



HAL
open science

Public Debt as Private Liquidity: Optimal Policy

George-Marios Angeletos, Fabrice Collard, Harris Dellas

► **To cite this version:**

George-Marios Angeletos, Fabrice Collard, Harris Dellas. Public Debt as Private Liquidity: Optimal Policy. 2022. hal-03763095

HAL Id: hal-03763095

<https://hal.science/hal-03763095>

Preprint submitted on 29 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Debt as Private Liquidity: Optimal Policy*

George-Marios Angeletos

MIT

Fabrice Collard

Toulouse School of Economics

Harris Dellas

University of Bern

May 2022

Abstract

We study optimal policy in an economy in which public debt serves as collateral or buffer stock. Issuing more public debt can raise welfare by easing the underlying friction; but it also reduces the premium that the private sector is willing to pay for this service, which in turn raises interest rates. This trade off shapes the optimal long-run quantity of public debt. It justifies larger deficits during financial crises. It may subsume other considerations, such as whether public debt crowds out, or in, capital. And it clarifies the circumstances under which a low risk-free rate represents an opportunity of cheap borrowing by the government.

*This paper supersedes an earlier paper, entitled “Optimal Public Debt Management and Liquidity Provision” (Angeletos et al, 2016). We are grateful to Behzad Diba for his collaboration on the early stages of this project; to Pedro Teles and Per Krusell for discussing our paper in conferences; to numerous seminar participants for their comments; and, last but not certainly not least, to Harald Uhlig and three anonymous for their extensive feedback. Angeletos also thanks the University of Bern, the Study Center Gerzensee, and the Swiss Finance Institute for their hospitality.

1 Introduction

The persistently low real interest rates on US public debt have prompted calls for higher deficit spending to exploit the low cost of borrowing (e.g., Delong, 2011, 2015; Krugman, 2016, 2020). This argument seems particularly tempting in turbulent times, such as the Great Recession or the COVID pandemic: a flight to safety during such episodes appears to reduce the government’s cost of borrowing exactly when fiscal stimulus is needed the most. However, this argument finds no place in the basic, complete-markets, Ramsey paradigm (Barro, 1979; Lucas and Stokey, 1983): in this benchmark, there is never a real opportunity for “cheap” government borrowing, and tax smoothing is optimal regardless of the level and the cyclicity of interest rates. Our objective in this paper is to revisit this issue in the presence of financial frictions and to draw the implications for optimal fiscal policy.

Financial frictions are not only conspicuous in the real world. They also endow public debt with a special function: government borrowing contributes to the aggregate supply of assets that private agents—households, firms, financial institutions—can use as collateral or buffer stock, which in turn can ease the frictions (Woodford, 1990; Aiyagari and McGrattan, 1998; Holmström and Tirole, 1998). Corroborating this idea, Krishnamurthy and Vissing-Jorgensen (2012) document that the spread between government bonds and high-grade corporate bonds, an empirical measure of the market’s valuation of the “services” provided by public debt, decreases with the quantity of public debt.

This mechanism allows government borrowing to be cheap in a sense that we will make precise, but does not provide an unlimited “free lunch.” Instead, it gives rise to a novel policy trade off: the flip side of easing the friction and improving market allocations is an increase in the cost of government borrowing. Our contribution is to show how this trade off, in conjunction with the desire to smooth tax distortions, shapes optimal policy in both short and the long run.

Framework. We work primarily with a reduced-form Ramsey problem, which is similar to that in Barro (1979), except that it lets public debt be non-neutral for two reasons in addition to distortionary taxation. First, it carries a lower interest rate than the underlying social discount rate by an amount that reflects the *private* value of the liquidity, collateral, or other “convenience” services provided by it. Second, for given taxes welfare also depends on public debt, reflecting the corresponding *social* value of these services.

In its most literal interpretation, our problem corresponds to an economy with public debt in the utility function. But the same reduced form also characterizes two other economies: one in which public debt plays a dual role as a vehicle for life-cycle savings (Diamond, 1965) and as a buffer stock against idiosyncratic risk (Aiyagari and McGrattan, 1998); and another in which public debt serves as collateral (Holmström and Tirole, 1998). Although the offered micro-foundations are highly stylized, they help illustrate how the trade off we emphasize may transcend various applications. They also corroborate the following interpretations: the “service” provided by public debt is to ease a financial friction within the private sector; the “scarcity” of this service is inversely related to the private sector’s own capacity to cre-

ate assets or “inside” liquidity; and the private value of this service may differ from its social counterpart because of pecuniary externalities in collateral constraints.

Main lessons. The optimal policy is dictated by the interplay of three forces: (i) the desire to smooth the tax distortion over time; (ii) the desire to ease the financial friction and improve the allocation of resources within the private sector; and (iii) the desire to maximize the rents, or “seigniorage,” that the government can extract from the private sector in the form of a low interest rate on its debt.

While the first force reigns supreme in the textbook Ramsey paradigm (Barro, 1979; Lucas and Stokey, 1983), here it has to be balanced against the other two. The second force, in particular, calls for raising public debt towards a certain “bliss” point, defined by the level of public debt that satiates the economy’s demand for collateral or, more generally, that equates the interest rate on public debt with the underlying social discount rate. The third force, on the other hand, involves a Laffer curve: the relevant rents are maximized at a quantity of public debt strictly less than the bliss point.

Because taxes are smooth in a steady state, one may expect the long run to be shaped solely by the other two forces. This intuition has to be qualified in two ways. First, there is a fixed-point relation between the steady-state quantity of public debt and the relative importance of the aforementioned forces. Second, the desire to smooth taxes acts as an adjustment cost that, in certain cases, may help rationalize multiple locally determinate steady states, each with its own basin of attraction. Notwithstanding these qualifications, the aforementioned intuition applies: the long-run quantity of public debt trades off the desire to ease the financial friction against the desire to depress interest-rate costs.

The same trade off also figures in the optimal response to shocks, except that now the desire to smooth taxes naturally gains prominence. Consider, e.g., an unanticipated, uninsured, positive shock to government spending (a “war”). In the textbook 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, the optimal response is front-loaded: taxes increase relatively more early on in order to keep interest rates on debt low, enabling a smaller tax burden later on.

Consider next an inefficient recession, modeled as a labor-wedge shock that not only reduces output but also justifies a fiscal stimulus in the form of a payroll tax cut. Suppose further that the financial friction is aggravated. This has an ambiguous effect on the trade-off we emphasize: while it encourages liquidity provision, it also raises the rents that can be extracted by preserving the shortage of collateral, which 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 argument 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 of the apparent drop in the government’s cost of borrowing. As already mentioned, this argument has no place in the textbook paradigm. But it makes sense insofar as a lower interest rate is a symptom of higher idiosyncratic risk, a shortage of collateral, or a flight to safety.

Additional points and related literature. The mechanisms we emphasize in this paper are consistent with an empirical literature that argues that the liquidity premium on public debt moves countercyclically and also decreases with its quantity (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood and Vayanos, 2014). Note, though, that this literature has focused on the interest-rate spread between government bonds and high-grade corporate bonds, which is only a small fraction of the relevant theoretical object: the “premium” that shapes optimal policy is the wedge between the interest rate paid on public debt and the underlying social discount rate, or equivalently between the former and the interest rate that would have obtained under complete markets.

These and a few additional lessons are made possible by nesting a handful of micro-founded, albeit stylized, examples in our reduced form. As we move across these examples, the precise market friction and the corresponding “service” of public debt changes—and so do a few other “details,” such as the possibility that public debt crowds out physical capital or inside liquidity (Aiyagari and McGrattan, 1998; Azzimonti and Yared, 2019). But the essence of the problem remains the same, and so do our results.

Turning to a technical aspect of our paper, we observe that the policy problem at hand is non-convex. As a result, the standard first-order approach fails and there generally exist multiple optimal steady states. We address these challenges by adapting the optimal control methods of Skiba (1978). We also explain why the approach taken in Aiyagari and McGrattan (1998) overestimates the tax burden of the services provided by public debt and, thereby, underestimates its optimal long-run quantity.

Last but not least, although our paper is rooted in the earlier works of Woodford (1990), Aiyagari and McGrattan (1998) and Holmström and Tirole (1998), our insights extend to the recent literature spurred by Blanchard (2019). The issues that preoccupy this literature, such as “ $r - g$ ” and debt sustainability (Brunnermeier, Merkel and Sannikov, 2021; Mehrotra and Sergeyev, 2020; Reis, 2021; Barro, 2021), inequality (Aguiar, Amador and Arellano, 2021), and the ZLB constraint (Mian, Straub and Sufi, 2022), are outside our scope. Nonetheless, insofar as these works emphasize financial frictions, they naturally inherit the trade off highlighted here. Related are also Canzoneri and Diba (2005) and Andolfatto and Martin (2018), which allow public debt to provide money-like services but do not share our focus on financial frictions and the relevant trade off; and especially Bassetto and Cui (2021), which investigates how this trade off relates to capital taxation, in an economy similar to Example 3 of our paper.

2 The Reduced-Form Policy Problem

This section introduces the reduced-form policy problem. As anticipated above, this is basically the same as that in Barro (1979), except for the introduction of a “service” provided by public debt.

Time is continuous, $t \in [0, \infty)$. Let $s(t)$, $b(t)$ and $r(t)$ denote, respectively, tax revenue, the stock of public debt, and the interest rate on it. The government’s flow budget constraint is given by

$$\dot{b}(t) = r(t)b(t) + g - s(t),$$

where g is government spending (time-invariant for simplicity). Welfare is given by

$$\mathcal{W} = \int_0^{+\infty} e^{-\rho t} [U(s(t)) + V(b(t))] dt,$$

where $\rho > 0$ is the social discount rate, $U(\cdot)$ captures the social cost of taxation, and $V(\cdot)$ captures the social value of the liquidity or other services provided by public debt. Finally, the interest rate satisfies

$$r(t) = \rho - \pi(b(t)),$$

where $\pi(\cdot)$ captures the premium that private agents pay in equilibrium for the aforementioned services.¹

Let $\bar{s} > 0$ represent the maximal feasible tax revenue, corresponding to the peak of the Laffer curve for taxes; let $\underline{s} \leq 0$ be an arbitrary lower bound; and let $\bar{b} \equiv \frac{\bar{s}-g}{\rho} > 0$ be the maximal sustainable level of debt. These are fixed parameters, like g and ρ .

The planner's problem is as follows:

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

$$\max \int_0^{+\infty} e^{-\rho t} [U(s(t)) + V(b(t))] dt \quad (1)$$

$$\text{subject to } \dot{b}(t) = (\rho - \pi(b(t)))b(t) + g - s(t) \quad \forall t \quad (2)$$

with initial condition $b(0) = b_0$, for some $b_0 \in [0, \bar{b}]$.

We finally impose the following restrictions on the reduced form:

Main Assumptions. [A0] U, V , and π are continuously differentiable.²

[A1] U is concave in s , with a maximum attained at $s = 0$.

