}$. In this case, $V'(b) = \pi(b) = 0$ and the ODE system reduces to

$$\begin{aligned}\dot{b} &= \rho b - S(\lambda) \\ \dot{\lambda} &= 0\end{aligned}$$

implying that λ and hence the tax rate is perfectly smoothed, so that b stays put at b_0 . This is the celebrated Barro tax smoothing result.

Let us now consider $b_0 < b_{bliss}$. Let us first assume that $\gamma(b_{bliss}) > \psi(b_{bliss})$ and define $\lambda_{bliss} = \psi(b_{bliss})$. Using the fact that with satiation $\pi(b) = 0$, the approximate local dynamics around the satiation

²⁴One cannot rule out $\lambda_0 < 0$ for sufficiently low b_0 . When this is the case, the negative λ signals the high value that the planner attaches to issuing public debt. In fact, if it were feasible for b to jump, the planner would let b jump to the point where λ turns non-negative, and only thereafter we she follow the blue path in the figure. By the same token, if we allow the planner to make non-negative lump-sum transfers, these transfers will not affect the solution in the region where $\lambda > 0$, but would help speed up the accumulation of debt in the region where $\lambda < 0$.

point are given by

$$\dot{X}(t) = \mathbf{J}X(t) \text{ with } \mathbf{J} = \begin{pmatrix} \rho & -\frac{\rho}{\psi'(b_{bliss})} \\ V''(\bar{b}) - \lambda_{bliss}\pi'(b_{bliss})(\sigma(b_{bliss}) - 1) & 0 \end{pmatrix}$$

Note that $\text{Tr}(\mathbf{J}) = \rho > 0$ so that the two eigenvalues of \mathbf{J} sum up to a positive number. The determinant of \mathbf{J} is given by

$$\det(\mathbf{J}) = \frac{\rho}{\psi'(b_{bliss})} (V''(b_{bliss}) - \psi(b_{bliss})\pi'(b_{bliss})(\sigma(b_{bliss}) - 1))$$

By assumption, $\gamma(b_{bliss}) > \psi(b_{bliss})$, we have

$$\det(\mathbf{J}) < \frac{\rho}{\psi'(b_{bliss})} (V''(b_{bliss}) - \gamma(b_{bliss})\pi'(b_{bliss})(\sigma(b_{bliss}) - 1))$$

At b_{bliss} , both $V'(b)$ and $\pi(b)$ are zero, therefore $\gamma(b_{bliss})$ obtains from L'Hôpital's rule as

$$\lim_{b \rightarrow b_{bliss}} \gamma(b) = \frac{V''(b_{bliss})}{\pi'(b_{bliss})(\sigma(b_{bliss}) - 1)}$$

implying that $\det(\mathbf{J}) < 0$. Furthermore, the discriminant of the polynomial associated with the eigenvalue problem is strictly positive, $\Delta = \rho^2 - 4 \det(\mathbf{J}) > 0$. Taken together, these results imply that the two eigenvalues are real, add up to a positive number and are of opposite sign. The local dynamics around the point $(b_{bliss}, \lambda_{bliss})$ therefore satisfy a saddle path property. It is also easy to show that the eigenvector associated to the stable eigenvalue is given by

$$\mathbf{v} = \left(\frac{\rho}{\psi'(b_{bliss})}, \frac{\rho + \sqrt{\Delta}}{2} \right)$$


and is not degenerate as $\psi'(b) > 0$. In other words, starting from $b(0) \in \{b_{bliss} - \varepsilon; \varepsilon > 0\}$, there exists a unique path taking the economy to satiation. This establishes the first part of the proposition.

Let us now consider a situation where $\gamma(b_{bliss}) < \psi(b_{bliss})$. In this case, the inequality established for the determinant of \mathbf{J} is reversed and $\det(\mathbf{J}) > 0$. The two eigenvalues have the same sign and sum up to a positive number, and are therefore positive. $(b_{bliss}, \lambda_{bliss})$ is not locally stable and starting from $b < b_{bliss}$, there exists no path leading the economy towards it. \square

Let us now consider Figure 8. In this case, γ and ψ intersect exactly once, at $b = b^* \in (b_{seign}, b_{bliss})$. Let $\lambda^* \equiv \psi(b^*)$ denote the shadow cost of taxation associated with balancing the budget when $b = b^*$. By construction, the pair (b^*, λ^*) identifies the unique steady state of the ODE system (27)-(29) to the left of the satiation point (i.e., within regions L and M). As is clear from the figure, this steady state is saddle-path stable. In particular, for any $b_0 < b_{bliss}$, we can find a continuous path that satisfies conditions (27)-(29) and that asymptotically converges to (b^*, λ^*) . Exactly the same arguments as in Figure 7 guarantee that this path is the unique candidate for optimality, and hence also the optimal path, as long as $b_0 < b_{bliss}$.

A crucial difference from the case in Figure 7 is that the economy now converges to a steady state characterized by a debt level that is strictly lower than the satiation level: Proposition 9 applies with $b^* <$

Figure 8: Benchmark, with $\psi_{bliiss} > \gamma_{bliiss}$.



b_{bliiss} . Consequently, the sign of the drift in debt and taxes now depends on the initial position: if $b_0 < b^*$, then debt and taxes increase monotonically over time, whereas the converse is true if $b_0 \in (b^*, b_{bliiss})$.

Another important difference concerns the behavior of the system in the region to the right of the satiation point. In the previous case, the Barro-like plan of keeping taxes and debt constant over time was the *unique* candidate for optimality throughout region H, that is, for all $b_0 > b_{bliiss}$. This is no longer true. Instead, as it is evident in the figure, for any $b_0 \in [b_{bliiss}, b_{skiba})$, there is an additional candidate for optimality: the path indicated with blue color in the figure.

This path lets b fall over time, crossing b_{bliiss} in finite time and asymptotically converging to b^* . Accordingly, the economy goes through two phases. In the first phase, which is defined by the time interval over which b remains above b_{bliiss} , λ stays constant over time, which means that tax smoothing applies. Although this resembles Barro (1979), there is a key difference: the constant value of λ exceeds $\psi(b)$ throughout this phase, which means that taxes are smoothed at a level that is higher than what is required for balancing the budget (in turn explaining why debt falls over time). In the second phase, which starts as soon as b has crossed b_{bliiss} from above, debt continues to fall, but tax smoothing no longer holds, for the reasons explained earlier on.

By construction, the path described above satisfies the ODE system (27)-(29) at all t and asymptotically converges to (b^*, λ^*) , which means that it also satisfies the transversality condition. This verifies that, as long as it exists, this path is a candidate for optimality. But so is the Barro-like plan of “staying put” at the point of the graph of ψ that corresponds to the initial level of debt, that is, at $(b, \lambda) = (b_0, \lambda_0)$ with $\lambda_0 = \psi(b_0)$. *How can we tell which path is better?*

To address this question, we use an elementary but powerful result from optimal-control theory. Below, we first state the result, which holds true for any configuration of the planner’s problem. We then

use it to complete the characterization of the particular benchmark under consideration.

For any b_0 , let $\mathcal{P}(b_0)$ be the set of all the paths for (b, λ) that start from b_0 , satisfy the ODE system in all t , and also satisfy the transversality condition at infinity. Since these conditions are necessary for optimality, the optimal path is necessarily contained in $\mathcal{P}(b_0)$. More generally, we can reduce the planner's problem to that of choosing a path $\mathcal{P}(b_0)$. Next, note that any path in $\mathcal{P}(b_0)$ is associated with a different initial value for the costate and let $\Lambda(b_0)$ be the set of such initial values for the costate. Choosing a path in $\mathcal{P}(b_0)$ is therefore equivalent to choosing an initial value λ_0 in $\Lambda(b_0)$. The following result is helpful for evaluating the welfare associated with any candidate path.

Lemma 4 (Skiba, 1978, Brock and Dechert, 1983). *For any b_0 and any $\lambda_0 \in \Lambda(b_0)$, the path in $\mathcal{P}(b_0)$ that starts from initial point (b_0, λ_0) yields a value that is equal to $\mathcal{H}(b_0, \lambda_0) / \rho$.*

Proof. See Brock and Dechert (1983). □

For any given b_0 , the above result allows one to rank the candidate paths in $\mathcal{P}(b_0)$ by simply inspecting how the value of the Hamiltonian, $\mathcal{H}(b_0, \lambda_0)$, varies as λ_0 varies within the set $\Lambda(b_0)$. But now note that $\mathcal{H}(b, \lambda)$ is strictly convex in λ , as it is defined as the upper envelop of functions that are linear in λ . It follows that, whenever $\mathcal{P}(b_0)$ is not a singleton, the optimal path is necessarily the path that starts with λ_0 either at the maximal or the minimal value inside $\Lambda(b_0)$. This property is instrumental for identifying the optimal path starting from any given initial level of debt, not only in the benchmark under consideration, but also in the more general case studied later.

Let us now go back to Figure 8. Pick any $b_0 \geq b_{bliss}$ and suppose there exists a continuous path that satisfies the ODEs and asymptotically converges to b^* . As already noted, this path is a candidate for optimality. But so is the Barro-like plan that keeps b and λ constant for ever at, respectively, b_0 and $\psi(b_0)$. Note, next, that the first plan is associated with a higher λ_0 (i.e., higher taxes) than the second, because the first runs a surplus whereas the second balances the budget. Finally, note that, along any candidate path, $\mathcal{H}_\lambda(b, \lambda) = \dot{b}$. For the path that leads the economy to b^* , we have that $\dot{b} < 0$ at $t = 0$, and hence $\mathcal{H}_\lambda(b_0, \lambda_0) < 0$. For the Barro-like plan, instead, $\dot{b} = 0$ and hence $\mathcal{H}_\lambda(b_0, \lambda_0) = 0$. Since \mathcal{H} is convex, this means that the Barro-like plan attains the minimum of \mathcal{H} over the set of candidate paths. It follows that, whenever the path that takes the economy to b^* exists, this path strictly dominates the Barro-like, and it is the optimal one.

The preceding argument *supposes* the existence of such a path. Whether such a path exists or not depends on the initial level of debt, b_0 . In the figure, it is evident that this is the case if and only if b_0 is lower than the threshold b_{skiba} . *But how is this threshold defined in the first place, and what guarantees its own existence?*

Consider $b_0 = b_{bliss}$. If we initiate the ODE system with a starting value $\lambda(0)$ slightly above $\psi_{bliss} = \psi(b_{bliss})$, which means that we run a sufficiently small enough surplus, then the resulting path for b never reaches b^* . By contrast, if we start with $\lambda(0)$ far above $\psi(b_{bliss})$, debt falls below b^* in finite time.

Finally, note the path of b induced by the ODE system is continuous and monotonic in $\lambda(0)$. It follows that there exists a critical value $\lambda_{skiba} \in (\psi_{bliiss}, \infty)$ such that, if we start with $\lambda(0) = \lambda_{skiba}$, then and only then the economy converges asymptotically to b^* .

By continuity, this kind of path also exists for b_0 above but close enough to b_{bliiss} . Furthermore, because the planner's Euler condition dictates $\dot{\lambda} = 0$ (tax smoothing) throughout region H, the plan under consideration keeps λ constant as long as b is above b_{bliiss} . It follows that the portion of this path that is to the right of the satiation point is flat at the level λ_{skiba} .

Define next $b_{skiba} \in (b_{bliiss}, \bar{b})$ as the level of debt that balances the budget when taxes are set at the level corresponding to λ_{skiba} ; that is, $b_{skiba} \equiv \psi^{-1}(\lambda_{skiba})$. Note that ψ is continuous and monotone, $\lambda_{bliiss} > \psi(b_{bliiss})$, and $\lim_{b \rightarrow \bar{b}} \psi(b) = \infty$; this verifies that b_{skiba} exists and is necessarily strictly between b_{bliiss} and \bar{b} . It is then immediate that a continuous path that satisfies the ODEs and that converges to b^* exists if and only if $b_0 < b_{skiba}$, as illustrated in the figure.

We thus have the following complement to Proposition 10.

Proposition 11. *Let Assumptions **B1** and **B2** hold and suppose $\psi_{bliiss} > \gamma_{bliiss}$. There exist unique points $b^* \in (b_{seign}, b_{bliiss})$ and $b_{skiba} \in (b_{bliiss}, \bar{b})$ such as the optimal debt level converges monotonically to b^* if $b_0 < b_{skiba}$, whereas it stays constant at b_0 for ever if $b_0 \geq b_{skiba}$. Optimal taxes exhibit a positive drift as long as $b \in (b_{seign}, b^*)$, a negative drift as long as $b \in (b^*, b_{bliiss})$, and are smoothed as long as $b > b_{bliiss}$.*

Proof. The discussion preceding the proposition in the main text establishes the existence of b_{skiba} by using a continuity argument. Here we analyze the stability of the steady state (b^*, λ^*) .

The linear approximation of the system of the ODEs around a stationary point (b^*, λ^*) is given by

$$\dot{X}(t) = \begin{pmatrix} \rho + \bar{\omega} V'(b^*)(\sigma(b^*) - 1) & -S'(\lambda^*) \\ V''(b^*) - \lambda^* \bar{\omega} V''(b^*)(\sigma(b^*) - 1) - \lambda^* \bar{\omega} V'(b^*)\sigma'(b^*) & -\bar{\omega} V'(b^*)(\sigma(b^*) - 1) \end{pmatrix} X(t) = \mathbf{J}X(t)$$

where $\bar{\omega} \equiv \pi(b)/V'(b)$ and $X(t) \equiv (b(t) - b^*, \lambda(t) - \lambda^*)'$. Using the definitions of the functions $\psi(b)$, $\gamma(b)$ and their respective derivatives, the matrix \mathbf{J} , evaluated at (b^*, λ^*) , is

$$\mathbf{J} = \begin{pmatrix} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) & -\frac{1}{\psi'(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \\ \gamma'(b^*) \frac{V'(b^*)}{\gamma(b^*)} & -\frac{V'(b^*)}{\gamma(b^*)} \end{pmatrix}$$

First note that the trace of matrix \mathbf{J} is given by $\rho > 0$, implying that the two eigenvalues of \mathbf{J} sum up to a positive number. The determinant of the \mathbf{J} matrix, evaluated at (b^*, λ^*) , is

$$\det(\mathbf{J}) = \frac{V'(b^*)}{\gamma(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \left(\frac{\gamma'(b^*)}{\psi'(b^*)} - 1 \right)$$

Given that $b^* < b_{bliiss}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ and $V'(b^*) > 0$. Finally, from Lemma 2, we know that $\gamma'(b) < 0$ for $b \in (b_{seign}, b_{bliiss}]$. Therefore, given that $\psi'(b) > 0$, $\det(\mathbf{J}) < 0$ and hence the two eigenvalues are distributed around 0. Therefore, (b^*, λ^*) a saddle path stable.

Note that, the stable root of the system is given by

$$\mu = \frac{\rho - \sqrt{\Delta}}{2}$$

where $\Delta = \rho^2 - 4 \frac{V'(b^*)}{\gamma(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \left(\frac{\gamma'(b^*)}{\psi'(b^*)} - 1 \right) > 0$ is the discriminant of the polynomial. Hence the eigenvector, (v_1, v_2) , associated to this eigenvalue satisfies

$$\left(\frac{\rho}{2} + \frac{V'(b^*)}{\gamma(b^*)} + \frac{\sqrt{\Delta}}{2} \right) v_1 - S'(\lambda^*) v_2 = 0$$

Consider the eigenvector is $\left(S'(\lambda^*), \frac{\rho}{2} + \frac{V'(b^*)}{\gamma(b^*)} + \frac{\sqrt{\Delta}}{2} \right)$. Given that $V'(b^*) > 0$, $\gamma(b^*) > 0$ (since $\sigma(b) > 1$) and $S'(\lambda^*) > 0$ in the upward sloping part of the Laffer curve, both components of the vector are positive. The co-movement result follows: For any $\varepsilon > 0$, starting from $b_0 = b^* - \varepsilon$ (resp. $b_0 = b^* + \varepsilon$), the economy will converge to (b^*, λ^*) increasing (resp. decreasing) both debt and taxes along the transition path. \square

For practical purposes, we think it is appropriate to restrict $b_0 < b_{bliss}$, so that the financial distortion is present in the initial period. Under this restriction, the combination of Propositions 10 and 11 generates the following two key lessons.

The first lesson is that the economy can belong in one of two classes. In the one, debt converges to b_{bliss} , which means that the planner extinguishes the financial distortion in the long run. In the other class, the opposite is true: the planner preserves the financial distortion in the long run. We will study below whether and how this taxonomy extends to the general case. For now, we wish to emphasize that both classes feature a deviation from tax smoothing along the transition.

The second lesson is that the condition $\psi_{bliss} > \gamma_{bliss}$ is *both* necessary and sufficient for an economy to belong in the second of the aforementioned two classes. In order to derive an interpretation of this condition recall that $\psi(b)$ measures the value of λ implied by balancing the budget; that $\gamma(b)$ identifies the value of λ that balances the planner's conflicting objectives: when $\lambda > \gamma(b)$, then and only then the value the planner attaches to interest-rate manipulation (or seigniorage) outweighs the value of collateral creation (or liquidity provision); and finally that $\psi_{bliss} \equiv \psi(b_{bliss})$ and $\gamma_{bliss} \equiv \lim_{b \downarrow b_{bliss}} \gamma(b)$. It follows that $\psi_{bliss} > \gamma_{bliss}$ if and only if $\Omega_b(b, \lambda) < 0$ for (b, λ) close enough to $(b_{bliss}, \psi(b_{bliss}))$, which leads to the following simple interpretation.

Fact 1. $\psi_{bliss} > \gamma_{bliss}$ if and only if, in the neighborhood of b_{bliss} , the benefit of relaxing the government budget by depressing the interest rate on public debt exceeds the cost of the financial distortion.

The proof of Proposition 9 is then completed by noting that $\psi_{bliss} > \gamma_{bliss}$ if and only if g is high enough, a property that holds even outside our benchmark and that is proved in Lemma 5 below.

But: *Do the lessons obtained above apply outside the benchmark under consideration?* We address this question next.

B.3 Beyond the Benchmark

The benchmark studied above has two key properties: $\pi(b)b$ is singled-peaked, so that the phase diagram can be organized in the three regions described above; and γ is decreasing over the region M, so that it can intersect at most once with ψ . If we modified the benchmark by allowing either for a non-monotone σ or for $V' \neq \pi$ but maintained the aforementioned properties, then the preceding arguments go through and Propositions 10 and 11 continue to hold.

What if the aforementioned properties do not hold, as it may be the case for certain micro-foundations? There is a plethora of possibilities. To make progress, we will continue for a moment to assume that $\pi(b)b$ is single-peaked, which preserves the tripartite structure of the phase diagram, but will let $\gamma(b)$ be non-monotone over region M.²⁵ In this case, the graphs of γ and ψ may intersect multiple times. Clearly, any such intersection identifies a steady-state point of the ODE system. *What are the local dynamics around each of these points? Starting from a given initial b_0 , how many paths are candidates for optimality? And what are the properties of the optimal path?*

There is a multitude of possible answers to these questions. To illustrate, consider the case in which γ and ψ happen to intersect three times, giving rise to three steady-state points for the ODE system within region M. Figures 9, 10 and 11 below illustrate three phase diagrams that are consistent with this case. The three diagrams feature similar configurations of the γ and ψ functions and similar local dynamics around each of the three steady states, but different global dynamics and different types of optimal policies. We go over each of these three possibilities one by one.

Consider Figure 9. In order to simplify the exposition, we truncate region L, where $b < b_{seign}$, γ is negatively valued, and there can be no steady state; we thus focus on region M, where $b \in (b_{seign}, b_{bliss})$ and where γ and ψ intersect three times. Denote the level of debt at the three intersection points by b_L^* , b_M^* , and b_H^* (for, respectively, “low”, “medium”, and “high”). Because γ goes to infinity in the neighborhood of b_{seign} , we know that γ must intersect ψ from above at b_L^* and b_H^* , and from below at b_M^* . This is useful to note, because, as shown in the next proposition, the relation between the slope of γ and that of ψ dictates the local stability properties of the ODE system around any steady state.

Proposition 12. *Consider any (b^*, λ^*) such that $\lambda^* = \gamma(b^*) = \psi(b^*)$, that is any steady-state point of the ODE system in the region to the left of the satiation point. There exists a finite scalar $\chi > 0$ such that the local dynamics around that steady-state point are*

- (i) *saddle-path stable if $\gamma'(b^*) < \psi'(b^*)$;*
- (ii) *explosive with real eigenvalues if $\psi'(b^*) < \gamma'(b^*) < \psi'(b^*) + \chi$;*
- (iii) *explosive with imaginary eigenvalues (i.e. with cycles) if $\gamma'(b^*) > \psi'(b^*) + \chi$.*

²⁵Recall that γ is necessarily decreasing in a neighborhood to the right of b_{seign} , because $\sigma(b) \downarrow 1$ and $\gamma(b) \uparrow \infty$ as $b \downarrow b_{bliss}$. Allowing for a non-monotone γ therefore means that γ is increasing over a portion of region M. This in turn can happen when the elasticity σ and/or that the ratio π/V' is decreasing over a subset of (b_{seign}, b_{bliss}) .

Proof. The linear approximation of the system of the ODEs around a stationary point $(b^\#, \lambda^\#)$ is given by

$$\dot{X}(t) = \begin{pmatrix} \rho + \pi(b^*)(\sigma(b^*) - 1) & -S'(\lambda^*) \\ V''(b^*) - \lambda^* \pi'(b^*)(\sigma(b^*) - 1) - \lambda^* \pi(b^*)\sigma'(b^*) & -\pi(b^*)(\sigma(b^*) - 1) \end{pmatrix} X(t) = \mathbf{J}X(t)$$

with $X(t) = (b(t) - b^*, \lambda(t) - \lambda^*)'$. Using the definitions of the functions $\psi(b)$, $\gamma(b)$, $\psi'(b)$ and $\gamma'(b)$, we can rewrite the matrix \mathbf{J} , evaluated at (b^*, λ^*) as

$$\mathbf{J} = \begin{pmatrix} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) & -\frac{1}{\psi'(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \\ \gamma'(b^*) \frac{V'(b^*)}{\gamma(b^*)} & -\frac{V'(b^*)}{\gamma(b^*)} \end{pmatrix}$$

First note that the trace of matrix \mathbf{J} is given by $\rho > 0$, implying that the two eigenvalues of \mathbf{J} sum up to a positive number. The determinant of the \mathbf{J} matrix, evaluated at (b^*, λ^*) , is

$$\det(\mathbf{J}) = \frac{V'(b^*)}{\gamma(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \left(\frac{\gamma'(b^*)}{\psi'(b^*)} - 1 \right)$$

Given that $b^* < b_{bliss}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ and $V'(b^*) > 0$. Therefore, the position of $\gamma'(b^*)/\psi'(b^*)$ with respect to 1 determines the sign of the determinant, and hence the position of the two eigenvalues around 0. Note that a steady state only exists in regions where $\sigma(b^*) > 1$ and hence $\gamma(b^*) > 0$. When $\gamma'(b^*) < \psi'(b^*)$, $\det(\mathbf{J}) < 0$ and hence the two eigenvalues are distributed around 0. Therefore, a saddle path exists (recall that $\text{Tr}(\mathbf{J}) = \rho > 0$), hence proving the first statement. In the opposite situation the two eigenvalues have positive real part, hence establishing the explosiveness part of the proposition.

The emergence of cycles is related to the real vs complex nature of the eigenvalues. This is established by looking at the discriminant, Δ , of the characteristic polynomial:

$$\Delta = (\text{Tr}\mathbf{J})^2 - 4 \det \mathbf{J} = \rho^2 - 4 \frac{V'(b^*)}{\gamma(b^*)} \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \left(\frac{\gamma'(b^*)}{\psi'(b^*)} - 1 \right)$$

The two roots are complex if the discriminant is negative

$$\Delta < 0 \iff \gamma'(b^*) > \psi'(b^*) + \chi \text{ with } \chi \equiv \frac{\rho^2 \psi'(b^*)}{4 \left(\rho + \frac{V'(b^*)}{\gamma(b^*)} \right) \frac{V'(b^*)}{\gamma(b^*)}}$$


Therefore establishing the condition for the emergence of complex vs real explosive eigenvalues. \square

This result restricts the *local* dynamics of the ODE system in the neighborhood of any steady state point, *i.e.* around the intersections of γ and ψ . Consistent with this result, Figure 9 imposes that the lowest and the highest steady states (b_L^* and b_H^*) are saddle-path stable, while letting the middle one (b_M^*) feature explosive cycles.

Notwithstanding these restrictions on the local dynamics, there remain three distinct possibilities with regard to the *global* dynamics. Figure 9 considers one of these possibilities.