[A2] There exists a $b_{\text{bliss}} \in (0, \bar{b})$ such that $V'(b), \pi(b) > 0$ if $b < b_{\text{bliss}}$, and $V'(b), \pi(b) \leq 0$ otherwise.

Assumptions A0 and A1 require little explanation: the former is technical; the latter captures the distortionary effects of taxation. In particular, Barro (1979) is nested by letting $\pi(b) = 0$ for all b (equivalently, $r = \rho$). What is new, therefore, is only Assumption A2. This assumption represents the role of public debt as an essential vehicle for life-cycle saving (Diamond, 1965), as a buffer stock against idiosyncratic risk (Aiyagari and McGrattan, 1998), or as a form of collateral (Holmström and Tirole, 1998).

3 A Detour: Stylized Micro-foundations

In this section, we corroborate the interpretation of our reduced-form problem, and shed additional light on the economics behind it, with the help of a few micro-founded examples. For the purposes of these examples, we momentarily switch from continuous time to discrete. Readers eager to see the main results can jump to Section 4.3.

¹Note that V is the *total* social value of public debt, V' is the corresponding *marginal* value, and π is the equilibrium counterpart of V' . We will later explain why π and V' can, but do not have to, coincide.

²To be precise, we allow V and π to be non-differentiable at $b = b_{\text{bliss}}$. Such a kink obtains in the second example from Section 3, because a borrowing constraint switches from binding to non-binding as b crosses b_{bliss} from below.

Example 1. Diamond (1965) meets Aiyagari and McGrattan (1998)

There are overlapping generations of two-period-lived households. Households work when they are young, consume in both periods of life, and do not care about future generations. In the version considered here, saving takes place only in a risk-free bond, whose net supply is controlled by the government; an extension discussed in Section 6 adds capital accumulation, without affecting the results. More crucially, individuals are subject to uninsurable idiosyncratic risk; this introduces a precautionary motive for saving and lets public debt play a similar role to that in Aiyagari and McGrattan (1998).

Setup. Individuals are indexed by (i, t) , where $t \in \{0, 1, \dots\}$ stands for their birth date and $i \in [0, 1]$ for their idiosyncratic identity. Life-time utility is given by

$$\mathcal{U}_{it} = c_{it}^y - v(h_{it}) + \delta \mathbb{E}_{it} [u(c_{it+1}^o)],$$

and the budget constraints in the two periods of life are given by

$$c_{it}^y + q_t a_{it} = (1 - \tau_t) w_t h_{it} \quad \text{and} \quad c_{it+1}^o = e_{it+1} + a_{it}.$$

c_{it}^y and c_{it+1}^o are the individual's consumption when young and old, respectively; h_{it} is her labor supply; a_{it} is her saving in the risk-free asset; q_t is the asset's price (i.e., the reciprocal of the gross interest rate between t and $t+1$); w_t is the wage; and e_{it+1} is a random endowment received when old. The latter is i.i.d. across (i, t) and takes values $e_{it} \in \{k - \epsilon, k + \epsilon\}$, with respective probabilities φ and $1 - \varphi$, for some $k > 0$, $\epsilon \in (0, k)$, and $\varphi \in (0, 1)$. Finally, u and v satisfy the following properties: $u' > 0$, $u'' < 0$, $u''' > 0$ (preferences exhibit risk aversion and prudence); $v' > 0$ and $v'' > 0$ (the disutility of labor is convex); $\lim_{c \rightarrow 0} u'(c) = \lim_{h \rightarrow \infty} v'(h) = \infty$ and $\lim_{c \rightarrow \infty} u'(c) = \lim_{h \rightarrow 0} v'(h) = 0$ (the usual Inada conditions hold).

Labor is employed by a representative competitive firm, which produces the final good according to a linear returns technology: $y_t = Ah_t$, for some fixed $A > 0$. It follows that, in equilibrium, the pre-tax wage is given by $w_t = A$ and firm profits are zero.

The model is completed by the government (the Ramsey planner). Its budget constraint is given by

$$b_{t-1} + g = q_t b_t + \tau_t w_t h_t \tag{3}$$

where b_{t-1} is the inherited stock of debt, $q_t b_t$ are the proceeds from new debt issuance, $s_t \equiv \tau_t w_t h_t$ is tax revenue, and g is the exogenous government spending. Finally, its objective is to maximize the following welfare measure:

$$\mathcal{W} \equiv \sum_{t=0}^{\infty} \beta^t \int \mathcal{U}_{it} di,$$

where $\beta \in (0, 1)$ is the social discount factor.

The demand for safe assets. Since young households are identical, they make the same choices: for all i and t , $c_{it}^y = c_t^y$, $h_{it} = h_t$, and $a_{it} = a_t$. Furthermore, optimal saving solves the following FOC:

$$q_t = \delta \mathbb{E}_t[u'(c_{i,t+1}^o)] = \delta \{\varphi u'(k - \epsilon + a_t) + (1 - \varphi)u'(k + \epsilon + a_t)\}. \quad (4)$$

Together with market clearing in the bond market ($a_t = b_t$), this gives

$$q_t = Q(b_t; \epsilon, k), \quad (5)$$

where the function Q is defined by the right-hand side of (4) replacing a_t with b_t . This function can be viewed interchangeably as the equilibrium price function for government debt and as the inverse of the aggregate demand for it. Risk aversion and prudence ($u'' < 0 < u'''$) imply the following properties:

Proposition 1. *The demand for public debt is downward slopping ($Q_b < 0$); it shifts up with idiosyncratic risk ($Q_\epsilon > 0$); and it shifts down with k ($Q_k < 0$).*

These properties translate to our reduced form by letting $\pi(b) \equiv Q(b) - \beta$ and help proxy the demand for safe assets in a larger class of models. In particular, $Q_\epsilon > 0$ and $Q_b < 0$ mimic two core predictions of the incomplete-markets literature: that the equilibrium risk-free rate falls with more idiosyncratic risk (Bewley, 1980; Aiyagari, 1994) and increases with the aggregate supply of the “buffer stock” provided in the form of government debt (Aiyagari and McGrattan, 1998). Similarly, $Q_k < 0$ mimics the role of physical capital and/or that pledgeable income à la Holmström and Tirole (1998); these points will become clear in Example 2 below and also in the extension with capital discussed in Section 6. Finally, note that here $\pi(b)$ can turn negative for b high enough, but this is not the case in Examples 2 and 3 below: there, $\pi(b)$ is positive whenever a collateral constraint binds and zero otherwise.

Let us now switch attention from the *private* value of public debt to the *social* counterpart. By the resource constraint ($c_t^y + \int c_{it}^o di = Ah_t + k$) and aggregation of the old-age budgets ($\int c_{it}^o di = b_{t-1} + k$), we have $c_t^y = Ah_t - b_{t-1}$ and

$$\mathcal{W} \equiv \sum_{t=0}^{\infty} \beta^t \int \mathcal{U}_{it} di = \sum_{t=0}^{\infty} \beta^t \{Ah_t - v(h_t) + V(b_{t-1}; \epsilon, k)\},$$

where

$$V(b; \epsilon, k) \equiv \varphi \delta u(k - \epsilon + b) + (1 - \varphi) \delta u(k + \epsilon + b) - b$$

defines the (total) social value of public debt. Note then that $\pi(b) \equiv Q(b)/\beta - 1 = V'(b)$, that is, the (marginal) private and social values coincide. But this does not have to be true in general—indeed, it is not true in Examples 2 and 3 below, because of a pecuniary externality.

Reduced form. The Inada conditions on u imply that $\lim_{b \rightarrow \infty} Q(b) = 0$ and $Q(0) > \beta$ for sufficiently high levels of risk (ϵ close to k). It follows that there exists a $b_{\text{bliss}} > 0$, defined by $Q(b_{\text{bliss}}) = \beta$, such that $V'(b) = \pi(b) > 0$ for $b < b_{\text{bliss}}$ and $V'(b) = \pi(b) < 0$ for $b > b_{\text{bliss}}$. That is, the example satisfies Assumption A2. The reduced form is completed by letting $U(s)$ measure the joint utility from consumption and labor obtained in equilibrium when tax revenue equals s and verifying that U satisfies Assumption A1.

Proposition 2. *The optimal path for taxes and public debt solves the following problem:*

$$\max_{\{s_t, b_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t [U(s_t) + V(b_{t-1})] \quad (6)$$

$$\text{subject to } Q(b_t)b_t = b_{t-1} + g - s_t \quad (7)$$

which is nested in our reduced form. Furthermore, Assumptions A0–A2 are readily satisfied.

We have so far emphasized the function of government debt as a buffer stock against idiosyncratic risk. But in the present example, government debt has one more function: similarly to Diamond (1965), it regulates the allocation of aggregate consumption between the young and the old. It follows that, in the present example, Q and V encapsulate *both* of these functions. The example provided below shuts down the second function, while also recasting the first one in a way that bridges the analysis to Holmström and Tirole (1998). Put together, these examples illustrate that, while the distinction between the three functions of public debt under consideration— collateral, buffer stock, or vehicle for life-cycle savings— may be important in general, it is not essential for our own purposes.

Example 2. An infinite-horizon variant with a Holmström and Tirole (1998) flavor

As already mentioned, our second example seeks to stylize the role of public debt as collateral. It does so by introducing a trading/financial friction and by letting private assets ease that friction.

Setup. Households are infinitely-lived and there are two edible goods, consumed in different halves of each period. The first good is the (exogenous) fruit of a tree that becomes ripe in the “morning” and whose trade is subject to a financial friction. The second good is the (endogenous) output of a representative firm, which is produced in the afternoon with the labor of the households. Each good has to be consumed in the respective sub-period, or else it perishes.³

Let h_{it} , x_{it} , and c_{it} denote, respectively, labor supply, consumption of the morning good, and consumption of the afternoon good. The household’s expected life-time utility is given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \{c_{it} + \theta_{it} \log x_{it} - v(h_{it})\} \right], \quad (8)$$

and its budget constraint at the end of each period is given by

$$c_{it} + p_t x_{it} + q_t a_{it} = a_{it-1} + (1 - \tau_t) w_t h_{it} + p_t e_{it}, \quad (9)$$

where $\beta \in (0, 1)$, v satisfies the same properties as in the previous example, a_{it} is the household’s saving in the risk-free asset, q_t is its price, p_t is the price of the morning good, and e_{it} and θ_{it} are the household’s idiosyncratic endowment of and taste for the morning good. These shocks serve the purpose of

³As made clear in the proof of Proposition 3, the afternoon good maps to the standard consumption good in the textbook Ramsey paradigm, while the morning good helps introduce the friction of interest and thereby the functions π and V of our reduced form.

inducing a desire for trading the morning good between the agents. For simplicity, we assume only two idiosyncratic states: either the household has a high taste for the morning good but no endowment, $(\theta_{it}, e_{it}) = (1 + \epsilon, 0)$, or it has a low taste and a positive endowment, namely $(\theta_{it}, e_{it}) = (1, 1/(1 - \varphi))$, with respective probabilities φ and $1 - \varphi$, for some $\varphi, \epsilon \in (0, 1)$.⁴

In addition to (9), the household faces a liquidity, or collateral, constraint in the morning of each period. Let $z_{it} \equiv p_t(x_{it} - e_{it})$ denote the household's net trade of the morning good. 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:

$$z_{it} \equiv p_t(x_{it} - e_{it}) \leq \xi y_t^{\text{def}} + a_{it-1} \quad (10)$$

where y_t^{def} denotes the income received in the (off-equilibrium) event of default.⁵ The presence of a_{it-1} on the right hand side of this constraint explains the precise sense in which holdings of the risk-free asset—and thereby public debt—serve as collateral in our example.⁶

Finally, the production side of the economy is the same as in the previous example, the government's budget constraint is the same as well, and its objective is the maximization of (8) behind the veil of ignorance, i.e., for a random i .