In Figure 9, we have imposed the following property on the global dynamics: both the stable arm that leads to b_L^* from above and the one that leads to b_H^* from below cycle back to b_M^* . It follows that there exist

Figure 9: Beyond the Benchmark: Rich Dynamics and Multiple Steady States



values \tilde{b} and $\tilde{\tilde{b}}$, as indicated in the figure, such that the following is true within region M. Whenever $b_0 < \tilde{b}$, $\Lambda(b_0)$ is a singleton and the unique candidate for optimality is the saddle path that leads to b_L^* . Whenever $b > \tilde{\tilde{b}}$, $\Lambda(b_0)$ is again a singleton, but now the unique candidate is the saddle path that leads to b_H^* . Finally, whenever $b_0 \in [\tilde{b}, \tilde{\tilde{b}}]$, there are multiple paths that are candidates for optimality. For instance, if we take $b_0 = \hat{b}$ as indicated in the figure, one candidate is obtained by setting $\lambda_0 = \hat{\lambda}_1 \equiv \max \Lambda(b_0)$ and letting debt decrease monotonically towards b_L^* ; another candidate is obtained by setting $\lambda_0 = \hat{\lambda}_2 = \min \Lambda(b_0)$ and letting debt increase monotonically towards b_H^* ; and yet another candidate is obtained by setting $\lambda_0 = \hat{\lambda}_3$ and letting debt to cycle twice around \hat{b} before eventually converging to b_H^* . The closer b_0 is to b_M^* , the larger the number of candidates; when b_0 is exactly b_M^* , there is actually a countable infinity of candidates.

At first glance, the task of comparing candidate paths seems daunting. Fortunately, Lemma 4 and the convexity of the Hamiltonian with respect λ guarantee that only the paths associated with the extremes of $\Lambda(b_0)$ can be optimal. For any $b_0 \in [\tilde{b}, \tilde{\tilde{b}}]$, we can thus rule out cycles and restrict attention to just two candidate paths, namely the paths that let b converge monotonically either to b_L^* or to b_H^* . To rank these two candidate paths, we proceed as follows.

First, recall that the value of any candidate path is given by the Hamiltonian as described in Lemma 4; that the Hamiltonian is convex in λ ; and that its derivative is given by $\mathcal{H}_\lambda = \dot{b}$. Next, consider the value of \dot{b} at each of the two candidate paths. For all $b_0 \in [\tilde{b}, \tilde{\tilde{b}}]$, the path that leads to the lowest steady state starts from a point above the graph of ψ , meaning that $\dot{b} < 0$. But as b_0 gets closer to $\tilde{\tilde{b}}$, the starting points gets closer to the graph of ψ , meaning that value of \dot{b} gets closer to 0. In the knife-edge case in which $b_0 = \tilde{\tilde{b}}$, this path is associated with $\dot{b} = 0$. Conversely, the path that leads to the highest steady state is associated with $\dot{b} > 0$ for all $b_0 \in (\tilde{b}, \tilde{\tilde{b}}]$, and with $\dot{b} = 0$ in the reverse knife-edge case in which $b_0 = \tilde{b}$.

Combining these observations, we obtain the following properties. When $b_0 = \bar{b}$, the path that leads to b_L^* features $\mathcal{H}_\lambda = \dot{b} < 0$, whereas the path that leads to b_H^* features $\mathcal{H}_\lambda = \dot{b} = 0$. By the convexity of \mathcal{H} , the latter path is dominated. Conversely, when $b_0 = \tilde{b}$, it is the former path that now features $\mathcal{H}_\lambda = \dot{b} = 0$ and that is therefore dominated. By continuity,²⁶ the path that leads to b_L^* is therefore optimal for b_0 close enough to \bar{b} , whereas the path that leads to b_H^* is optimal for b_0 close enough to \tilde{b} . Finally, the assumption that U is convex in s guarantees that the optimal path for b is monotone. It follows that there exists a threshold $\hat{b} \in (\bar{b}, \tilde{b})$ such that the *unique* optimal path is the path leading to the lowest steady state whenever $b_0 < \hat{b}$ and it is the path leading to the higher steady state whenever $b_0 > \hat{b}$. See Figure 9 for an illustration: the bold segments of the two stable arms indicate the optimal selection among the two candidate paths.²⁷

So far, we focused on region M. In region H ($b_0 \geq b_{bliss}$), the analysis is similar to Figure 8. That is, there is a threshold $b_{skiba} \in (b_{bliss}, \bar{b})$ such that, as long as $b_0 \in (b_{bliss}, b_{skiba})$, there are two candidate paths, the one leading to b_H^* and the Barro-like one, and the former dominates the latter, whereas the latter is the only candidate for $b_0 \geq b_{skiba}$. Finally, in region L ($b_0 < b_{seign}$), there is a unique candidate path, one leading to b_L^* .

The kind of optimal policy illustrated in Figure 9 has the following properties: (i) whenever $b_0 < \hat{b}$, debt converges monotonically to b_L^* ; (ii) whenever $b_0 \in (\hat{b}, b_{skiba})$, debt converges monotonically to b_H^* ; and (iii) whenever $b_0 \geq b_{skiba}$, debt stays constant at b_0 for ever. Comparing this result to our earlier benchmark, we see that one key property survives whereas another is lost: as in our benchmark, it is true that there exists a threshold $b_{skiba} > b_{bliss}$ such that debt converges to a steady-state level below b_{bliss} whenever the economy starts below b_{skiba} ; but unlike our benchmark, the steady-state level is not the same for all initial conditions.

We now turn to two variants of the case studied in Figure 9. One of these variants is illustrated in Figure 10, the other in Figure 11. These variants maintain the same qualitative configuration for the functions γ and ψ , the same steady-state points, and the same local dynamics around them, but perturb the global dynamics. One of the stable arms is now allowed to extend throughout region M instead of cycling back to b_M^* . This path then emerges as the optimal path for *all* initial conditions: in the case seen in Figure 10, it is optimal to converge to b_H^* for all $b_0 < b_{skiba}$; and in the case seen in Figure 11, it is optimal to converge to b_L^* .

Let us fill in the details, starting with Figure 10. Unlike Figure 9, the stable arm corresponding to the highest steady state no longer cycles back to b_M^* ; instead, it extends past b_L^* . This has the following important implication. If we consider $b_0 = b_L^*$, then there are two candidate optimal plans, namely the

²⁶Here, we take for granted the continuity of the value of each candidate path with respect to b_0 ; for a general proof of this property, see Dechert and Nishimura (1981).

²⁷In the optimal-control literature, *any* threshold level of the state variable at which the solution switches from one to another candidate path, such as the threshold \hat{b} here, is often referred to as a “Skiba point”. In our paper, we reserve the notation b_{skiba} to refer only to the highest such threshold.

Figure 10: Optimal to Converge to b_H^* for all $b_0 < b_{skiba}$





Figure 11: Optimal to Converge to b_L^* for all $b_0 < b_{skiba}$



plan of staying put at b_L^* and the plan that leads to b_H^* . The former plan is dominated because it features $\mathcal{H}_\lambda = \dot{b} = 0$, whereas the latter features $\mathcal{H}_\lambda = \dot{b} > 0$. By continuity, the saddle path that leads to b_L^* is dominated also for any b_0 in an open neighborhood of b_L^* . But then the path leading to b_L^* can *never* be optimal: if the economy were to follow this path starting from any initial point b_0 , the economy would enter the aforementioned neighborhood in finite time; at that point, switching paths would increase welfare, which contradicts the optimality of the original path. We conclude that, contrary to what happens in Figure 9, the path that leads to b_H^* in Figure 10 is now the optimal path for all $b_0 < b_{skiba}$.

Consider next Figure 11. This illustrates a diametrically opposite scenario from that shown in Figure 10: it is now the stable arm that leads to b_L^* that fails to cycle back to b_M^* , extends past b_H^* , and dominates throughout. What the two scenarios share in common that distinguishes from the scenario depicted in Figure 9 is the following: even though the ODE system continues to admit multiple saddle-path stable steady states, the optimal policy now features a unique and globally stable steady state in the region to the left of the satiation point, that is, optimal debt converges monotonically to the *same* long run value b^* for all initial values $b_0 \leq b_{bliss}$.

These findings illustrate the following more general points and qualify some of the properties of the benchmark model. To the extent that the ODE system admits multiple steady states below b_{bliss} , any such point represents a point of indifference between the desire to depress the interest rate on public debt and the desire to improve liquidity and efficiency; this is our earlier observation that $\Omega_b = 0$ at any such point. Furthermore, to the extent that such a point is locally saddle-path stable, it is optimal to converge to it over time if the economy starts in a small enough neighborhood of this point and if in addition the planner is precluded from moving outside that neighborhood. In this regard, the *local* optimality of the steady state can be understood by inspecting the trade off between collateral creation and interest rate manipulation, as what we did in our benchmark. However, once the planner is free to move from one steady state to another, such local intuitions are no longer sufficient. Moreover, as we show below, there is no guarantee that the steady state can be rationalized as either a global or a local maximum of Ω , despite the fact that it satisfies $\Omega_b = 0$.


The number of possible scenarios would increase if we allowed γ and ψ to intersect more than three times. Yet an additional layer of complexity emerges if the assumption that $\pi(b)b$ is single-valued is relaxed. The tripartite structure of the phase diagram is then lost. Instead, the phase diagram now looks like the outcome of patching together *multiple* pairs of L and M regions from our earlier examples. However, as explained next, this complication does not change the big picture.

Suppose that $\pi(b)b$ has N local extrema, denoted by $\{b_1, b_2, b_3, \dots, b_N\}$, with $\underline{b} < b_1 < b_2 < \dots < b_N < b_{bliss}$, where N is an arbitrary finite number. First, note that $\sigma(b)$ crosses 1 whenever b crosses any of these points. Next, note that the last point, namely b_N , is necessarily a local maximum, because after that point $\pi(b)b$ falls to zero as b approaches b_{bliss} . It follows that $\sigma(b)$ is higher than 1 when $b \in (b_N, b_{bliss})$, lower than 1 when $b \in (b_{N-1}, b_N)$, higher than 1 when $b \in (b_{N-2}, b_{N-1})$, and so on. By the same token, γ

is positively valued $b \in (b_N, b_{bliss})$, negatively valued than 1 when $b \in (b_{N-1}, b_N)$, positively valued when $b \in (b_{N-2}, b_{N-1})$, and so on.

We illustrate this in Figure 12. As anticipated above, the phase diagram now looks like the product of patching together multiple pairs of L and M regions from our earlier examples. But the earlier lessons survive in the following sense: if the economy starts inside any of the L regions, it is optimal to exit this region in finite time and thereafter converge asymptotically either to an intersection point of γ and ψ within one of the M regions or to satiation.

Figure 12: Multiple Regions



Notwithstanding all the complexity, we can thus establish the following result, which offers a qualified generalization of Proposition 11 in our benchmark.

Proposition 13. *Suppose $\psi_{bliss} > \gamma_{bliss}$. There exists a threshold $b_{skiba} > b_{bliss}$ such that, for every $b_0 < b_{skiba}$, the optimal policy lets debt converge monotonically to a point strictly below b_{bliss} .*

Proof. By a similar argument as in Dechert and Nishimura (1981), the optimal path for b is monotone, for any initial condition. Because b is bounded between \underline{b} and \bar{b} , this also means that b converges. The limit point may depend on the initial level of debt. Nevertheless, it is necessarily contained either in the set B^* or in the interval $[b_{bliss}, \bar{b})$.

Let $b^\ddagger \in (0, b_{bliss})$ be the *last* local maximum of $\pi(b)b$.²⁸ By construction of b^\ddagger , $\gamma(b) > 0$ for all $b \in (b^\ddagger, b_{bliss})$ and $\lim_{b \downarrow b^\ddagger} \gamma(b) = +\infty > \psi(b^\ddagger)$. By the assumption that $\gamma_{bliss} < \psi_{bliss}$ along with the continu-

²⁸Because $\pi(b)b$ is strictly positive for all $b \in (0, b_{bliss})$ and converges to zero as b approaches either 0 from above or b_{bliss} from below, we know that there exists $\epsilon > 0$ such that $\pi(b)b$ is increasing for $b \in (0, \epsilon)$ and decreasing for $b \in (b_{bliss} - \epsilon, b_{bliss})$. Because the derivative of $\pi(b)b$ is $-(\sigma(b) - 1)\pi(b)$, the aforementioned property means that $\sigma(b) < 1$ for $b \in (0, \epsilon)$ and $\sigma(b) > 1$ for $b \in (b_{bliss} - \epsilon, b_{bliss})$. By the continuity of σ , then, the threshold b^\ddagger exists and is strictly between 0 and b_{bliss} .

ity and differentiability of γ and ψ , there exists at least one point $b^* \in (b^\dagger, b_{bliss})$ such that $\gamma(b^*) = \psi(b^*)$ and $\gamma'(b^*) < \psi'(b^*)$, that is, a steady-state point in which γ intersects ψ from above. If there are multiple such points, consider the highest one. By Proposition 12, we know that this steady state is saddle-path stable. Similarly to Figure 8, the following is therefore true: there exists a threshold $b_{skiba} > b_{bliss}$ and a scalar $\epsilon > 0$ such that, whenever $b_0 \in (b^* - \epsilon, b_{skiba})$, there exists path that satisfies the ODE system at all t and that asymptotically leads to b^* . Clearly, this path is a candidate for optimality for all $b_0 \in (b^* - \epsilon, b_{skiba})$. Furthermore, this path dominates the Barro-like plan for all $b_0 \in [b_{bliss}, b_{skiba})$. Finally, there is no candidate path that leads to satiation when $b_0 < b_{bliss}$, thanks again to the assumption that $\gamma_{bliss} < \psi_{bliss}$. \square

All these facts obtain by applying the same arguments as in our benchmark. What is different is that we no longer know (i) whether the path that leads to b^* ceases to exist for b_0 low enough and (ii) whether this path is itself dominated by another candidate path in a region of b_0 . Notwithstanding these possibilities, any other candidate path must itself be a saddle path leading to one of the intersection points of γ and ψ . By construction of b^* , any other such point is strictly below b^* . It follows that, no matter the initial level of debt and no matter which candidate path is the optimal one, debt converges to a point that does not exceed b^* , which proves the claim.²⁹

B.4 The condition $\psi_{bliss} > \gamma_{bliss}$

In the preceding analysis, the condition $\psi_{bliss} > \gamma_{bliss}$ played a crucial role: it guaranteed that it is optimal to lead the economy to a steady state below satiation not only for all initial levels of debt below b_{bliss} , but also over a range of initial levels above it. This generalized the related insight from the main text.

As already explained, the condition $\psi_{bliss} > \gamma_{bliss}$ has a simple interpretation: it means that, in the neighborhood of b_{bliss} , the shadow cost of taxation is sufficiently high so that the marginal value of depressing the interest rate on public debt outweighs the marginal cost of the financial distortion. Consistent with this interpretation, it is straightforward to show this case obtains when the level of government spending is sufficiently high.³⁰

Lemma 5. *Suppose $\gamma_{bliss} < \infty$. There exists a threshold \hat{g} such that $\psi_{bliss} > \gamma_{bliss}$ if and only if $g > \hat{g}$.*

Proof. Note that ψ_{bliss} is continuous and increasing in g as long as $g < g_{max}$ and diverges to $+\infty$ as $g \rightarrow g_{max}$. This is because a higher g requires higher taxes to balance the budget, and the marginal cost


²⁹This argument mirrors Theorem 2 in Brock and Dechert (1983). Applied to our setting, this theorem states that, whenever the policy rule of the costate features a discontinuous jump, this jump is downward. By the same token, as we move from higher to lower levels of debt, the costate can only jump upwards, which means that lower levels of debt are necessarily associated with convergence to weakly lower steady states.

³⁰In fact, the threshold \hat{g} in the lemma can be *negative* in some economies, implying that, in these economies, this result obtains for *all* positive levels of government spending.

of these taxes explodes to infinity as we approach the peak of the Laffer curve. Furthermore, $\psi_{bliss} = 0$ if and only if $g = -\rho b_{bliss} < 0$. Finally, note that γ_{bliss} is (i) invariant to g ; (ii) positive for the reasons offered above; and (iii) finite by assumption. It then follows that there exists a threshold \hat{g} , necessarily less than g_{max} and possibly negative, such that $\psi_{bliss} > \gamma_{bliss}$ if and only if $g > \hat{g}$. \square

This generalizes the related point made in the main text. The only subtlety is the following. In the benchmark studied in the main text, $\psi_{bliss} > \gamma_{bliss}$ (and by the same token $g > \hat{g}$) was both sufficient and necessary for $b_{skiba} > b_{bliss}$ and, equivalently, for the existence of a steady state below satiation. Sufficiency was established in Proposition 11, necessity in Proposition 10. In the more general case allowed here, sufficiency remains valid by Proposition 13, but necessity may not apply.

Figure 13: No Satiation Despite $\psi_{bliss} < \gamma_{bliss}$ (or g low enough)



We illustrate this in Figure 13. As in our benchmark (see Figure 7 in particular), letting $\gamma_{bliss} > \psi_{bliss}$ guarantees the local existence of a candidate path that leads to satiation: for some $\epsilon > 0$ and all $b_0 \in (b_0 - \epsilon, b_{bliss})$, there exists a path that satisfies the ODEs at all dates and that asymptotically converges to b_{bliss} . But unlike what was true in our benchmark, this type of path does not exist for sufficiently low b_0 . What is more, for all $b_0 < b_{bliss}$, there happens to exist another candidate optimal path, namely the one that leads to a steady state below b_{bliss} . Finally, note that the path leading to b_{bliss} features an initial value for \dot{b} that is arbitrarily close to 0 when b_0 is close enough to b_{bliss} , whereas the path leading to b_L^* features a \dot{b} bounded away from zero. Using once again Lemma 4, the convexity of \mathcal{H} in λ , and the fact that $\mathcal{H}_\lambda = \dot{b}$, we infer that the latter path dominates the former for b_0 in a neighborhood of b_{bliss} . But this also means that the path leading to satiation can not be optimal for any initial b_0 . Instead, there

again exists a $b_{skiba} > b_{bliss}$ such that for all $b_0 < b_{skiba}$ it is optimal to converge either to b_L^* or to some point further below.

To sum up, away from the benchmarks studied in the main text, g high enough may not be necessary for the existence of a steady state below satiation. But it is always sufficient for this to be true, and this is what we think of as the most interesting scenario.

B.5 Complete Characterization

Building on the preceding results, we can now offer a characterization of the optimal policy that nests all possible scenarios. To this goal, we henceforth let

$$B^\# \equiv \{b \in (\underline{b}, b_{bliss}] : \gamma(b) = \psi(b) \text{ and } \gamma'(b) \leq \psi'(b)\}$$

be the set of the points at which γ intersects ψ from above. As shown in Proposition 12, these points identify the saddle-path stable steady states of the ODE system.³¹ Depending on primitives, $B^\#$ may be empty, or may contain an arbitrary number of elements.³² Regardless of this, we have the following result.

Theorem 2. *In every economy, there exists a threshold $b_{skiba} \in [\underline{b}, \bar{b}]$ and a set $B^* \subseteq B^\#$ such that the following are true along the optimal policy:*

- (i) *If either $b_0 \in B^*$ or $b_0 > \max\{b_{bliss}, b_{skiba}\}$, debt stays constant at b_0 for ever.*
- (ii) *If $b_0 < b_{skiba}$ and $b_0 \notin B^*$, then debt converges monotonically to a point inside B^* .*
- (iii) *If $b_{skiba} < b_{bliss}$ and $b_0 \in (b_{skiba}, b_{bliss})$, debt converges monotonically to b_{bliss} .*

Proof. We prove this result with the help of Theorem 2 from Brock and Dechert (1983). Consider the optimal policy rule for the co-state variable, namely the correspondence from any given b_0 to the optimal value for λ_0 . Denote this correspondence by Λ^{opt} . Note that this is a selection from the correspondence Λ (which was defined in the context of Lemma 4). To illustrate, consider Figure 9. In this example, the aforementioned correspondence is given by the combination of three segments: the thick green line on the left of \hat{b} , plus the solid blue line between \hat{b} and b_{skiba} , plus the segment of the graph of the $\dot{b} = 0$ locus that rests on the right of b_{skiba} . As it is evident in this example, the correspondence λ^* is single-valued and continuous for all b_0 other than \hat{b} ; the discontinuity at \hat{b} reflects a switch in the optimal selection among different candidate paths. Moving beyond this specific example, the policy rule for the co-state can feature multiple such discontinuities. Any such discontinuity, however, has

³¹In knife-edge cases in which a steady state of the ODE system features $\gamma'(b) = \psi'(b)$, we can not be sure of saddle-path stability. Clearly, such knife-edge cases are degenerate. In any event, they do not affect the validity of the result stated below, because this result allows B^* to be a *strict* subset of $B^\#$.

³²We wish to think of the empirically relevant case as one in which $B^\#$ contains either a single or a “small” finite number of points. At the present level of abstraction, however, the best we can say is that $B^\#$ is generically countable.

to involve a jump in a specific direction: applied to our setting, Theorem 2 from Brock and Dechert (1983) states that, at any point \hat{b} such that $\lim_{b \uparrow \hat{b}} \Lambda^{opt}(b) \neq \lim_{b \downarrow \hat{b}} \Lambda^{opt}(b)$, it is necessarily the case that $\lim_{b \uparrow \hat{b}} \Lambda^{opt}(b) > \lim_{b \downarrow \hat{b}} \Lambda^{opt}(b)$,³³ In other words, as we move from higher to lower levels of debt, the co-state can only jump upwards, which means that the rate of taxation and the level of government surpluses must also jump upwards. It then follows that lower initial conditions are necessarily associated with convergence to lower steady states, which in turn is the key to the result.

Thus suppose there exists an initial point $b_0 = \tilde{b}_0$ such that it is optimal to converge to a point $b^* < b_{bliss}$. Clearly, b^* must be inside $B^\#$. Next, consider the set of points at which the policy rule of the co-state features a discontinuity and let \hat{b} be the highest such point below b^* ; if no such point exists, just let $\hat{b} = \underline{b}$. When $b_0 \in (\hat{b}, \tilde{b}_0)$, debt converges to b^* . When instead $b_0 < \hat{b}$ (which, of course, is relevant only insofar as $\hat{b} > \underline{b}$), debt converges to a point that is below \hat{b} , and hence also below b^* , but still inside $B^\#$. It follows that there exists a point $b_{skiba} \geq b^*$ such that, when $b_0 \leq b_{skiba}$, then and only then it is optimal to converge to a point inside $B^\#$.

The above argument presumed the existence of an initial point at which it became optimal to converge to a point below b_{bliss} . If no such initial point exists, we simply let $b_{skiba} = \underline{b}$. This completes the proof of part (ii) of our theorem.

To prove part (iii), recall from Proposition 13 that $\psi_{bliss} > \gamma_{bliss}$ is sufficient for $b_{skiba} > b_{bliss}$. It follows that $b_{skiba} < b_{bliss}$ is possible only insofar as $\psi_{bliss} < \gamma_{bliss}$, which in turn guarantees the existence of a candidate path that converges to b_{bliss} for any $b_0 \in [\hat{b}, b_{bliss})$ and some $\hat{b} < b_{bliss}$. Clearly, $\hat{b} \leq b_{skiba}$. By definition of b_{skiba} , the optimal path is one of the candidate paths that converge to a point inside $B^\#$ if and only if $b_0 < b_{skiba}$. Therefore, for any $b_0 \in [b_{skiba}, b_{bliss})$, either the path that leads to b_{bliss} is the unique candidate path, or it dominates any of the candidate paths that lead to a point inside $B^\#$.

Turning to part (i), note that this contains two subparts: one regarding $b_0 \in B^*$, and another regarding $b_0 \geq \max\{b_{skiba}, b_{bliss}\}$. Once part (ii) of the theorem is established, the first of the aforementioned two subparts is trivial: it merely identifies B^* as the set of the steady states of the optimal policy that happen to lie below b_{bliss} . The second subpart, on the other hand, is proved by the following variant of the proof of part (ii). As long as $b_0 \geq b_{bliss}$, there necessarily exists a Barro-like candidate path that keeps the level of debt constant at its initial value and the premium at zero for ever. Whenever another candidate path exists, it converges to a point inside $B^\#$. By definition of b_{skiba} , such a path is optimal if and only if $b_0 < b_{skiba}$. It follows that, whenever $b_0 \geq \max\{b_{bliss}, b_{skiba}\}$, either the aforementioned Barro-like path is the unique candidate path or it dominates any alternative path. \square

The point b_{skiba} is a threshold in the state space such that it is optimal to satiate the private sector's demand for collateral—and eliminate the financial distortion—in the long run if and only if the initial

³³At first glance, the original version of Theorem 2 in Brock and Dechert (1983) appears to state the opposite; the apparent contradiction is resolved by noting that our co-state variable is defined with the opposite sign than theirs.

level of public debt exceeds this threshold. The set B^* , on the other hand, identifies the set of the steady-state points of the optimal policy—aka the optimal steady states—that lie below the satiation point.

When B^* is a singleton, debt converges to the unique point in B^* for all $b_0 < b_{skiba}$. When instead B^* contains multiple points, each such point is associated with a basin of attraction around it, and the union of all these basins equals $[\underline{b}, b_{skiba})$.