Reduced form. The collateral constraint (10) can bind only under the high taste shock. When public debt is absent ($b_t = 0$), this in turn happens if and only if the “inside collateral” is sufficiently scarce, in the sense that $\xi < \bar{\xi}$ for some $\bar{\xi} > 0$. Assuming that this is true (or else the friction is inactive and policy problem is trivial), and after some tedious derivations, we can reach the following result.

Proposition 3. (i) *Proposition 2 extends to the present economy. That is, the planner's problem is again nested in our reduced form and Assumptions A0-A2 continue to hold.*

(ii) *For $b < b_{\text{bliss}}$, the financial friction is binding and the private value of collateral is higher than the social one, implying that $\pi(b) > V'(b) > 0$. For $b > b_{\text{bliss}}$, on the other hand, the financial friction is not binding and $\pi(b) = V'(b) = 0$. Finally, the threshold b_{bliss} and the value of π for all $b < b_{\text{bliss}}$ increase with ϵ , the amount of idiosyncratic risk, and decrease with ξ , the pledgeable income.*

⁴This particular specification of the idiosyncratic risk facilitates part (ii) of Proposition 3, but is not needed for part (i), which is all we really need. The linearity of preferences in the afternoon good, on the other hand, plays a similar role as in Lagos and Wright (2005): it guarantees that the cross-sectional distribution of wealth is not a relevant state variable.

⁵For simplicity, we let income not be taxed in the event of default. It follows that $y_t^{\text{def}} \equiv w_t h_t^{\text{def}}$, with $h_t^{\text{def}} \equiv (v')^{-1}(w_t)$.

⁶Applying the same logic to inter-period borrowing yields the additional constraint $-a_{it} \leq \xi y_t^{\text{def}}$. In the present version of Example 2, this constraint never binds in equilibrium, because $a_{it} = b_t$ by equilibrium and $b_t \geq 0$ by assumption. We can thus ignore it and focus on (10). An earlier version introduced additional heterogeneity so as to accommodate “inside liquidity” (i.e., issuance of private debt by some agents, which in turn could be used as buffer stock or collateral by other agents, in place of public debt). This allowed the constraint $-a_{it} \leq \xi y_t^{\text{def}}$ to bind sometimes but did not change the essence.

Part (i) verifies that our reduced-form approach is still applicable. Part (ii) sheds additional light on the role of debt in the present example. $\pi(b)$ is now directly related to the Lagrange multiplier on (10). Put differently, $\pi(b)$ measures the equilibrium value of collateral, echoing Holmström and Tirole (1998). This value is strictly positive for $b < b_{\text{bliss}}$, but becomes null for $b > b_{\text{bliss}}$, because after that point (10) does not bind for either the low or the high taste shock and the economy's demand collateral is satiated.

The last property differs from its counterpart in our first example, where $\pi(b)$ was *negative* for $b > b_{\text{bliss}}$, reflecting excessive saving by the young. In other words, the region $b > b_{\text{bliss}}$ has changed from “harmful excess” in the previous example to “harmless satiation” in the present one. Another difference is that the present example features a wedge between the private and the social value of public debt, due to the pecuniary externality operating via p_t in (10). Intuitively, agents fail to internalize how their collateral pushes up the price of the morning good, which in turn tightens the constraints of others. The accommodation of these differences further illustrates the versatility of our reduced form.

Example 3. From consumer frictions to production frictions

In Online Appendix A we study a variation of the previous example, in which a constraint similar to (10) impedes the efficient allocation of capital and labor across firms. This changes the usage of collateral and lets b_t enter the economy's aggregate production function. The reduced form is a bit more complicated, but the essence remains the same.

4 Optimal Policy

This section contains our main results. We first shine the spotlight on the key policy trade off. We next characterize the optimal long-run target for public debt and the transitional dynamics towards it.

4.1 Preamble: Liquidity provision versus interest-rate suppression

To build intuition, consider momentarily a two-period version of the problem. Suppose further that any debt issued at $t = 1$ must be retired at $t = 2$. Under these simplifications, the problem reduces to

$$\begin{aligned} & \max_{s_1, b_1, s_2} \{ [U(s_1) + V(b_0)] + \beta [U(s_2) + V(b_1)] \} \\ \text{s.t.} \quad & Q(b_1)b_1 + s_1 = b_0 + g_1 \quad \text{and} \quad s_2 = b_1 + g_2. \end{aligned}$$

Let λ_1 and $\beta\lambda_2$ be the respective Lagrange multipliers. Then, the optimal debt issuance at $t = 1$ is

$$b^\circ = \arg \max_{b_1} \{ \lambda_1 Q(b_1)b_1 + \beta V(b_1) - \beta\lambda_2 b_1 \}. \quad (11)$$

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 $t = 2$.

Replace $Q(b) = \beta(1 + \pi(b))$ and let $\lambda_1 = \lambda_2 = \lambda$, which, as it will become clear later, helps proxy for the steady state of the infinite-horizon problem. Condition (11) then reduces to

$$b^\diamond = b^\diamond(\lambda) \equiv \operatorname{argmax}_b \Omega(b, \lambda), \quad \text{where} \quad \Omega(b, \lambda) \equiv V(b) + \lambda \pi(b)b. \quad (12)$$

The first term in Ω 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 shadow value of the total “rents” that the government extracts from the private sector for providing these services. This rent reminds of seigniorage in monetary models; it emerges because the financial friction depresses the interest rate on public debt relative to the underlying social discount rate; and it explains—indeed *defines*—the sense in which government borrowing can be “cheap” when interest rates are low.

By Assumption A2, $\pi(b)b > 0$, when $b \in (0, b_{\text{bliss}})$ and $\pi(b)b \leq 0$ otherwise. It follows that $b_{\text{seig}} \equiv \operatorname{argmax}_b \pi(b)b$ is necessarily both positive and strictly lower than the bliss point. In other words, one can think of the graph of $\pi(b)b$ as a “debt Laffer curve,” whose peak is attained at $b = b_{\text{seig}}$.

Had the government cared only about maximizing rents, it would have set $b = b_{\text{seig}}$. Had it cared only about easing the friction in the private sector, it would have set $b = b_{\text{bliss}} (\equiv \operatorname{argmax}_b V(b))$. For any $\lambda \in (0, \infty)$, the optimal debt issuance strikes a balance between these two goals, i.e., $b_{\text{seig}} < b^\diamond < b_{\text{bliss}}$. The stronger the need for fiscal space, as measured by λ , the closer b^\diamond is to b_{seig} .

Clearly, such fiscal considerations are present only when $\lambda > 0$, i.e., when taxation is distortionary. That is, distortionary taxation is *necessary* for satiation to be suboptimal—but it may not be *sufficient*. To see this, suppose that $\pi(b)$ was invariant to b_{bliss} , and zero thereafter. Then, b^\diamond would have equaled b_{bliss} , no matter λ . This underscores the following point:

Observation. *On the margin, the cost of more liquidity provision is not higher debt burden per se but rather higher interest rates: preserving the shortage of aggregate collateral, $b^\diamond < b_{\text{bliss}}$, is optimal because, and only because, it helps suppress the government’s cost of borrowing. The higher the need for fiscal space, as measured by λ , the higher the importance of this consideration, and the smaller b^\diamond .*

Consider now how b^\diamond depends on the severity of the financial friction. In Example 2, b^\diamond turns out to increase in ϵ , the amount of idiosyncratic risk (see Online Appendix C.4). At first glance, this property may seem intuitive: shouldn’t the government provide more liquidity/collateral when the demand for it is high? The answer is, however, subtler, because the size of idiosyncratic risk, or the severity of the financial friction, affects both the value of providing liquidity and the value of depressing interest rates. To see this, suppose that we scale both $V(b)$ and $\pi(b)$ by a scalar $\delta > 1$. Clearly, this scales $\Omega(b)$ by δ as well, but leaves b^\diamond unchanged. We therefore conclude:

Observation. *In general, a tighter financial friction can have an ambiguous effect on the government’s provision of liquidity/collateral (i.e., on optimal debt issuance), because it affects both $V(b)$ and $\pi(b)$.*

To sum up, the above analysis has highlighted the trade off between liquidity provision and interest-rate suppression. But it has abstracted from the interaction of these two objectives with that of tax smoothing, the distinction between the short run and the long run, and the optimal dynamics. We address these issues below.

4.2 Balancing the three objectives

Return to our infinite-horizon, continuous-time problem, let λ denote its costate, and consider its Hamiltonian: $H(s, b, \lambda) \equiv U(s) + V(b) + \lambda [s - R(b)b - g]$. This can be rewritten as

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

where $\Omega(b, \lambda)$ has the same meaning and definition as in our two-period heuristic above. But whereas there we treated λ as an exogenous constant, here we recognize that the path of λ is endogenous and is the mirror image of the optimal path of taxes. Indeed, since s ought to maximize the Hamiltonian, we have that $\lambda = -U'(s)$, that is, λ measures the tax distortion. Furthermore, the optimal path for λ must satisfy the planner's Euler condition, which can be written as

$$\dot{\lambda} = \Omega_b(b, \lambda). \tag{13}$$

In a steady state, this condition reduces to $\Omega_b = 0$, which suggests that a steady state of our problem is akin to the solution of the two-period problem used earlier. We verify and qualify this intuition at the end of this section. Away from a steady state, on the other hand, condition (13) equates Ω_b , the net effect of public debt on welfare and interest rates, with $\dot{\lambda}$, the change in the tax distortion. This underscores how the optimal policy balances the two objectives emphasized above—liquidity provision and interest-rate suppression—with the traditional objective of smoothing the tax distortion over time.

Intuitively, when $\Omega_b > 0$, there is value to increasing public debt, which requires raising taxes tomorrow relative to today. And the converse is true when $\Omega_b < 0$. 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 acts as an adjustment cost that slows down convergence to the steady state.

4.3 Full characterization

At this point, we have not required $V(b)$ and $\pi(b)b$ to be concave over the region $[0, b_{\text{bliss}})$. But even if we do so, concavity is a lost cause in the neighborhood of a satiation point: around it, $\pi(b)b$ turns from positive to zero, which is necessarily non-concave. As a result, there generally exist multiple paths for debt and taxes that satisfy the budget constraint, the Euler condition, and the transversality condition; each of these paths represents a local maximum; and *additional* arguments are necessary to identify the global maximum. Skiba (1978) provides such arguments. The general case is studied in Online Appendix D. Here, we sharpen the analysis by making the following simplifications:

Auxiliary Assumptions. [B0] For $b > b_{\text{bliss}}$, $V'(b) = \pi(b) = 0$.