Clearly, B^* has to be a subset of $B^\#$, but the two need not coincide: it is possible that the planner never finds optimal to converge to some, or even any of the points in $B^\#$. For instance, whereas $B^* = B^\#$ in Figures 8 and 9, B^* is a strict subset of $B^\#$ in Figures 10 and 11.

Finally, it is generally possible that $B^* = \emptyset$, meaning that satiation obtains in the long run regardless of initial conditions. But as already explained, this scenario is possible only if g is low enough, or $\psi_{bliss} < \gamma_{bliss}$. Conversely, g high enough suffices for the economy to admit at least one steady state below satiation—and this is the case we find most interesting.

We conclude with the following clarification: Proposition 2 identifies a set of possible scenarios for the optimal policy, but does not specify whether each of these scenarios does obtain for some economies. The next result completes the picture by offering a complete taxonomy of all the economies under consideration and of all possibilities that *do* obtain for some specification of U , V and π .

Theorem 3. *Any economy belongs to one of the following three non-empty classes:*

- (i) *Economies in which $B^* = \emptyset$ and $b_{skiba} = \underline{b}$.*
- (ii) *Economies in which $B^* \neq \emptyset$ and $b_{skiba} \in (\underline{b}, b_{bliss})$.*
- (iii) *Economies in which $B^* \neq \emptyset$ and $b_{skiba} > b_{bliss}$.*

Furthermore, g high enough is sufficient for an economy to belong to the last class.

Proof. That any economy must belong to one of these three classes follows from Theorem 2. That the first and the third classes are not empty follows from the examples we have already considered; an example of the second class was provided above. Finally, the claimed sufficiency of the condition $\psi_{bliss} > \gamma_{bliss}$ follows from Proposition 13. \square

B.6 Local Dynamics and Local Comparative Statics

We conclude this Appendix with two additional results. The first result establishes that, in a neighborhood of any steady state below satiation, debt and taxes co-move along the transition to it. The second result offers a general result on the comparative statics of the model.

Proposition 14. *For any $b^* \in B^*$ there exists $\epsilon > 0$ such that the following is true: if $b_0 \in (b^* - \epsilon, b^*)$, then both debt and taxes increase over time; and if $b_0 \in (b^*, b^* + \epsilon)$, then both debt and taxes decrease over time.*

Proof. By the definition of $b^* \in B^*$ and $b^* < b_{bliss}$, we know that the point $(b^*, \lambda^* \equiv \psi(b^*))$ is locally stable. Similarly to Proposition 12, the local dynamics are given by

$$\dot{X}(t) = \begin{pmatrix} \rho + \pi(b^*)(\sigma(b^*) - 1) & -S'(\lambda^*) \\ V''(b^*) - \lambda^* \pi'(b^*)(\sigma(b^*) - 1) - \lambda^* \pi(b^*)\sigma'(b^*) & -\pi(b^*)(\sigma(b^*) - 1) \end{pmatrix} X(t) = \mathbf{J}X(t)$$

we know from proposition 12, that the eigenvalue associated with the stable arm is given by $\mu = \frac{\rho - \sqrt{\Delta}}{2}$ with $\Delta > 0$ (see proof of Proposition 12). It is then straightforward to obtain the eigenvector $\mathbf{v} = (v_1, v_2)$ satisfying

$$\left(\frac{\rho}{2} + \frac{V'(b^*)}{\gamma(b^*)} + \frac{\sqrt{\Delta}}{2} \right) v_1 - S'(\lambda^*) v_2 = 0$$

An eigenvector is $\left(S'(\lambda^*), \frac{\rho}{2} + \frac{V'(b^*)}{\gamma(b^*)} + \frac{\sqrt{\Delta}}{2} \right)$. Given that $V'(b^*) > 0$, $\gamma(b^*) > 0$ (since $\sigma(b^*) > 1$) and $S'(\lambda^*) > 0$ in the upward sloping part of the Laffer curve, both components of the vector are positive. The comovement result follows: For any $\varepsilon > 0$, starting from $b_0 = b^* - \varepsilon$ (resp. $b_0 = b^* + \varepsilon$), the economy will converge to (b^*, λ^*) increasing (resp. decreasing) both debt and taxes along the transition path. \square

Proposition 15. *Let $v(\cdot) = \omega\pi(\cdot)$ and hold $\sigma(\cdot)$ constant. For any $b^* \in B^*$, b^* increases with a small enough increase in ω , a small enough decrease in g , or a small enough increase in $\pi(\cdot)$.*

Proof. Any $b^* \in B^*$ is such that $\gamma(b^*) = \phi(b^*)$. Therefore, it inherits the comparative statics of the γ and ϕ functions described in Section B.1. \square

These two results together imply that, at least for small changes in the primitives of the economy, the relevant trade off, the nature of transitional dynamics, and the comparative statics of the optimal long-run quantity of debt are the same as those discussed in the main text.

C Additional Results

C.1 Private versus social value of liquidity

Consider the micro-founded model of Section 2. Let $a(\theta, b)$ and $P(b)$ denote the allocation of the bond and the price of the afternoon good that obtains from solving the planning sub-problem (36)-(44) reported in Section D and, to simplify, let $a(\theta, b) > -\phi$ for all θ . From the definition of $V(\cdot)$ together with the fact that the aggregate net trade of the morning good is zero in equilibrium, we have that

$$V(b) = \int \mathcal{U}(a(\theta, b), \theta, P(b)) \varphi(\theta) d\theta$$

and therefore

$$V'(b) = \int [\mathcal{U}_a(\cdot) a_b(\theta, b) + \mathcal{U}_p(\cdot) P_b(b)] \varphi(\theta) d\theta.$$

From the household's Euler condition (43), we have that $Q(b) = \beta + \mathcal{U}_a(\cdot)$ for all θ . From the bond market clearing condition (Equation 38), we have $\int a_b(\theta, b)\varphi(\theta)d\theta = 1$. Using these last two relations, we infer that

$$V'(b) = \pi(b) + e(b),$$

where $\pi(b) \equiv Q(b) - \beta$ is the market premium and $e(b) \equiv \int \mathcal{U}_p(\cdot)\varphi(\theta)d\theta P_b(b)$ is the relevant externality. Finally, it can be shown that $\int \mathcal{U}_p(\cdot)\varphi(\theta)d\theta$ and $P_b(b)$ are strictly negative and strictly positive when the collateral constraint binds with positive probability, and zero otherwise. The intuition is simple: as long as the constraint binds, a higher b means a higher P because it facilitates a more efficient allocation of the morning good. A higher price has a negative aggregate welfare effect because it tightens the constraint and distorts the allocation. As long as the constraint binds, we therefore have $e(b) < 0$, or equivalently $\pi(b) > V'(b)$.

C.2 Allowing for state-contingent debt

We now discuss how our analysis qualifies the insights of Lucas and Stokey (1983). Relative to Barro and AMSS, the key difference in Lucas and Stokey (1983) is the availability of state-contingent debt. This makes it feasible for the government to completely insulate its budget against any shock. But is it desirable to do so?


The answer to this question is unambiguously “yes” in Lucas and Stokey (1983). This is because the transfers implemented by state-contingent debt are non-distortionary, so that the planner necessarily prefers them to any variation in the distortionary tax. This also explains why Lucas and Stokey (1983) find that the tax distortion is smoothed, not only across dates, but also across states; or, by the same token, why the optimal allocation is history-independent, in sharp contrast to the unit-root persistence predicted by Barro and AMSS.

The answer differs in our setting. When state-contingent debt is available, our planner maintains the option to equate the shadow cost of taxation across different histories of shocks, exactly as in Lucas and Stokey (1983). But unlike that environment, the planner no longer finds it optimal to do so. Instead, he finds it optimal to deviate from tax smoothing across states, in a manner that resembles the departure from smoothing taxes across dates in the deterministic model.

The rationale is simple. In order to eliminate variation in the shadow cost of taxation, the planner would have to endure a non-trivial variation in the aggregate collateral, or liquidity, of the private sector. Starting from this reference point, a small mean-preserving reduction in the variation of the value of government liabilities leads to a second-order welfare loss in terms of increased variation in the cost of taxation but to a first-order welfare gain in terms of reduced variation in the social value of liquidity and/or seigniorage collected. It follows that the optimal policy accommodates some variation in the tax distortion in order to smooth the supply of liquidity to the private sector. But this also means that the economy behaves *as if* the planner did not have access to a complete set of state contingent debt

instruments: the optimal tax and the optimal allocation depend on the history of fiscal shocks *as if* those were (partially) uninsurable.

We illustrate this property in Figure 14 using a persistent war. This is exactly the same as in the bottom of Figure 3, except that now debt is allowed to be state-contingent. The black lines give the impulse responses of the market value of debt and the tax rate in our model; the orange lines give their Lucas-Stokey counterparts, i.e., those that obtain in the absence of the financial friction. In both cases, the market value of debt jumps down in response to the war, reflecting the state-contingency of the debt burden. But the drop is more modest in the presence of the financial friction (black line), reflecting the planner's desire to limit the reduction in aggregate collateral. By the same token, the planner in our setting opts to raise more taxes during the war, while in the Lucas-Stokey benchmark the tax rate does not change at all.



To sum up, once public debt is non-neutral for the reasons we have accommodated in this paper, the difference between Barro/AMSS and Lucas-Stokey is attenuated: the qualitative response of the optimal tax and the optimal allocation is the same whether the government has access to state-contingent debt or not.

C.3 On the Friedman Rule

Our analysis departs from that in the Friedman-rule literature by allowing all types of government-issued assets, rather than a subset of them, to facilitate private liquidity. This assumption seems both appropriate for the issues we are addressing and realistic (see Krishnamurthy and Vissing-Jorgensen (2012) for corroborating evidence). To elaborate on the role played by this assumption, we now consider a modification of our baseline model that helps nest the case studied in the Friedman-rule literature.

Suppose that the government enacts a law that prohibits the use of corporate bonds as collateral in morning transactions. This restriction adds a constraint to the planner's sub-problem defined in (36)-

(44) and results in a change in the functions V and π . By shutting the private supply of collateral down the law may reduce V and increase π , so it has ambiguous welfare implications. But its effect on the price of corporate bonds is unambiguous: they are now priced at the discount factor, while government bonds command a premium over the discount factor. Our model is now directly comparable to those in the Friedman rule literature, with government bonds playing the role of money and corporate bonds the role of the non-money asset.

Suppose next that the government can not only borrow in the money-like asset (here, government bonds), but also invest freely in the non-money asset (here, corporate bonds). Then, public debt is given by $b = m + n$, where m is the stock of government bonds and $-n$ is the quantity of corporate bonds held by the government. The budget constraint is given by

$$\dot{m} + \dot{n} = [\rho - \pi(m)]m + \rho n + g - s,$$

or equivalently

$$\dot{b} = \rho b - \pi(m)m + g - S(\tau), \quad (31)$$

where $\pi(m)m$ is seigniorage and $S(\tau)$ is tax revenue. The following is evident: For any given $b, \pi(m)m, g$, the government can vary the mixture of taxes and new debt issued that satisfies its budget without affecting either the level of private sector liquidity or the interest rate on public debt. Moreover, the latter is now equal to the discount rate.

Therefore, when the government varies b , it does not face the key trade off present in our model. By the same token, the optimal supply of liquidity is disentangled from the optimal dynamics of debt and taxes, and the latter are determined in exactly the same fashion as in Barro (1979).

To see this more clearly, integrate (31) over time to obtain the familiar intertemporal budget constraint:

$$b_0 + G = \int_0^{+\infty} e^{-\rho t} [\pi(m)m + S(\tau)] dt. \quad (32)$$

where $G \equiv \int_0^{+\infty} e^{-\rho t} g dt$ is the present value of government spending. The planner's problem reduces to finding the paths of m and τ that maximize ex ante welfare,

$$\int_0^{+\infty} e^{-\rho t} [U(\tau) + V(m)] dt,$$

subject to the single integral constraint in (32). Let λ^* denote the Lagrange multiplier on the intertemporal budget. It is then immediate that the optimal supply of liquidity is given by

$$m^* = \operatorname{argmax}_m \Omega(m, \lambda^*), \quad (33)$$

where $\Omega(m, \lambda) \equiv V(m) + \lambda \pi(m)m$ measures "liquidity plus seigniorage". Depending on primitives, m^* may or may not coincide with satiation; that is, the Friedman rule may or may not apply. Regardless of this, however, tax smoothing obtains and the optimal fiscal policy is determined in exactly the same fashion as in Barro (1979).

Another, subtler point is that the m^* characterized above *necessarily* attains the global maximum of $\Omega(m, \lambda)$. By contrast, this is not necessarily true in our context. In particular, under the Auxiliary Assumptions introduced in Section 3.3, it is true that b^* is unique and maximizes $\Omega(b, \lambda^*)$. But without these assumption, it is possible that there are multiple steady states, or even a unique steady state that attains a local *minimum* of Ω . And while we don't view this possibility as practically relevant, it does reinforce our point about how the optimal provision of liquidity is intertwined with the transitional dynamics, or the desire to smooth taxes.

C.4 Relation to Aiyagari and McGrattan (1998)

In this Appendix we discuss why the solution strategy followed in Aiyagari and McGrattan (1998) both fails to recognize this trade off and offers a distorted answer to the question of interest.

That paper allows for more realistic micro-foundations than ours, including concave utility and an empirically calibrated labor-income risk. The role played by public debt is fundamentally similar (it eases the underlying borrowing constraint), but the wealth heterogeneity is a relevant state variable for aggregate outcomes, forcing the authors not only to rely on numerical simulations but also to take a certain short-cut. Instead of solving the problem of a Ramsey planner who chooses the dynamic path of taxes and debt so as to maximize ex-ante utility, they restrict taxes and debt to be constant over time, abstract from transitional dynamics, and maximize welfare in steady state.

Replicating this strategy in our framework means maximizing $U(s) + V(b)$ subject to $R(b)b = g + s$. Let \hat{b} denote the debt level that solves this problem and let $\hat{\lambda}$ be the associated Lagrange multiplier. Clearly,

$$\hat{b} = \arg \max_b \{V(b) - \hat{\lambda}R(b)b\}, \quad (34)$$

This underscores how the Aiyagari-McGrattan approach treats the interest payments on public debt, $R(b)b$, as a cost. But as first highlight in Section 4.1, the component ρb of these interest-rate payments is *not* a cost. Accordingly, the *truly* optimal steady state satisfies

$$b^* = \arg \max_b \{V(b) + \lambda^* \pi(b)b\}, \quad (35)$$

which underscores that the correct planning problem treats debt issuance as a profit-generating business to the tune of $\pi(b)b$.

In summary, the solution strategy Aiyagari and McGrattan (1998) not only abstracts from transitional dynamics (or, relatedly, the optimal response to shocks) but also offers a distorted perspective on optimal long-run quantity of public debt. At the same time, Aiyagari and McGrattan (1998) allow for an interesting economic effect that our main analysis abstract from: that public debt may crowd out capital accumulation by offering a substitute form of buffer stock. We explain why this possibility, or even the opposite one, does not fundamentally change the policy problem in Section 4.1 in the main text.

Finally, note that discussion here presumes, like the analysis in Section 3.3, that the Auxiliary Assumptions hold. As explained in Appendix B.3, the economy may feature multiple steady states when these assumptions do not hold. In these circumstances, the Aiyagari-McGrattan approach will never detect this multiplicity, for it is the (generically) unique solution to a static optimization problem.

D Main Text Proofs

Proof of Proposition 1. The proof of part (i) of the proposition proceeds in two steps. The first step show how to represent the individual's problem as one in which the level of assets enters the utility function. The second solves for the equilibrium in the morning and afternoon markets and shows how to express the resulting welfare and the equilibrium price of public debt as functions of the quantity of public debt.

Step 1. Let us start from the primitive formulation of the individual's problem:

$$\begin{aligned} & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (c_{it} + \theta u(x_{it}) - v(h_{it})) \right] \\ \text{s.t. } & c_{it} + p_t x_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} + p_t \bar{e} \\ & p_t (x_{it} - \bar{e}) \leq \xi w_t h_{it}^{def} + a_{it} \\ & -a_{it+1} \leq \xi w_{t+1} h_{it+1}^{def} \end{aligned}$$

Assuming a zero tax rate when there is default, the labor supply in the event of default solves

$$v'(h_{it}^{def}) = (1 - \xi) w_t$$

because the marginal utility of the afternoon consumption good, and hence the Lagrange multiplier associated to the budget constraint, is 1. Using the fact that the equilibrium wage rate is A , the two financial constraints can be written as

$$\begin{aligned} z_{it} & \leq \phi + a_{it} \\ a_{it} & \geq -\phi \end{aligned}$$

where $\phi \equiv \xi A (v')^{-1}((1 - \xi)A)$ and $z_{it} = p_t (x_{it} - \bar{e})$. Defining $\tilde{c}_{it} = c_{it} + z_{it}$, we have

$$\begin{aligned} & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (c_{it} + \theta u(x_{it}) - p_t (x_{it} - \bar{e}) - v(h_{it})) \right] \\ \text{s.t. } & \tilde{c}_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} \\ & p_t (x_{it} - \bar{e}) \leq \phi + a_{it} \\ & a_{it+1} \geq -\phi \end{aligned}$$

Consider now the sub problem of determining the demand for the morning good. This problem is purely static and is given by

$$\begin{aligned} & \max_x [\theta u(x) - p(x - \bar{e})] \\ \text{s.t. } & p(x - \bar{e}) \leq \phi + a \end{aligned}$$

which gives $x = \mathcal{X}(a, \theta, p)$ and an indirect utility net of the cost of purchasing $\mathbf{u}(a, \theta, p) = \theta u(\mathcal{X}(a, \theta, p)) - p\mathcal{X}(a, \theta, p)$. Defining the discounted expected indirect utility of the morning good as

$$\mathcal{U}(a, \theta, p) \equiv \beta \int \mathbf{u}(a, \theta', p) \varphi(\theta' | \theta) d\theta'$$

and using it in the optimization program allows us to write it as

$$\begin{aligned} & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (\tilde{c}_{it} - v(h_{it}) + \mathcal{U}(a_{it+1}, \theta_t, p_{t+1})) \right] \\ \text{s.t. } & \tilde{c}_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} \\ & a_{it+1} \geq -\phi \end{aligned}$$

The utility $\mathcal{U}(a_{it+1}, \theta_t, p_{t+1})$ will be used in the next part of the proof. This concludes the first part of the proof.

Step 2. Because the utility is linear in the afternoon good, the optimal savings decision of every agent in any period t is independent of her initial asset position. This implies that, for any t , the set of implementable next-period wealth distribution, the next-period allocation of the morning good, and the prices q_t and p_{t+1} is independent of the period- t wealth distribution. It follows that we can split the planner's in two parts: an "inner" problem for each period t , where the planner solves for the best implementable allocation of the risk-free asset and the morning good, taking as given the aggregate quantity of public debt; and an "outer" problem over all t , where the planner solves for the optimal path of public debt and taxes.

Fix a $t \geq 0$ and consider the period- t subproblem. This can be represented as follows:

$$\max_{(p, q) \in \mathbb{R}_+^2, (x, a): [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+ \times [-\phi, +\infty)} \int \theta u(x(\theta)) \varphi(\theta) d\theta \quad (36)$$

$$\text{subject to} \quad \int x(\theta) \varphi(\theta) d\theta = \bar{e} \quad (37)$$

$$\int a(\theta_-) \varphi(\theta_-) d\theta_- = b \quad (38)$$

$$\phi + a(\theta_-) - p(x(\theta) - \bar{e}) \geq 0 \quad \forall (\theta, \theta_-) \quad (39)$$

$$\theta u'(x(\theta)) \geq p \quad \forall \theta \quad (40)$$

$$[\theta u'(x(\theta)) - p] [\phi + a(\theta_-) - p(x(\theta) - \bar{e})] = 0 \quad \forall (\theta, \theta_-) \quad (41)$$

$$a(\theta_-) + \phi \geq 0 \quad \forall \theta_- \quad (42)$$

$$\beta + \mathcal{U}_a(a(\theta_-), \theta_-, p) \leq q \quad \forall \theta_- \quad (43)$$

$$[\mathcal{U}_a(a(\theta_-), \theta_-, p) - \pi] [a(\theta_-) + \phi] = 0 \quad \forall \theta_- \quad (44)$$

In this problem, $x(\theta)$ stands for $x_{i,t+1} = x(\theta_{i,t+1})$, $a(\theta_-)$ stands for $a_{i,t+1} = a(\theta_{i,t})$, p stands for p_{t+1} , q stands for q_t , and b stands for b_{t+1} . Letting the planner choose the functions (x, a) means that we let her choose the cross-sectional allocation of the risk-free asset and the morning good during period $t + 1$.

This choice is not free as the planner must respect the feasibility and implementability constraints stated in conditions (37) through (44): conditions (37) and (38) are the resource constraint for the morning good and the clearing condition for the asset market; conditions (39)-(44) are the household's optimality conditions for morning consumption and asset holdings, together with the associated collateral constraints and complementary slackness conditions. Finally, note that (36) is simply the ex-ante utility of the morning good. It follows that the solution to (36)-(44) identifies the best cross-sectional allocation of asset holdings and morning-good consumption among those that can be implemented as an equilibrium whenever $b_{t+1} = b$, along with the corresponding prices.³⁴

Now take any path for $\{\tau_t, b_t\}$ that is part of an equilibrium. If there exists a unique equilibrium with this path for taxes and debt, the above problem simply returns the associated allocation of the morning good and the risk-free asset. And if there exist multiple such equilibria, the above problem selects the best one (i.e., the one that maximizes ex-ante utility).

For any b , let $P(b)$ be the resulting value for p ; $Q(b)$ be the resulting value for q ; and let $\pi(b) \equiv Q(b) - \beta$. Next, note that welfare (ex-ante utility) is given, from step 1, by

$$\mathcal{W} \equiv \mathbb{E}_0 \left[\sum \beta^t (c_{it} + \theta_{it} u(x_{it}) - v(h_{it})) \right].$$

By the preceding argument we have that $\mathbb{E}_0[\theta_{it} u(x_{it})]$ equals $\frac{1}{\beta} V(b_t)$ along the best implementable allocation. Strictly speaking, the last statement is valid for $t \geq 1$ but not for $t = 0$. This is because the wealth distribution in period 0 is exogenous and does not have to coincide with the one obtained by the solution to (36)-(44) when $b = b_0$. That is, if we let V_0 denote the value of $\mathbb{E}[\theta_{i0} u(x_{i0})]$ attained at the period-0 equilibrium allocation of the morning good, whatever this is, we have that, in general, $V_0 \neq V(b_0)$. To simplify the notation, we impose $V_0 = V(b_0)$. This is a completely innocuous constant for our results, because b_0 is fixed and the restriction $V_0 = V(b_0)$ does not affect the optimal choice of $\{\tau_t, b_{t+1}\}_{t=0}^{\infty}$. In other words,

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \theta_{it} u(x_{it}) \right] = V_0 + \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^t \frac{V(b_t)}{\beta} \right] = V_0 + \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t V(b_{t+1}) \right]$$

In addition, we know that $\mathbb{E}[c_{it}]$ equals aggregate consumption, c_t , and all agents supply the same amount of labor, $h_{it} = h_t$, due to the quasi-linearity in preferences. We infer that, once we have solved the aforementioned subproblem, we can express welfare as

$$\mathcal{W} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [c_t - v(h_t) + V(b_{t+1})] \quad (45)$$

which completes the part (i) of the proof.

³⁴One potentially confusing point in the definition of the above problem is the following: the problem allows the planner to choose the allocation of the morning good; but it also uses the function $\mathcal{U}_a(a(\theta_-), \theta_-, p)$, which itself embeds the individual's optimal consumption of the morning good. Are the two elements consistent? Yes, because the individual's optimality and feasibility conditions are themselves included in the constraints of the problem.

The properties reported in part (ii) can be proved by considering the first best allocation of the morning good.

$$\theta u'(x^{\text{FB}}(\theta)) = p^{\text{FP}} \quad (46)$$

$$\int x^{\text{FB}}(\theta)\phi(\theta)d\theta = \bar{e} \quad (47)$$

Note that at the first best allocation, debt is priced at the discount fact such that $q = \beta$ and accordingly $\pi(b) = 0$. Let us consider the agent with the highest type, $\bar{\theta}$. At this allocation the borrowing constraint has to be satisfied for this individual

$$\phi + a(\theta_-) + p^{\text{FP}}(x^{\text{FB}}(\bar{\theta}) - \bar{e}) \geq 0$$

Integrating over the previous period type, θ_- , we get

$$\phi + \int a(\theta_-)\varphi(\theta_-)d\theta_- + p^{\text{FP}}(x^{\text{FB}}(\bar{\theta}) - \bar{e}) \geq 0$$

Using the market clearing condition

$$\int a(\theta_-)\varphi(\theta_-)d\theta_- = b$$

The condition rewrites

$$\phi + b + p^{\text{FP}}(x^{\text{FB}}(\bar{\theta}) - \bar{e}) \geq 0$$

Then there exists a debt level, b_{bliss} , such that

$$\phi + b_{\text{bliss}} + p^{\text{FP}}(x^{\text{FB}}(\bar{\theta}) - \bar{e}) = 0$$

Note that the optimal consumption decision (46) implies that the constraint is slack for any $\theta < \bar{\theta}$. In other words, for any $b \geq b_{\text{bliss}}$, the constraint never binds for any type $\theta \in (\underline{\theta}, \bar{\theta})$ implying that $\pi(b) = 0$ and $V(b) = V_{\text{bliss}} \equiv V(b_{\text{bliss}})$. By the same token, for any $b < b_{\text{bliss}}$, there exists $\tilde{\theta} \in (\underline{\theta}, \bar{\theta})$ such that the constraint binds for any $\theta \in [\tilde{\theta}, \bar{\theta})$, implying that $\pi(b) > 0$ and $V(b) < V_{\text{bliss}}$.