[B1] For $b < b_{\text{bliss}}$, the ratio $V'(b)/\pi(b)$ is a constant ω .

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

[B3] The government's need for tax revenue, as parameterized by g , is sufficiently large.

As discussed in Section 4.5, these assumptions are *not* strictly needed for our main insights. Still, let us explain what they mean and why they represent a useful benchmark. B0 is motivated by our second example: it identifies b_{bliss} with the supply of collateral beyond which the financial friction ceases to bind and the only remaining distortion is taxation. B1 nests the case in which pecuniary externalities are absent and the private and social values of liquidity are equated (i.e., $\pi = V'$). B2 guarantees that the debt Laffer curve is single-peaked. Finally, B3 implies a sufficiently large shadow value for depressing the government's cost of borrowing. Together, these assumptions suffice for the following result:

Theorem 1. *There exist unique thresholds $(b_{\text{skiba}}, b^*, s^*)$, with $b_{\text{seig}} < b^* < b_{\text{bliss}} < b_{\text{skiba}}$ and $0 < s^* < \bar{s}$, such that the following properties hold:*

- (i) For $b_0 < b_{\text{skiba}}$, optimal debt and taxes converge monotonically to, respectively, b^* and s^* .
- (ii) For $b_0 \geq b_{\text{skiba}}$, optimal debt and taxes stay constant at their initial levels.

This result identifies b^* as the unique steady-state level of debt below satiation and $[0, b_{\text{skiba}})$ as its basin of attraction. The threshold b_{skiba} , in turn, is an example of the ‘‘Skiba points’’ (Dechert and Nishimura, 1981) that emerge in non-convex, dynamic-optimization problems. As the problem's state variable crosses such a point, the globally optimal plan switches from one local optimum to another. Here, this means that tax smoothing does not hold and the economy converges to b^* if $b_0 < b_{\text{skiba}}$, whereas tax smoothing holds and the economy stays put at its initial position if $b_0 > b_{\text{skiba}}$.

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

Consider the $\dot{b} = 0$ locus, which corresponds to balanced budget and is given by $\lambda = -U'(S(b))$ with $S(b) \equiv g + (\rho - \pi(b))b$. This locus is upward sloping because a higher level of debt requires a higher rate of taxation for the budget to be balanced.⁷

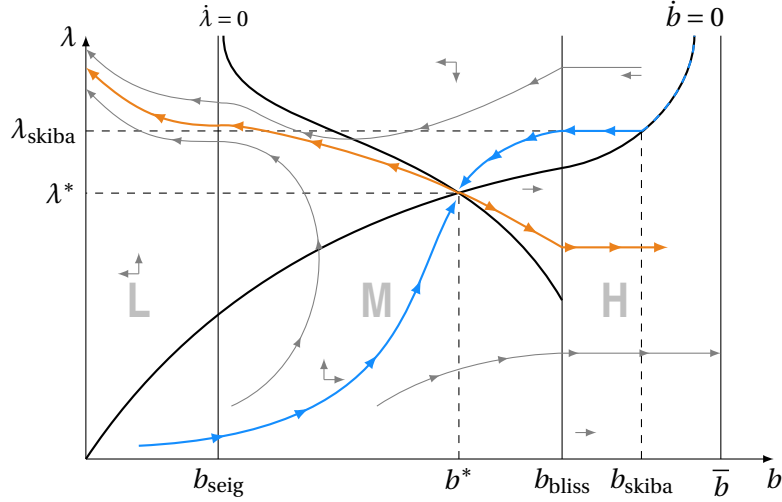
Next, consider the $\dot{\lambda} = 0$ locus, which captures tax smoothing. There are three scenarios to consider, corresponding to the regions L, M and H in the figure.

In region L, which corresponds to the upward-sloping portion of the debt Laffer curve ($b < b_{\text{seig}}$), increasing b raises both $V(b)$ and $\pi(b)b$. In this sense, there is a ‘‘free lunch’’ in this region. But the desire to smooth taxes means that this lunch has to be eaten slowly. Formally, the $\dot{\lambda} = 0$ locus does not exist,⁸ and instead $\dot{\lambda} > 0$ necessarily, reflecting the unambiguous optimality of issuing more debt.

⁷To be precise, this curve is upwards sloping if and only if the interest-rate bill, $R(b)b = (\rho - \pi(b))b$, is increasing in b . Although this may fail in region L, it does not affect the result, because it is optimal to leave region L and enter region M.

⁸Note that is true as long as non-negative lump-sum transfers are allowed, because this restricts $\lambda \geq 0$. Otherwise, a portion of the $\dot{\lambda} = 0$ locus exists in the negative territory of region L.

Figure 1: Phase Diagram and the Optimal Path.



In region M, which corresponds to the downward-sloping portion of the Laffer curve ($b_{\text{seign}} < b < b_{\text{bliss}}$), increasing b raises $V(b)$ at the expense of reducing $\pi(b)b$, so there is no more a free lunch. Instead, the trade off we have emphasized 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 tax revenue, 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 for λ , denoted by $\gamma(b)$ and defined by $\Omega_b(b, \gamma(b)) \equiv 0$, such that the following is true: $\dot{\lambda} = 0$ if $\lambda = \gamma(b)$; $\dot{\lambda} < 0$ (equivalently, taxes are falling) if $\lambda > \gamma(b)$; and $\dot{\lambda} > 0$ (equivalently, taxes are increasing) if $\lambda < \gamma(b)$.

Note that $\lambda = \gamma(b)$ is the inverse of the mapping $b = b^\diamond(\lambda)$ defined as in the previous subsection. As noted there, $b^\diamond(\lambda)$ decreases in λ because a higher λ favors interest-rate suppression (a b closer to b_{seign}) over liquidity provision (a b closer to b_{bliss}). It follows that $\gamma(b)$ is also decreasing, that is, the $\dot{\lambda} = 0$ locus is downward sloping. The $\dot{b} = 0$ locus, on the other hand, is upward sloping, reflecting the fact that higher debt maps to higher taxes under balanced budget. 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 .

Finally, consider region H, which corresponds to levels of debt above satiation. In this region, we have that $V(b)$ is flat and $\pi(b)b$ is zero, so $\dot{\lambda} = \Omega_b(b, \lambda) = 0$ for every λ . That is, the locus of $\dot{\lambda} = 0$ is now *all* of region H. Although this property may sound peculiar, it actually mirrors the textbook Ramsey paradigm (Barro, 1979; Lucas and Stokey, 1983). In that benchmark, both the liquidity-provision and the interest-rate concerns are absent, so $\dot{\lambda} = \Omega_b = 0$ over the entire phase diagram. Here, an analogous property holds in the portion of the phase diagram above satiation.

This also explains why, in region H, there exists a continuum of *apparently* optimal steady states, corresponding to the segment of the $\dot{b} = 0$ locus inside that region. All these steady states represent

local maxima: for any $b_0 > b_{\text{bliss}}$, the “Barro-like” plan that smooths the tax distortion and keeps $b(t)$ at b_0 for ever trivially satisfies both the Euler condition and the transversality condition. However, for $b_0 \in (b_{\text{skiba}}, b_{\text{bliss}})$, this plan is *not* the global maximum: it is dominated by an alternative plan, illustrated in Figure 1 by the segment of the saddle path that starts inside region H and enters in region M.

Debt falls gradually along this plan, crossing b_{bliss} in finite time and converging asymptotically to b^* . This plan requires a departure from tax smoothing (higher taxes early, lower taxes later), which is costly. But it allows the government to make a profit in terms of a positive liquidity premium, or equivalently a lower cost of borrowing, once debt falls below b_{bliss} . Provided that this happens fast enough, which is the case when $b_0 \in (b_{\text{bliss}}, b_{\text{skiba}})$, the sacrifice in terms of tax smoothing is justified and this plan dominates the Barro-like alternative. The converse is true if $b_0 > b_{\text{skiba}}$.

4.4 The optimal long-run quantity of public debt

Suppose that $b_0 < b_{\text{skiba}}$. This rules out the uninteresting scenario of a never binding financial friction..

Proposition 4. *The optimal long-run levels of debt and taxes solve the following fixed point:*

$$b^* = \arg \max_{b \in [0, b_{\text{bliss}}]} \{V(b) + \lambda^* \pi(b)b\} \quad (14)$$

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

This connects the steady state to the earlier, two-period exercise: b^* balances the desire to ease the financial friction with the value of extracting “seigniorage,” or suppressing the government’s cost of borrowing. A key difference, however, is that here the solution takes into account the fixed-point relation between the b^* that maximizes Ω and the Lagrange multiplier λ^* that appears inside Ω .

It is worth contrasting this result to the predictions of the textbook Ramsey paradigm (Barro, 1979; Lucas and Stokey, 1983). In the latter, the long-run level of debt is indeterminate, in the sense that it moves in tandem with the initial level of debt, by direct implication of the optimality of tax smoothing. Here, instead, the long-run level of debt is uniquely pinned down by the trade off between liquidity provision and interest-rate suppression.

Finally, it is useful to clarify how our result relates to Aiyagari and McGrattan (1998). Here, we characterize the optimal policy path starting from an arbitrary initial condition b_0 , establish that this path converges to a point b^* that is itself invariant to b_0 , and also characterize this point. By contrast, Aiyagari and McGrattan (1998) maximize welfare over steady-state policies, which in our context means the following: impose budget balance in all periods, and hence also $b_t = b_0$ for all t ; and then maximize welfare over b_0 . As explained in Online Appendix C.3, this over-estimates the opportunity cost of liquidity provision and, therefore, under-estimates the optimal long-run quantity of public debt.⁹

⁹Aiyagari and McGrattan (1998)’s strategy was probably dictated by the following complication: in their setting, unlike ours, the distribution of wealth is a non-redundant state variable for aggregate dynamics and hence for the planner’s problem. Here,

4.5 Relaxing Auxiliary Assumptions B0-B3

We conclude this section by relaxing Auxiliary Assumptions B0-B3. A detailed treatment can be found in Online Appendix D. Here, we summarize the main ideas.

Consider first B3. This requires that the value of fiscal space, as parameterized by g , be sufficiently high. If we relax this assumption, it becomes possible that (14) admits a corner solution at $b^* = b_{\text{bliss}}$. That is, satiation can be optimal in the long run, despite distortionary taxation. While logically possible, and also reminiscent of the literature on the Friedman rule, this scenario seems empirically implausible in our context.

Consider next B0. This lets us interpret b_{bliss} as a satiation point beyond which the private and social values of liquidity become zero. If we instead let π and V' turn negative for $b > b_{\text{bliss}}$, a scenario of “harmful excess,” the only change in Figure 1 is that Region H ceases to exist: convergence to a steady state below b_{bliss} is now guaranteed for every initial position and for any g . Clearly, this only reinforces our focus on the case of $b^* < b_{\text{bliss}}$.

Finally, consider B1 and B2. Without them, there can exist multiple steady states below b_{bliss} , each one with its own basin of attraction. Intuitively, the “adjustment cost” of a long-lasting departure from tax smoothing helps justify remaining at one steady state when another, seemingly superior, steady state exists but is sufficiently far away (see Online Appendix D for a detailed explanation). Nonetheless, the local dynamics around any such steady state remain the same as those around b^* in Figure 1. By the same token, our upcoming characterization of the optimal response to shocks remains the same as well (provided, of course, that shocks are small enough).

5 Optimal Response to Shocks

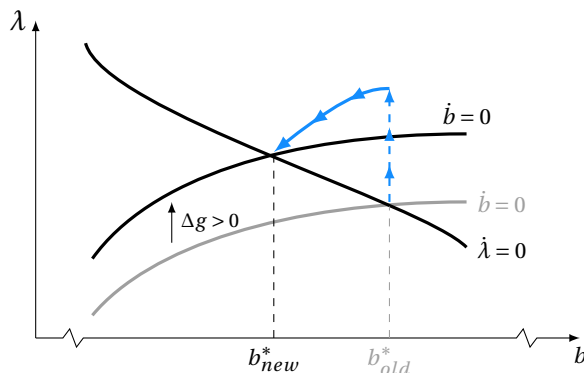
We now study how the tripartite trade off between liquidity provision, interest-rate suppression and tax smoothing shapes the optimal policy response to shocks. The main ideas are exposed by studying the comparative dynamics of the phase diagram. For further illustration, we also use the numerical, non-linear solution of a stochastic example.¹⁰

Government spending. Figure 2 considers an unexpected, once and for all, increase in g . Prior to the change, the economy rests at the steady-state point b_{old}^* . The increase in g causes the $\dot{b} = 0$ locus to shift upwards, reflecting the need for higher taxes. By contrast, the $\dot{\lambda} = 0$ locus does not move, because g does not enter π and V , and hence it does not enter the planner’s Euler condition either. As a result, we solve the correct policy problem but in a simpler model. Future research may attempt to solve the correct policy problem in the richer setting.

¹⁰Throughout this section, we let public debt be risk-free, as in Barro (1979) and Aiyagari et al. (2002). The case of fully state-contingent debt is considered in Online Appendix C.1. As in Lucas and Stokey (1983), this 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.

the steady-state level of debt drops from b_{old}^* to b_{new}^* and the optimal dynamic response is as follows: taxes initially increase by more than the increase in g , in order to allow debt and interest rates to fall; and because of this, taxes in the long run increase by less than the increase in g .

Figure 2: Permanent Increase in Government Spending



Proposition 5. *Suppose that government spending increases permanently by Δg . The optimal taxes increase by more than Δg in the short run and by less than Δg in the long run.*

Compare this result to Barro (1979) and Aiyagari et al. (2002), henceforth referred to as “Barro/AMSS.” 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.

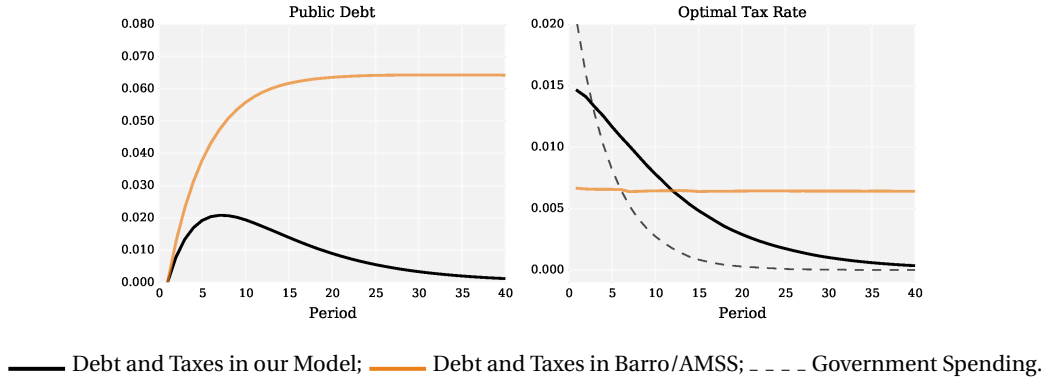
The same logic applies to “wars,” namely, temporary fiscal shocks. We illustrate this in Figure 3, using an example in which g follows a symmetric two-state Markov process. In the Barro/AMSS benchmark (orange lines), the war leaves a permanent mark on debt and taxes, reflecting the supremacy of tax smoothing. In our setting (black lines), instead, the economy eventually reverts to its initial position, reflecting the existence of a well-defined long-run target for debt. Finally, the accumulation of debt during the war is less pronounced than that in Barro/AMSS, because doing so allows the planner to moderate the increase in interest rates, which would have further tightened the budget.¹¹

Flight to safety. 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 $\pi(b) = \delta \tilde{\pi}(b)$ and $V(b) = \delta \tilde{V}(b)$, for fixed $\tilde{\pi}$ and \tilde{V} , and consider an increase in δ . We think of this situation as a “flight to safety.”

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. If the $\dot{b} = 0$ locus had also stayed the same, the

¹¹If the war is sufficiently persistent, this mechanism becomes so strong that the level of debt actually falls, as in the example with a permanent change discussed above.

Figure 3: Optimal Response to a War



optimal response would have been to stay put, in line with the discussion in Section 4.1. But the $\dot{b} = 0$ locus actually shifts down, because the shock reduces interest-rate costs, thus also reducing the taxes needed to balance the budget. In this sense, the private sector’s flight to safety is akin to a bonanza of government revenue.

Proposition 6. *A flight to safety, modeled as an equiproportional increase in the social and private value of liquidity, is equivalent to a positive shock in government revenue. This reduces the shadow cost of liquidity provision, thus justifying more debt issuance.*

Of course, reality is more complicated than the scenario just described. A financial shock may also shrink the tax basis, which can have a countervailing effect. Still, our insights provide a rationale for why financial shocks may justify larger deficits than other shocks, a point we expand on below. They also qualify the conventional intuition that an increase in the demand for liquidity calls for an increase in the government’s provision of it: this intuition fails to take into account how such a shock may also raise the marginal return to interest-rate suppression, which pulls in the opposite direction.

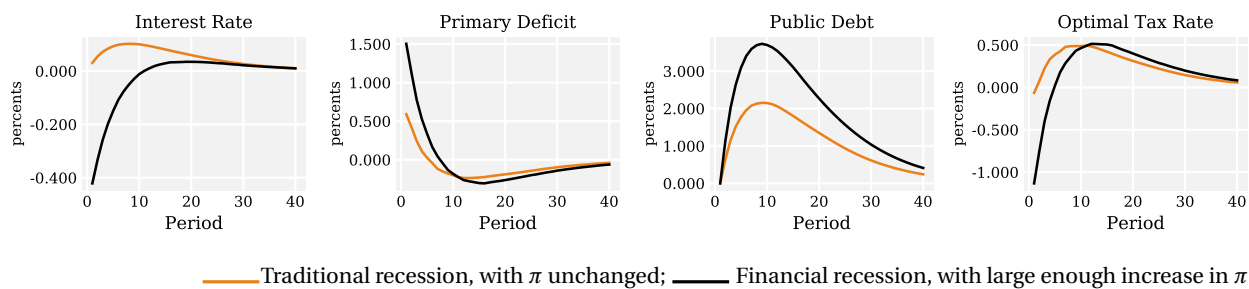
Traditional vs Financial Recessions. Let us represent a Keynesian/inefficient recession as an exogenous shock to the labor wedge. This naturally leads to lower aggregate output and tax revenue, and an increase in the deficit.¹² Next, let us distinguish between two flavors of such a recession: a “traditional” one, which leaves the functions π and V unaffected, and a “financial” recession, which shifts these functions up 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

¹²Formally, we modify the micro-founded examples of Section 3 by letting the equilibrium condition for labor be $v'(n_t) = (1 - \tau_t)(1 + \omega_t)$, where ω_t is an exogenous shock. This leaves V and π in our reduced form unaffected, but makes U a joint function of τ_t and ω_t . We then capture an (inefficient) recession as a temporary negative shock to ω_t .

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.” Clearly, the same is true for government spending if we endogenize g and let the recession raise its marginal value (proxing for the value of a fiscal stimulus in a New Keynesian economy during a liquidity trap).

Figure 4: IRFs to a Financial vs Traditional Recession



This provides a basis for the argument made by Paul Krugman, Brad DeLong and others that the drop in risk-free rates during a financial crisis 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 drop in the risk-free rate *per se*, but rather the extent to which this represents an increase in the wedge between it and the counterfactual rate that obtains in a frictionless world. Had r and ρ dropped together, leaving π the same, government borrowing would *not* be cheaper.

6 Discussion

In this section we return to the micro-foundations behind our reduced form so as to shed additional light on the following issues: the possibility that public debt crowds out (or, in) capital; the reason why public debt is non-neutral; and the precise liquidity premium that shapes optimal policy.

Crowding out (or in) capital. Revisit the first example in Section 3 and let the risk-free asset exist in two flavors: government bonds and physical capital. For simplicity, capital takes the form of a storage technology available to the young.¹³ If a household invests x units of the final good when young, she gets back $k = f(x)$ units of the final good when old, where f satisfies $f'(\cdot) > 0$, $f''(\cdot) < 0$, and $f(0) = 0$. This endogenizes the safe endowment k in the original model, without upsetting the equality of the private and social value of saving. Indeed, the property $\pi(b) = V'(b)$ continues to hold because there is no pecuniary externality, and so does Proposition 2, with the following change in V :

$$V(b) \equiv \max_k \{ \varphi \delta u(k - c + b) + (1 - \varphi) \delta u(k + c + b) - f^{-1}(k) - b \}.$$

¹³This simplifies the exposition by making sure that the returns to capital do not depend on future labor supply and thereby on future labor taxes. Also, we abstract from taxation of capital.

Furthermore, the equilibrium level of capital is given by the value of k that solves the above maximization problem. It is then immediate to see that public debt crowds out capital, similarly to Diamond (1965) and Aiyagari and McGrattan (1998). Summing up:

Proposition 7. *In the extension described above, public debt crowds out capital. Nonetheless, the reduced-form representation of the policy problem continues to hold, and so do all our paper's lessons.*

A similar result obtains if we add physical capital in our second example and let it serve as collateral in morning transactions. Together, these extensions illustrate how our insight can be robust to the possibility that public debt can crowd out other private sources of collateral." But note that such crowding out is *not* an additional, separate element of the costs and benefits of debt issuance, it is already subsumed in the trade off we have analyzed.