Q.E.D. ■

Proof of Proposition 2. Proposition 1 implicitly defines the optimal problem of the planner as

$$\begin{aligned} \max_{\{c_t, h_t, \tau_t, b_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (c_t - v(h_t) + V(b_{t+1})) \right] \\ q_t b_{t+1} = b_t + g_t - \tau_t A h_t \\ q_t = \beta + \pi(b_{t+1}) \\ c_t + g_t = A h_t \\ v'(h_t) = (1 - \tau_t) A \end{aligned}$$

Note that the only implementability constraint that has not already been incorporated in V and Q is the one for the supply of labor

$$v'(h_t) = (1 - \tau_t)A \quad (48)$$

Let us further simplify the planner's problem by solving (48) and the resource constraint for consumption and labor supply as functions of the tax rate. More specifically, let $H(\tau) \equiv (v')^{-1}\left(\frac{1-\tau}{A}\right)$ and $S(\tau) \equiv \tau AH(\tau)$ denote the equilibrium values of, respectively, labor supply and tax revenue, as functions of the tax rate. Note that it is straightforward to check that S is single-peaked—i.e., there is a Laffer curve—and attains its maximum value, \bar{s} , at $\tau = \bar{\tau}$ for some $\bar{\tau} \in (0, 1)$. For any $s \leq \bar{s}$, we thus have that, whenever the planner wishes to collect tax revenue equal to s , the tax rate that implements this goal is given by $\tau = T(s) \equiv \min\{\tau : S(\tau) = s\}$. Let $U(s) \equiv AH(T(s)) - v(H(T(s)))$ measure the resulting utility from consumption and leisure, as a function of tax revenue, and note that $U(s)$ is decreasing in s , reflecting the welfare cost of taxation. The problem simplifies to

$$\begin{aligned} \max_{\{s_t, \tau_t, b_{t+1}\}_{t=0}^{\infty}} \quad & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (U(s_t) + V(b_{t+1})) \right] \\ & q_t b_{t+1} = b_t + g_t - s_t \end{aligned}$$

where $q_t = \beta + \pi(b_{t+1})$.

Q.E.D. ■

Proof of Theorem 1. Assumption **B3** corresponds to the case $g > \hat{g}$ of Lemma 5, such that $\Psi_{bliss} > \gamma_{bliss}$. Then Proposition 11 applies and establishes part (i) of the theorem.

Q.E.D. ■

Proof of Proposition 3. Let us define $\Omega(b, \lambda^*) = V(b) + \lambda^* \pi(b)b$, where $\lambda^* = U'(s^*)$ and s^* solves $s^* + \pi(b^*)b^* = g + rb^*$. Note first that

$$\Omega_b(b, \lambda^*) = (\sigma(b) - 1) \pi(b) [\gamma(b) - \lambda^*]$$

Let us then recall that, in our benchmark a steady-state level, b^* below b_{bliss} exists if and only if $\psi_{bliss} > \gamma_{bliss}$, and it is then unique. Furthermore, the single-peakedness of $\pi(b)b$ guarantees that $\sigma(b) < 1$ and $\gamma(b) < 0$ for all $b < b_{seign}$, whereas $\sigma(b) > 1$ and $\gamma(b) > 0$ for all $b > b_{seign}$. Finally, the monotonicity of γ guarantees that $\gamma(b) > \gamma(b^*)$ for $b \in (b_{seign}, b^*)$, whereas $\gamma(b) < \gamma(b^*) = \lambda^*$. Together with the fact that $\gamma(b^*) = \psi(b^*) = \lambda^* > 0$, this implies that $\Omega_b(b, \lambda^*) > 0$ for all $b \in [\underline{b}, b^*)$ and $\Omega_b(b, \lambda^*) < 0$ for all $b \in [b^*, \bar{b})$, which proves that b^* solves $b^* = \operatorname{argmax}_b \Omega(b, \lambda^*)$ where $\lambda^* = \psi(b^*)$ and hence part (i).

Q.E.D. ■

Proof of Proposition 4. As stated in the text, the borrowing constraints are now affected by the presence of the capital stock. This implies that the program defined by the system 36–44 now includes the capital stock, and is completed by the first order condition on capital holdings. The solution of the problem

then gives both the price of capital, the level of capital and the price of bond as a function of public debt. The optimal labor decision implies that $h_t = H(\tau_t)$ such that the disutility of labor takes the form $U(\tau_t) = v(H(\tau_t))$. The budget constraint of the government then reads

$$q_t b_{t+1} = b_t + g_t - \tau_t A H(\tau_t) \iff Q(b_{t+1}) b_{t+1} = b_t + g_t - S(\tau_t)$$

The aggregate resource constraint of the model implies that $c_t = f(k_t) - k_{t+1} - g$. Hence, relying on the results of proposition 1, the welfare function writes

$$\mathcal{W}_t = \sum_{t=0}^{\infty} \beta^t (c_t - v(h_t) + V(b_{t+1}))$$

Using the aggregate resource constraint, we obtain

$$\begin{aligned} \mathcal{W}_t &= \sum_{t=0}^{\infty} \beta^t (f(k_t) - g_t - k_{t+1} - v(H(\tau_t)) + V(b_{t+1})) \\ &= \sum_{t=0}^{\infty} \beta^t (\beta f(k_{t+1}) - g_t - k_{t+1} - U(\tau_t) v) + f(k_0) \end{aligned}$$

using the result that since k_{t+1} is a function, $g_k(\cdot)$, of b_t , this rewrites

$$= \sum_{t=0}^{\infty} \beta^t (\beta g_c(b_t) - g_t - k_{t+1} - U(\tau_t) + V(b_{t+1})) + f(k_0)$$

where, $g_c(b_t) = \beta f(g_k(b_t)) - g_k(b_t)$ such that

$$= \sum_{t=0}^{\infty} \beta^t (\mathcal{V}(b_{t+1}) - U(\tau_t)) + f(k_0) + f(g_k(b_0)) - \frac{g_k(b_0)}{\beta}$$

where $\mathcal{V}(b) \equiv V(b) + \beta g_c(b) - g$. The crowding out effect is a direct consequence of the first order condition on capital and the properties of functions $f(\cdot)$ and $\pi(\cdot)$.

Q.E.D. ■

Proof of Proposition 5. The proposition is a direct consequence of the fact that $\frac{\dot{\lambda}}{\lambda} = \bar{\pi} > 0$.

Q.E.D. ■

Proof of Proposition 6. See discussion in main text.

Q.E.D. ■

Proof of Proposition 7. See discussion in main text.

Q.E.D. ■

Proof of Proposition 8. Let us start with the entrepreneur. He chooses his production plan by solving the following problem:

$$\begin{aligned} &\max_{k \geq 0, n \geq 0} [\theta f(k, n) + (1 - \delta)k - pk - wn] \\ \text{subject to } &z \leq \phi + a + \xi_k k + \xi_y \theta f(k, n) \\ &z = p(k - \kappa) \end{aligned}$$

Using the second constraint in the first one, and defining $x \equiv a + p\kappa$, as the net worth in period t , we obtain that the profit of the entrepreneur net of investment and labor costs is

$$\begin{aligned} \omega(x, p, w; \theta) &\equiv \max_{k \geq 0, n \geq 0} [\theta f(k, n) + (1 - \delta)k - pk - wn] \\ &\text{subject to } pk \leq \phi + x + \xi_k k + \xi_y \theta f(k, n) \end{aligned}$$

The production plan consists of the demand for labor, $n(x, p, w; \theta)$, and the demand for capital, $k(x, p, w; \theta)$. The aggregate quantities are

$$\mathbf{n}(x, p, w) = \int n(x, p, w; \theta) \varphi(\theta) d\theta \quad (49)$$

$$\mathbf{k}(x, p, w) = \int k(x, p, w; \theta) \varphi(\theta) d\theta \quad (50)$$

The problem of the household is

$$\begin{aligned} \max \quad & \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t (c_{it} - v(h_{it})) \right] \\ \text{s.t.} \quad & c_{it} + \kappa_{it+1} + q_t a_{it+1} = a_{it} + p_t \kappa_{it} + (1 - \tau_t) w_t h_{it} + \omega_{it} \end{aligned}$$

where we assumed that $a_{it} < \phi + \kappa_{it}$. ω_{it} denotes the profit received by household i . Note that (i) due to the linearity of the utility of consumption, all households supply the same amount of hours; and (ii) $\mathbb{E}[c_{it}]$ is aggregate consumption, c_t . Use the asset market clearing condition $\int a_{it} di = b_t$, let $\kappa_t \equiv \int \kappa_{it} di$ denote aggregate investment, and define

$$\Omega(x, p, w) \equiv \beta \int \omega(x, p, w; \theta) \varphi(\theta) d\theta.$$

The problem of the representative household can then be expressed as follows:

$$\begin{aligned} \max \quad & \sum_{t=0}^{\infty} \beta^t (c_t - v(h_t)) \\ \text{s.t.} \quad & c_t + \kappa_{t+1} + q_t b_{t+1} = b_t + p_t \kappa_t + (1 - \tau_t) w_t h_t + \Omega(x_t, p_t, w_t) \end{aligned}$$

where $x_t = b_t + p_t \kappa_t$. The first order conditions are given by

$$v'(h_t) = (1 - \tau_t) w_t \quad (51)$$

$$q_t = \beta(1 + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1})) \quad (52)$$

$$1 = \beta(1 + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1}) p_{t+1}) \quad (53)$$

where the last two conditions imply that $p_{t+1} = 1/q_t$, reflecting arbitrage between financial assets and physical capital. Notwithstanding this fact, the interest rate is lower than $1/\beta$ when $\Omega_x(\cdot) > 0$.

Clearing the labor and capital markets ($h_t = n_t$ and $k_t = \kappa_t$) implies

$$v'(\mathbf{n}(b_t + p_t k_t, p_t, w_t)) = (1 - \tau_t) w_t$$

$$k_t = \mathbf{k}(b_t + p_t k_t, p_t, w_t)$$

which can be solved for the wage $w(b_t, k_t, \tau_t)$ and the price of capital $p(b_t, k_t, \tau_t)$. Using them in the aggregate decisions for labor and capital, we have

$$h_t = H(b_t, \tau_t) \text{ and } k_t = K(b_t, \tau_t) \quad (54)$$

so that

$$w_t = W(b_t, \tau_t) \text{ and } p_t = P(b_t, \tau_t) \quad (55)$$

Likewise, using the resource constraint, we get

$$c_t = \theta f(k_t, n_t) + (1 - \delta)k_t - k_{t+1} - g = \tilde{C}(b_t, \tau_t) - k_{t+1} \quad (56)$$

Using (54) and (56) in the welfare function, we get

$$\sum_{t=0}^{\infty} \beta^t \left(\tilde{C}(b_t, \tau_t) - \frac{k_t}{\beta} - v(H(b_t, \tau_t)) \right) + \frac{K(b_0, \tau_0)}{\beta}$$

which can be written as

$$\sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t) + \frac{K(b_0, \tau_0)}{\beta}$$

Likewise, using the preceding results in (52), we get

$$q_t = Q(\tau_{t+1}, b_{t+1})$$

$$\tau_t w_t h_t - g = \tau_t W(b_t, \tau_t) H(b_t, \tau_t) - g = S(\tau_t, b_t)$$

and the government budget is

$$Q(\tau_{t+1}, b_{t+1}) b_{t+1} = b_t - S(\tau_t, b_t)$$

Hence, the problem of the central planner reduces to

$$\max_{\{\tau_t, b_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t)$$

$$s.t. \quad Q(\tau_{t+1}, b_{t+1}) b_{t+1} = b_t - S(\tau_t, b_t)$$

Q.E.D. ■

E A Simple Analytic Example

In this appendix, we use a simplified version of our model to obtain a sharper characterization of the functions (V, π) and of the optimal long-run level of public debt. This example uses log utility for the morning good and two types.