Finally, one can reverse the crowding-out property, again without affecting our results. We offer such an example in Online Appendix A. Compared to Aiyagari and McGrattan (1998) and the examples discussed so far, the key difference of that example is to let public debt ease a financial friction in production rather than in consumption. Intuitively, this lets aggregate TFP increase with "the aggregate supply of collateral," thus also letting the returns to investment increase with public debt issuance. The crowding-out property is thereby reversed, but the policy problem is not fundamentally changed: π and V remain "sufficient statistics" for the costs and benefits of public debt issuance.

Public debt as private collateral. In the micro-foundations that underly our reduced form, public debt is non-neutral only because private borrowing capacity is not reduced one to one with future tax obligations. To see this, modify Example 2 in Section 3 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 is perfectly offset by an equal reduction in pledgeable private income. The same point applies to, inter alia, Woodford (1990), Aiyagari and McGrattan (1998), Holmström and Tirole (1998), Brunnermeier, Merkel and Sannikov (2021) and Reis (2021): if financial constraints adjusted appropriately to future tax obligations, public debt would be neutral in all these papers, too. This clarifies why public debt issuance can ease the financial friction and can thereby have a causal effect on welfare and interest rates.

Public debt as money. Our framework brings to mind models with money in the utility function and the literature on the Friedman rule (e.g., Chari, Christiano and Kehoe, 1996). We share with this literature the idea that the government is the net supplier of a scarce "service," but depart from it in that we equate this supply with the sum of the government's liabilities instead of a particular component of it ("money"). This assumption is consistent not only with the aforementioned theoretical literature on public debt as buffer stock/collateral, but also with two key facts: the low interest rate on public debt (Blanchard, 2019); and the premium government bonds command over other similarly safe assets (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood and Vayanos, 2014).

In order to explain how this assumption matters for our results, revisit Example 2 from Section 3 and suppose that the government enacts a regulation that outlaws the use of private (“corporate”) bonds as collateral in morning transactions. Clearly, this eliminates the liquidity premium on private bonds ($r_{\text{priv}} = \rho$) while preserving the liquidity premium on government bonds ($r_{\text{gov}} = r = \rho - \pi < \rho$). But it also raises the possibility: that the government could borrow in the cheap, money-like asset (here, government bonds) and at the same time save in the high-return, non-money asset (here, corporate bonds) If such portfolio is unrestricted and costless, the government’s net borrowing is disentangled from its provision of liquidity: the interest rate the government must pay for any marginal *net* borrowing is simply ρ . The trade off we have emphasized then ceases to apply, tax smoothing reigns supreme, and the optimal policy is determined in exactly the same fashion as in Barro (1979). But as long as the aforementioned kind of arbitrage is sufficiently restricted or costly,¹⁴ our insights go through.

This discussion, which is made formal in Online Appendix C.2, highlights that the main, high-level difference between our analysis and the literature on the Friedman rule is the intertwining of liquidity provision and debt management. Another, subtler issue regards the conditions under which long-run “satiation” (i.e., $b_t \rightarrow b_{\text{bliss}}$) might be optimal in our framework notwithstanding this intertwining. This is not the case in our micro-founded examples and seems improbable more broadly, as it requires a corner solution in the trade off between liquidity provision and interest-rate suppression. But as illustrated in Online Appendix D, it cannot be ruled out for arbitrary V and π .¹⁵

Inside liquidity. In Example 2, all private agents were net buyers in the market for safe assets (i.e., they were net savers in the afternoon of each period), and the government was the only supplier. More generally, though, some private agents may also be on the supply side. For instance, this is the case in a variant of Example 2 that adds heterogeneity in the ability to pledge future income and issue private bonds in the afternoon of each period. In this variant, “outside liquidity” in the form of safe public debt may naturally crowd out “inside liquidity” in the form of safe private debt, similarly to, inter alia, Holmström and Tirole (1998), Stein (2012) and Azzimonti and Yared (2019). Nevertheless, this possibility, too, can be subsumed in our reduced form, leaving our main results unaffected.

Clearly, this echoes our discussion of the question of whether public debt crowds out physical capital. Perhaps more interestingly, the following complementary insight emerges about the regulation mentioned above, which taxes or outlaws the use of privately-issued assets as collateral: by crowding out inside liquidity, the regulation aggravates the financial friction, which in turn reduces private welfare but also suppresses the interest rate on government debt. This regulation, versions of which have indeed

¹⁴This could be not only for reasons outside the model (such a political constraints) but also for a reason inside the model: if households are subject to similar borrowing constraints in the afternoon as in the morning, there is a limit in how much private debt they can issue, which in turn translates to a limit on how much private assets the government itself can acquire.

¹⁵Understanding the mapping between the properties of V and π and the structure of the economy therefore seems a worthwhile exercise. An example is the recent paper by Bassetto and Cui (2021): in their setting, which has similar micro-foundations as Example 3 of our paper, satiation is ruled out by restrictions on capital/entrepreneurial taxation.

been used in practice, is therefore subject to a similar trade off as that highlighted in our analysis.

Which r or π ? Although we view the evidence in Krishnamurthy and Vissing-Jorgensen (2012) and Greenwood and Vayanos (2014) as supportive of the mechanisms we are after in this paper, it is important to recognize that the particular premium measured in these papers, namely the wedge between government bonds and similarly safe, high-grade, corporate bonds, does not map to π (nor to V') in our setting. In the examples of Section 3, private and government bonds were perfect substitutes, so that $r_{\text{priv}} = r_{\text{gov}} = r \equiv \rho - \pi$ and therefore $r_{\text{priv}} - r_{\text{gov}} = 0$. At the other extreme, in the variant mentioned above, where private bonds could not serve as collateral by law, we had $r_{\text{priv}} = \rho$ and $r_{\text{gov}} = \rho - \pi$, and therefore $r_{\text{priv}} - r_{\text{gov}} = \pi$. More generally, if we let private assets enter collateral constraint (10) with a discount relative to public debt, we can get $\pi \geq r_{\text{priv}} - r_{\text{gov}} \geq 0$. But no matter how large or small the spread $r_{\text{priv}} - r_{\text{gov}}$ might be, or how much public debt crowds out private safe assets, $\pi = \rho - r_{\text{gov}}$ remains the relevant “sufficient statistic” for the policy problem in all the cases.

A similar point applies to the difference between the expected return to capital, r_K , and the interest rate on government debt: if this difference reflects only a compensation for aggregate risk, it need not matter for optimal policy. This is precisely the case in the example in the beginning of this section where we let capital be risky: regardless of the size of this risk and the associated equity premium, $\pi = \rho - r_{\text{gov}}$ remains the relevant sufficient statistic for the policy problem.¹⁶

All in all, this discussion shifts the focus to the interpretation and measurement of ρ . This parameter represents more generally the risk-free rate that obtains in a counterfactual world without idiosyncratic risk and financial frictions. As such, ρ may reflect not only intertemporal preferences but also aggregate risk. For instance, one can read Barro (2021) as an example of how “disaster risk” can rationalize a low effective ρ within a frictionless, representative-agent model. As mentioned before, such a model leaves no space for “cheap” government borrowing: it restricts $r_{\text{gov}} = \rho$, no matter how low or volatile ρ may happen to be. But if the real world is ridden with financial frictions, the gap π between r_{gov} and ρ may be sizable. And while $r_{\text{priv}} - r_{\text{gov}}$ provides a lower bound, this lower bound is likely to be too conservative, for the reasons explained above.

7 Conclusion

We have studied optimal policy in a setting where public debt management helps not only smooth taxes, as in Barro (1979) and Lucas and Stokey (1983), but also regulate a financial friction, as in Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998). Issuing more public debt helps

¹⁶To illustrate, suppose that investing x in period t returns Ak in period $t + 1$, where $k = f(x)$ like before but now A is an i.i.d. random variable with finite support $\mathcal{A} \subset \mathbb{R}_+$ and probabilities $\{\lambda_A\}_{A \in \mathcal{A}}$. Then, the implicit equity premium is of course positive, but the entire analysis—including the property $\pi(b) = V'(b)$ and Propositions 2 and 7—goes through with the following inconsequential change in the definition of V : $V(b) \equiv \max_k \sum_{A \in \mathcal{A}} \lambda_A \{\varphi \delta u(Ak - \epsilon + b) + (1 - \varphi) \delta u(k + \epsilon + b) - f^{-1}(k) - b\}$.

ease the friction and improve welfare by raising the aggregate supply of safe assets that private agents can use as buffer stock or collateral. But the flip side of this benefit is a reduction in fiscal space: easing the friction depresses the fee that the government can charge for its provision of collateral/liquidity and, equivalently, it raises the government's cost of borrowing relative to the relevant social discount rate.

This trade off explains why the government may optimally limit its provision of liquidity, and may even enact regulations that restrict the use of other equally safe but privately-issued assets as collateral. It helps pin down a unique long-run target for debt and taxes. It necessitates a departure from tax smoothing in the short run, so as to help attain the aforementioned long-run targets. And it provides new insights on 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 present a greater opportunity for cheap borrowing.

These insights were delivered with the help of a reduced-form Ramsey problem, whose main difference from the textbook, complete-markets counterpart (Barro, 1979; Lucas and Stokey, 1983) was letting the quantity of government debt enter social welfare and the government's cost of borrowing, via some functions V and π , respectively. These functions were shown to be "sufficient statistics" for the benefits and costs of public debt in a few micro-founded examples. While highly stylized, these examples differed from one another in a number of interesting ways, such as: the precise function of public debt; a possible wedge between the relevant private and social values, due to pecuniary externalities operating via collateral constraints; the possibility that public debt crowds out, or perhaps crowds in, physical capital; the equity premium on capital; and the interest-rate differential between government bonds and comparable private assets.

Needless to say, we do not mean to suggest that one can always abstract from these issues or the relevant micro-foundations: our reduced form relied on strong, simplifying assumptions, which will not hold in richer, more realistic, applications. Still, by spelling out a number of examples where our reduced form and results remained *exactly* valid as we varied all the aforementioned "details," we hope to have accomplished three goals. First, to identify the relevant trade off and explain how it drives optimal policy. Second, to clarify when government borrowing is truly cheap and by what metric. And finally, to illustrate how our insights may transcend a variety of applications and micro-foundations. We hope that this prism will help guide future work on both the theoretical and the quantitative front.

References

- Aguiar, Mark, Manuel Amador, and Cristina Arellano.** (2021) “Micro risks and Pareto improving policies with low interest rates.” *NBER working paper 28996*.
- Aiyagari, S. Rao.** (1994) “Uninsured Idiosyncratic Risk and Aggregate Saving.” *The Quarterly Journal of Economics*, 109(3): 659–684.
- Aiyagari, S. Rao, Albert Marcet, Thomas Sargent, and Juha 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.
- Andolfatto, David, and Fernando M. Martin.** (2018) “Monetary policy and liquid government debt.” *Journal of Economic Dynamics and Control*, 89(C): 183–199.
- Azzimonti, Marina, and Pierre Yared.** (2019) “The Optimal Public and Private Provision of Safe Assets.” *Journal of Monetary Economics*, 102: 126–144.
- Barro, Robert J.** (1979) “On the Determination of the Public Debt.” *Journal of Political Economy*, 87(5): 940–971.
- Barro, Robert J.** (2021) “ r minus g .” *NBER Working Paper 28002*.
- Bassetto, Marco, and Wei Cui.** (2021) “A Ramsey theory of financial distortions.” *IFS Working Paper W21/05*.
- Bewley, Truman F.** (1980) “The Optimum Quantity of Money.” In *Models of Monetary Economics*, ed. J. Kareken and N. Wallace. Minneapolis, Minnesota: Federal Reserve Bank.
- Blanchard, Olivier.** (2019) “Public Debt and Low Interest Rates.” *American Economic Review*, 109(4): 1197–1229.
- Brock, William A., and W.D. Dechert.** (1983) “The Generalized Maximum Principle.” University of Wisconsin mimeo.
- Brunnermeier, Markus K., Sebastian Merkel, and Yuliy Sannikov.** (2021) “Debt As Safe Asset.” *Princeton University mimeo*.
- Canzoneri, Matthew B., and Behzad T. Diba.** (2005) “Interest rate rules and price determinacy: The role of transactions services of bonds.” *Journal of Monetary Economics*, 52(2): 329–343.

- Chari, V.V., Lawrence Christiano, and Patrick Kehoe.** (1996) “Optimality of the Friedman Rule in Economies with Distorting Taxes.” *Journal of Monetary Economics*, 37(2): 203–223.
- Correia, Isabel, and Pedro Teles.** (1999) “The Optimal Inflation Tax.” *Review of Economic Dynamics*, 2(2): 325–346.
- 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.
- Delong, Bradford.** (2011) “A Free Lunch for America.” *Project Syndicate*, (Sept. 29), <https://www.project-syndicate.org/commentary/a-free-lunch-for-america-2011-09>.
- Delong, Bradford.** (2015) “We Need Bigger Deficits Now.” *The Week*, (Jan. 11), <https://theweek.com/articles/493768/need-bigger-deficits>.
- Diamond, Peter A.** (1965) “National Debt in a Neoclassical Growth Model.” *The American Economic Review*, 55(5): 1126–1150.
- 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.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** (2012) “The Aggregate Demand for Treasury Debt.” *Journal of Political Economy*, 120(2): 233–267.
- Krugman, Paul.** (2016) “Time to Borrow.” *The New York Times*, (Aug. 8), <https://www.nytimes.com/2016/08/08/opinion/time-to-borrow.html>.
- Krugman, Paul.** (2020) “Learn to Stop Worrying and Love Debt.” *The New York Times*, (Dec. 03), <https://www.nytimes.com/2020/12/03/opinion/biden-republicans-debt.html>.
- Lagos, Ricardo, and Randall Wright.** (2005) “A Unified Framework for Monetary Theory and Policy Analysis.” *Journal of Political Economy*, 113(3): 463–484.
- 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.
- Mehrotra, Neil, and Dmitriy Sergeyev.** (2020) “Debt Sustainability in a Low Interest Rate World.” *mimeo*.
- Mian, Atif, Ludwig Straub, and Amir Sufi.** (2022) “A Goldilocks Theory of Fiscal Deficits.” *mimeo*.
- Reis, Ricardo.** (2021) “The constraint on public debt when $r < g$ but $g < m$.” *CEPR discussion paper 16664*.

Skiba, A. K. (1978) "Optimal Growth with a Convex-Concave Production Function." *Econometrica*, 46(3): 527–539.

Stein, J.C. (2012) "Monetary Policy as Financial Stability Regulation." *Quarterly Journal of Economics*, 127(1): 57–95.

Woodford, Michael. (1990) "Public Debt as Private Liquidity." *American Economic Review*, 80(2): 382–388.

Online Appendices

A An Example with Capital and Production Frictions

In this Online Appendix we develop Example 3 of Section 3. This translates the “morning trades” in Example 2 to reallocation of capital across entrepreneurs. Such reallocation is impeded by a financial friction, which in turn can be eased by an increase in the government’s provision of liquidity/collateral. This example therefore stylizes the production-side, liquidity-enhancing role of public debt found in, inter alia, Woodford (1990), Holmström and Tirole (1998), Bassetto and Cui (2021) and Brunnermeier, Merkel and Sannikov (2021).

Setup. 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, θ_{it} is an idiosyncratic productivity shock, and f 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 second example of Section 3, 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 .

Reduced form. Relative to Example 2 of Section 3, 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, as we illustrate momentarily, public debt can have an ambiguous effect on capital accumulation: it can crowd *out* capital for a similar reason as in Aiyagari and McGrattan (1998), but it

can also crowd *in* capital by raising production efficiency. Notwithstanding these novelties, the policy problem remains essentially the same. In particular, the following reduced form applies.

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:*

$$\max_{\{\tau_t, b_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t) \quad (15)$$

$$s.t. \quad Q(\tau_{t+1}, b_{t+1})b_{t+1} = b_t + g - S(\tau_t, b_t) \quad (16)$$

W, Q and S capture, respectively, the per-period welfare flow, the market price of public debt and the tax revenue. As we move from examples studied in the main text to the present model, the micro-foundations and functional forms of these objects change. In particular, τ_t and b_t jointly determined both the price of government debt and the tax revenue, because the tax distortion and the aggregate supply of collateral jointly determine all of the following: the allocation of capital and labor; aggregate productivity and income; the returns to capital; the shadow multiplier on the collateral constraint; and the tax base.¹⁷ Nevertheless, the strategy for obtaining the reduced-form representation is similar to that in Section 2. 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 regulates the bite of the financial friction, similarly to the main text.

Continuous-time analogue. Consider the following continuous-time policy problem, which is motivated by the above micro-foundations and generalizes the problem studied in the main text:

$$\max \int_0^{+\infty} e^{-\rho t} W(\tau, b) dt \quad (17)$$

$$s.t. \quad \dot{b} = [\rho - \pi(\tau, b)]b + g - S(\tau, b) \quad \forall t \quad (18)$$

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 $b < b_{\text{bliss}}(\tau)$, whereas $\pi(\tau, b) = W_b(b, \tau) = 0$ if $b \geq b_{\text{bliss}}(\tau)$; this allows for the possibility that the “satiation point” beyond which the friction ceases to bind may depend on the tax rate. Similarly, let $b_{\text{seign}}(\tau) \equiv \arg \max \{\pi(\tau, b)b + S(\tau, b)\}$; this is the analogue to the level of debt that maximized seigniorage in the models of Section 3, 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.

¹⁷Relatedly, before we could use s_t and τ_t interchangeably because there was a one-to-one mapping between them, but now s_t is a function of both τ_t and b_t , so it is easier to identify the control variable with τ_t .

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

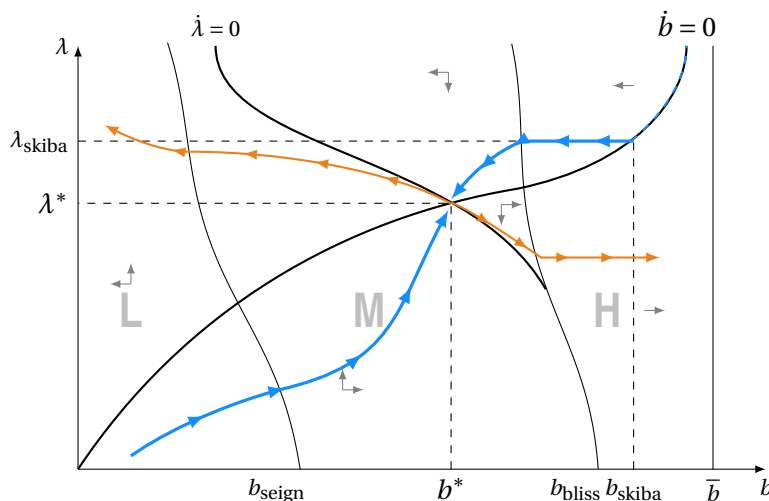
$$\dot{\lambda} = \Gamma(b, \lambda) \equiv \Omega_b(b, \lambda),$$

which has basically the same interpretation as its counterpart in Section 4.3. 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 $T(\lambda) = \operatorname{argmax}_{\tau} \{W(\tau, b) + \lambda[\pi(\tau, b)b + S(\tau, b)]\}$. We therefore obtain basically the same ODE system as in Section 3; some details are different but the essence remains the same.

Figure 5: Entrepreneurial Model

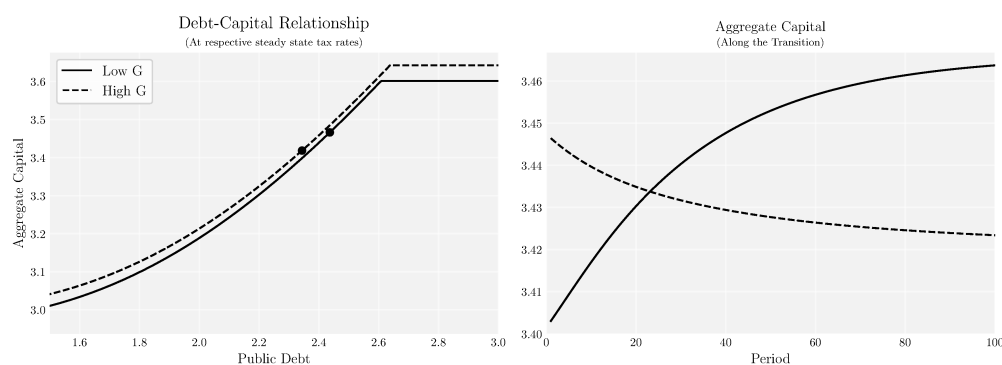


We illustrate these points in Figure 5. For this example, we let $[\pi(\tau, b)b + S(\tau, b)]$ be single-peaked in b . This guarantees that the phase diagram can be split in three regions, similar to regions L, M and H in Figure 1. The boundaries of these regions are now curved, rather than vertical, reflecting the fact that b_{seign} and b_{bliss} are allowed to vary with the rate of taxation, or equivalently with λ . The essence, however, remains the same: there is a unique steady state in which the financial friction does not bind, and the economy converges to it for all initial $b_0 < b_{\text{skiba}}$, for some b_{skiba} .

Crowding out or crowding in? We close this Online 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 two parameterizations of the micro-founded, discrete-time, model: one with relatively low level of government spending, and another with relatively high levels. For each parameterization, the left pannel of the figure considers the following exercise: we fix taxes at their optimal

Figure 6: Crowding in or out



steady-state levels, treat the quantity of public debt as a free variable, and report how the equilibrium amount of capital varies with it.¹⁸ This illustrates that, holding the tax distortion constant, public debt unambiguously crowds *in* capital. This is unlike Aiyagari and McGrattan (1998), because here public debt helps improve production efficiency and thereby raises the return to capital, which in turn encourages capital accumulation for given taxes.

Let us now zero in on the two dots in the left panel. These represent the *optimal* steady states for the two parameterizations. Note that the economy with higher g features a lower b^* for the reason discussed in the main text: the increased need for fiscal space encourages the government to squeeze liquidity out of the economy. It follows that this economy features a lower capital stock in steady state for two reasons: both because taxes are higher and because liquidity is (optimally) lower.