The household solves the problem

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [c_{i,t} + \theta_{i,t} u(x_{i,t}) - v(h_{i,t})] \quad (57)$$

subject to

$$c_{i,t} + p_{i,t} x_{i,t} + q_t b_{i,t+1} = b_{i,t} + (1 - \tau_t) w_t h_{i,t} + p_t \bar{e}_i \quad (58)$$

$$p_t(x_{i,t} - e_i) \leq \xi + b_{i,t} \quad (59)$$

where $u(x) = \log x$ and θ_{it} is i.i.d., drawn from the binary support $\{\theta_H, \theta_L\}$, for some $\theta_H > \theta_L > 0$. Let φ be the share of high types in the population and, to simplify the exposition, set $\theta_L = 1$, $\theta_H = \vartheta > 1$, $e_H = 0$, and $e_L = \bar{e}/(1 - \varphi)$. The rest of the notation is identical to that used in Section 2.1.

In equilibrium, the borrowing constraint (59) can bind at most for the high type. Letting μ_t be the associated multiplier, we can thus write the conditions that characterize the equilibrium in the market for the afternoon good in period t as follows:

$$\partial u'(x_{Ht}) = p_t(1 + \mu_t) \quad (60)$$

$$u'(x_{Lt}) = p_t \quad (61)$$

$$p_t x_{Ht} \leq \xi + b_t \quad (62)$$

$$\mu_t \geq 0 \quad (63)$$

$$\mu_t(\xi + b_t - p_t(x_{Ht} - \bar{e})) = 0 \quad (64)$$

$$\varphi x_{Ht} + (1 - \varphi)x_{Lt} = \bar{e} \quad (65)$$

The Euler condition for the optimal savings in period t , on the other hand, reduces to

$$\pi_t \equiv q_t - \beta = \beta \varphi \mu_{t+1} \geq 0 \quad (66)$$

Let (x_H^*, x_L^*) denote the first-best allocation. This solves

$$\partial u'(x_H^*) = u'(x_L^*) \quad (67)$$

$$\varphi x_H^* + (1 - \varphi)x_L^* = \bar{e} \quad (68)$$

and trivially satisfies $x_H^* > \bar{e} > x_L^*$ and $\partial x_H^*/\partial \vartheta > 0 > \partial x_L^*/\partial \vartheta$. In particular, using the assumption $u(x) = \log x$, we get

$$x_H^* = \frac{\vartheta \bar{e}}{\varphi \vartheta + 1 - \varphi} \quad \text{and} \quad x_L^* = \frac{\bar{e}}{\varphi \vartheta + 1 - \varphi}.$$

Clearly, this allocation can be attained in equilibrium if and only if

$$p_t = \vartheta u'(x_H^*) \quad \text{and} \quad p_t x_H^* \leq \xi + b_t.$$

Therefore, if we define

$$b_{bliss} \equiv \vartheta u'(x_H^*) x_H^* - \xi = \vartheta - \xi,$$

we immediately have that $b_t \geq b_{bliss}$ is sufficient for the borrowing constraint not to bind ($\mu_t = 0$) and the first best allocation to obtain.

And conversely, when $b_t < b_{bliss}$, the first best allocation is unattainable, the borrowing constraint binds, and the equilibrium yields

$$x_{Ht} = \frac{(b + \xi)\bar{e}}{\varphi(b + \xi - 1) + 1} e, \quad x_{Lt} = \frac{(\varphi + 1)\bar{e}}{\varphi(b + \xi - 1) + 1}, \quad \text{and} \quad \mu = \frac{\vartheta - b - \xi}{b + \xi + \varphi - 1}.$$

Using the definition of b_{bliss} , we can rewrite the above as

$$x_{Ht} = \frac{\bar{e}(b_t - b_{bliss} + \vartheta)}{\varphi(b_t - b_{bliss} + \vartheta) + 1 - \varphi}, \quad x_{Lt} = \frac{\bar{e}}{\varphi(b_t - b_{bliss} + \vartheta) + 1 - \varphi}, \quad \text{and} \quad \mu_t = \frac{(b_{bliss} - b_t)}{\vartheta - (b_{bliss} - b_t)},$$

which makes clear how the equilibrium allocation converges monotonically to the first-best counterpart, and how μ_t converges monotonically to 0 from above, as b_t converges to b_{bliss} from below.

Using these results, we then also have the following closed-form solution for (V, π) : for $b \geq b_{bliss}$, $\pi(b) = 0$ and $V(b) = V_{bliss}$, where

$$V_{bliss} \equiv \beta \{ \varphi \vartheta u(x_H^*) + (1 - \varphi) u(x_L^*) \} = \beta \{ v - (\varphi \vartheta + 1 - \varphi) \log(\varphi \vartheta + 1 - \varphi) \}$$

and $v \equiv (\varphi \vartheta + 1 - \varphi) \log(e)$; and for $b < b_{bliss}$,

$$\pi(b) = \beta \varphi \frac{(b_{bliss} - b)}{\vartheta - (b_{bliss} - b)} > 0$$

and

$$V(b) = \beta \{ v + \varphi \vartheta \log(\vartheta + b - b_{bliss}) - (1 - \varphi + \varphi \vartheta) \log(\varphi(\vartheta + b - b_{bliss}) + 1 - \varphi) \}.$$

We therefore reach the following result:

Lemma 6. *Suppose $\xi < \vartheta$. There exists a threshold $b_{bliss} > 0$, given by $b_{bliss} = \vartheta - \xi$, such that the following properties hold for all $b < b_{bliss}$:*

$$\pi(b) > 0, \quad \pi'(b) < 0, \quad \pi''(b) > 0,$$

$$V(b) < V_{bliss}, \quad V'(b) > 0, \quad V''(b) < 0.$$

For $b \geq b_{bliss}$, on the other hand, $\pi(b) = 0$ and $V(b) = V_{bliss}$. Finally, a tighter financial friction, or lower private collateral (lower ξ), increases b_{bliss} and uniformly raises both $V(b)$ and $\pi(b)$ for all $b < b_{bliss}$.

Proof. The properties of π and V with respect to b follow directly from their closed-form characterization. And the effect of ξ follows from the fact that $b_{bliss} = \vartheta - \xi$ along with the fact that, for any b_{bliss} and any $b < b_{bliss}$, π and V are increasing in b_{bliss} and otherwise invariant to ξ . \square

Consider now the problem introduced in Section 2.4, which as shown in Section 4.1 also characterizes the optimal steady state. In particular, consider the following two objects:

$$b_{seign} = \arg \max_b \pi(b)b$$

$$b^* = \arg \max_b \Omega(b, \lambda)$$

where $\Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b$ and $\lambda > 0$. The following result can then be shown.

Lemma 7. $\pi(b)b$ and $\Omega(b, \lambda)$ are strictly concave in $b \in [0, b_{bliss}]$ and their maxima satisfy at $0 < b_{seign} < b^* < b_{bliss}$.

Proof. Consider $g(b) \equiv \pi(b)b$ and note that

$$g'(b) = \pi(b) + \pi'(b)b \quad \text{and} \quad g''(b) = 2\pi'(b) + \pi''(b)b$$

Using the fact that $\pi''(b) = -\frac{2\pi'(b)}{b+\xi}$, we get that

$$g''(b) = 2\pi'(b) \frac{\xi}{b+\xi} < 0$$

which proves that $g(b) \equiv \pi(b)b$ is concave. Next, note that $g'(0) = \beta\varphi \frac{\vartheta-\xi}{\xi} > 0$ and $g'(b_{bliss}) = \pi'(b_{bliss})b_{bliss} = -\beta\varphi \frac{\vartheta-\xi}{\theta} < 0$. It follows that b_{seign} is the unique solution to $g'(b) = 0$ and is strictly between 0 and b_{bliss} .

Consider now $\Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b$. Its concavity follows directly from the concavity of $V(b)$, which was established in the previous result, and the concavity of $g(b) = \pi(b)b$, which was just established. It follows that b^* is the unique solution to $\partial \Omega(b, \lambda) / \partial b = 0$. Furthermore, because $g'(b_{seign}) = 0$, $g'(b_{bliss}) < 0$, $V'(b_{seign}) > 0$, and $V'(b_{bliss}) = 0$, we have that $\partial \Omega(b, \lambda) / \partial b > 0$ at $b = b_{seign}$ and $\partial \Omega(b, \lambda) / \partial b < 0$ at $b = b_{bliss}$, and therefore that b^* is strictly between b_{seign} and b_{bliss} . \square

This result echoes the properties we establish in Section 3 for the steady state. But because we now have a simple closed-form characterization of Ω , we can go a step further to study the comparative statics of b^* with respect to the underlying primitives. Using our closed-form solution for π and V along with the fact $b_{bliss} = \theta - \xi$, we can show that

$$\frac{\partial^2 \Omega}{\partial b \partial \theta} = \frac{\beta\varphi(1-\varphi)}{(b+\xi)(\varphi(b+\xi)+1-\varphi)} + \lambda \frac{\beta\varphi\xi}{(b+\xi)^2} > 0$$

Furthermore,

$$\left. \frac{\partial^2 \Omega}{\partial b \partial \lambda} \right|_{b=b^*} = \lambda g'(b^*) < 0$$

by the fact that $b^* > b_{seign}$. Applying the Implicit Function Theorem (IFT), we then have that

$$\frac{\partial b^*}{\partial \theta} > 0 \quad \text{and} \quad \frac{\partial b^*}{\partial \lambda} < 0.$$

Finally, consider how b^* varies with ξ . Note that

$$\frac{\partial^2 \Omega}{\partial b \partial \xi} = \frac{\partial^2 V}{\partial b \partial \xi} + \lambda \frac{\partial^2 g}{\partial b \partial \xi}$$

and $\frac{\partial^2 V}{\partial b \partial \xi} = V''(b) < 0$. But because $\frac{\partial^2 g(b)}{\partial b \partial \xi} = \frac{\beta \varphi \theta (b - \xi)}{(\beta + \xi)^2}$ changes sign with the position of b relative to ξ , the effect of ξ on b^* is generally ambiguous. In particular, we have found numerically that b^* is inversely U-shaped with respect to ξ .

Proposition 16. *The optimal quantity of public debt increases with the size of the liquidity shocks (θ), decreases with the value of fiscal space (λ), and is generally non-monotonic in the amount of private collateral (ξ).*

Although b^* can be decreasing in ξ , which means that more private collateral can crowd out the government-provided collateral, there is no complete crowding out: an increase in ξ always increases total collateral, $b^* + \xi$.³⁵ It then also follows that, at the optimal quantity of public debt, more private collateral depresses the liquidity premium ($\frac{\partial \pi(b^*)}{\partial \xi} < 0$), whereas the converse is true with an aggravation of liquidity needs ($\frac{\partial \pi(b^*)}{\partial \theta} > 0$).

To conclude, these findings complement the intuitions developed in Section 2.4. Strictly speaking, they do not apply to the steady state of the infinite-horizon model, because they treat λ as exogenous. But we can use the government budget evaluated at the steady state to obtain λ as an increasing function of b , an increasing function of g , and a decreasing function of π (and thereby a decreasing function of θ and an increasing function of ξ). We can then readily translate the result to the steady-state level of debt, modulo the replacement of λ with g . That is, the value of fiscal space is re-parameterized by g , but the comparative statics with respect to θ and ξ go through.

³⁵To see this, let $z \equiv b + \xi$ and re-express V , π , and Ω as functions of z instead of b :

$$\pi(b) = \tilde{\pi}(z) \equiv \beta \varphi \frac{\vartheta - z}{z}$$

$$V(b) = \tilde{V}(z) \equiv \beta \{v + \varphi \vartheta \log(z) - (1 - \varphi + \varphi \vartheta) \log(\varphi z + 1 - \varphi)\}$$

$$\Omega(b, \lambda) = \tilde{\Omega}(z, \lambda) \equiv \tilde{V}(z) + \lambda \tilde{\pi}(z)(z - \xi)$$

Because \tilde{V} and $\tilde{\pi}$ are invariant to ξ , it is immediate that $\frac{\partial^2 \tilde{\Omega}}{\partial z \partial \xi} = -\lambda \tilde{\pi}'(z) > 0$, which via the IFT implies that $z^* \equiv \operatorname{argmax}_z \tilde{\Omega}(z, \lambda) = b^* + \xi$ increases with ξ . In fact, because the property that V and π are invariant to ξ conditional on z applies generally, so does the result that z^* increases with ξ .