Let us finally turn to the right panel of Figure 6. This illustrates the aggregate capital dynamics along the *optimal* transition to b^* , starting from a $b_0 < b^*$. 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 pulls 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 is so strong that it dominates the crowding-out effect of taxes.

It's an open question which scenario is empirically relevant in the US context. But our analysis suggests that the answer to this question might not be as essential as previously thought: regardless of whether public debt crowds out or in capital, the main trade off remains that between easing the private sector's friction and easing the government's cost of borrowing. On the other hand, a recent paper by Bassetto and Cui (2021) suggests that this trade off may depend more critically on the interaction between rent-extraction and capital taxation. We leave a further exploration of this issue for future work.

¹⁸Both public debt and private capital are normalized by the steady-state level of output in the respective parameterization.

B Proofs

Proof of Proposition 1. The result follows directly from the derivations in the main text together with concavity and prudence ($u''(x) < 0 < u'''(x)$).

Q.E.D. ■

Proof of Proposition 2. The result follows directly from the main text.

Q.E.D. ■

Proof of Proposition 3. The household solves the problem

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [c_{i,t} + \theta_{i,t} \log(x_{i,t}) - v(h_{i,t})] \quad (19)$$

subject to

$$c_{i,t} + p_{i,t}x_{i,t} + q_t b_{i,t+1} = b_{i,t} + (1 - \tau_t)w_t h_{i,t} + p_t e_i \quad (20)$$

$$p_t(x_{i,t} - e_i) \leq \xi + b_{i,t} \quad (21)$$

where 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, H , in the population and, to simplify the exposition, set $\theta_L = 1$, $\theta_H = 1 + \epsilon$, $\epsilon > 0$, implying $\mathbb{E}[\theta] = 1 + \varphi\epsilon$ and $\mathbb{V}[\theta] = \varphi(1 - \varphi)\epsilon^2$. A greater ϵ is hence associated to a greater type dispersion. Endowments are set to $e_H = 0$, and $e_L = 1/(1 - \varphi)$, such that the total endowment, $\bar{e} = \varphi e_H + (1 - \varphi)e_L = 1$. The rest of the notation is identical to that used in Section 3.2.

In equilibrium, the borrowing constraint (21) 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:

$$\frac{1 + \epsilon}{x_{Ht}} = p_t(1 + \mu_t) \quad (22)$$

$$\frac{1}{x_{Lt}} = p_t \quad (23)$$

$$p_t x_{Ht} \leq \xi + b_t \quad (24)$$

$$\mu_t \geq 0 \quad (25)$$

$$\mu_t(\xi + b_t - p_t x_{Ht}) = 0 \quad (26)$$

$$\varphi x_{Ht} + (1 - \varphi)x_{Lt} = 1 \quad (27)$$

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

$$\pi_t \equiv q_t - \beta = \beta\varphi\mu_{t+1} \geq 0 \quad (28)$$

The first-best allocation, (x_H^*, x_L^*) , solves

$$\frac{1 + \epsilon}{x_{Ht}^*} = \frac{1}{x_{Lt}^*} \quad (29)$$

$$\varphi x_{Ht}^* + (1 - \varphi)x_{Lt}^* = 1 \quad (30)$$

and trivially satisfies $x_H^* > 1 > x_L^*$ and $\partial x_H^* / \partial \epsilon > 0 > \partial x_L^* / \partial \epsilon$. In particular, we get

$$x_H^* = \frac{1 + \epsilon}{1 + \varphi \epsilon} \quad \text{and} \quad x_L^* = \frac{1}{1 + \varphi \epsilon}.$$

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

$$p_t = \frac{1 + \epsilon}{x_{Ht}^*} \quad \text{and} \quad p_t x_H^* \leq \xi + b_t.$$

Let us define b_{bliss} such that the constraint exactly binds. b_{bliss} is unique and is given by $b_{\text{bliss}} \equiv 1 + \epsilon - \xi$. We immediately have that $b_t \geq b_{\text{bliss}}$ is sufficient for the borrowing constraint not to bind ($\mu_t = 0$) and the first best allocation to obtain. Conversely, when $b_t < b_{\text{bliss}}$, the first best allocation is unattainable. This establishes the first part of (ii) in Proposition 3.

When the borrowing constraint binds, still assuming $u(x) = \log x$, the equilibrium yields

$$x_{Ht} = \frac{b_t + \xi}{1 - \varphi + \varphi(b_t + \xi)}, \quad x_{Lt} = \frac{1}{1 - \varphi + \varphi(b_t + \xi)}, \quad \text{and} \quad \mu_t = \frac{1 + \epsilon - (b_t + \xi)}{b_t + \xi}.$$

Using the definition of b_{bliss} , this rewrites as

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

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 the premium π :

$$\pi(b) = \begin{cases} \beta \varphi \frac{(b_{\text{bliss}} - b)}{1 + \epsilon - (b_{\text{bliss}} - b)} & \text{for } b < b_{\text{bliss}} \\ 0 & \text{for } b \geq b_{\text{bliss}} \end{cases}$$

and the social value of debt $V(b) \equiv \beta [\varphi(1 + \epsilon) \log(x_H) + (1 - \varphi) \log(x_L)]$

$$V(b) = \begin{cases} \beta [\varphi(1 + \epsilon) \log(1 + \epsilon - (b_{\text{bliss}} - b)) - (1 + \varphi \epsilon) \log(1 + \varphi \epsilon - \varphi(b_{\text{bliss}} - b))] & \text{for } b < b_{\text{bliss}} \\ V_{\text{bliss}} \equiv \beta [\varphi(1 + \epsilon) \log(1 + \epsilon) - (1 + \varphi \epsilon) \log(1 + \varphi \epsilon)] & \text{for } b \geq b_{\text{bliss}} \end{cases}$$

We therefore reach the following result:

Lemma 1. *Suppose $\xi < 1 + \epsilon$, we have*

- (i) *There exists a unique threshold $b_{\text{bliss}} > 0$, given by $b_{\text{bliss}} = 1 + \epsilon - \xi$, such that the following properties hold for all $b < b_{\text{bliss}}$:*

$$\begin{aligned} \pi(b) &> 0, & \pi'(b) &< 0, & \pi''(b) &> 0, \\ V(b) &< V_{\text{bliss}}, & V'(b) &> 0, & V''(b) &< 0. \end{aligned}$$

Furthermore, for all $b \in (0, b_{\text{bliss}})$, we have $\pi(b) > V'(b)$. For $b \geq b_{\text{bliss}}$, on the other hand, $\pi(b) = 0$ and $V(b) = V_{\text{bliss}}$.

(ii) A tighter financial friction, or lower private collateral (lower ξ), increases b_{bliss} and uniformly raises both $V'(b)$ and $\pi(b)$ for all $b < b_{\text{bliss}}$.

(ii) Greater type dispersion (ϵ) increases b_{bliss} and uniformly raises both $V'(b)$ and $\pi(b)$ for all $b < b_{\text{bliss}}$.

Proof. The properties of π and V with respect to b , and hence (i), follow directly from their closed-form characterization. Part (ii) follows from the fact that $b_{\text{bliss}} = 1 + \epsilon - \xi$ along with the fact that, for any b_{bliss} and any $b < b_{\text{bliss}}$, π and V' are decreasing in b (hence increasing in b_{bliss}) and otherwise invariant to ξ . Part (iii) obtains directly by differentiating $\pi(b)$ and $V'(b)$ with respect to ϵ for any $b < b_{\text{bliss}}$. \square

Note that, for any $b < b_{\text{bliss}}$, the marginal social value of debt, $V'(b)$ is given by

$$V'(b) = \beta \left(\frac{\varphi(1 + \epsilon)}{1 + \epsilon - (b_{\text{bliss}} - b)} - \frac{\varphi(1 + \varphi\epsilon)}{1 + \varphi\epsilon - \varphi(b_{\text{bliss}} - b)} \right)$$

and can, after some algebraic manipulation, be rewritten as

$$V'(b) = \pi(b) + e(b) \text{ with } e(b) \equiv - \frac{\beta\varphi^2(b_{\text{bliss}} - b)}{1 + \varphi\epsilon - \varphi(b_{\text{bliss}} - b)}$$

where $e(b)$ captures the negative pecuniary externality at work in the model. 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 $\pi(b) > V'(b)$.

Q.E.D. ■

Proof of Theorem 1. Assumption B3 corresponds to the case $g > \hat{g}$ of Lemma 7, such that $\Psi_{\text{bliss}} > \gamma_{\text{bliss}}$. Then Proposition 12 applies and establishes part(i) of the theorem.

Q.E.D. ■

Proof of Proposition 4. Let us define $\Omega(b, \lambda^*) = V(b) + \lambda^* \pi(b)b$, where $\lambda^* = U'(s^*)$ and $s^* = g + r b^* - \pi(b^*)b^*$. 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 $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$, and it is then unique. Furthermore, the single-peakedness of $\pi(b)b$ guarantees that $\sigma(b) < 1$ and $\gamma(b) < 0$ for all $b < b_{\text{seign}}$, whereas $\sigma(b) > 1$ and $\gamma(b) > 0$ for all $b > b_{\text{seign}}$. Finally, the monotonicity of γ guarantees that $\gamma(b) > \gamma(b^*)$ for $b \in (b_{\text{seign}}, b^*)$, 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 $b \in [b^*, b_{\text{bliss}})$, which proves that $b^* = \arg\max_b \Omega(b, \lambda^*)$.

Q.E.D. ■

Proof of Propositions 5 to 7. Follow directly from the discussions in the 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

$$\mathbf{n}(x, p, w) = \int n(x, p, w; \theta) \varphi(\theta) d\theta \quad (31)$$

$$\mathbf{k}(x, p, w) = \int k(x, p, w; \theta) \varphi(\theta) d\theta \quad (32)$$

The problem of the household is

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

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:

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

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

$$v'(h_t) = (1 - \tau_t) w_t \quad (33)$$

$$q_t = \beta(1 + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1})) \quad (34)$$

$$1 = \beta(1 + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1}) p_{t+1}) \quad (35)$$

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

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

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

$$h_t = H(b_t, \tau_t) \text{ and } k_t = K(b_t, \tau_t) \quad (36)$$

so that

$$w_t = W(b_t, \tau_t) \text{ and } p_t = P(b_t, \tau_t) \quad (37)$$

Likewise, using the resource constraint, we get

$$c_t = \theta f(k_t, n_t) + (1 - \delta)k_t - k_{t+1} - g = \tilde{C}(b_t, \tau_t) - k_{t+1} \quad (38)$$

Using (36) and (38) 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 (34), we get

$$\begin{aligned} q_t &= Q(\tau_{t+1}, b_{t+1}) \\ \tau_t w_t h_t - g &= \tau_t W(b_t, \tau_t) H(b_t, \tau_t) - g = S(\tau_t, b_t) \end{aligned}$$

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

$$\begin{aligned} \max_{\{\tau_t, b_t\}_{t=0}^{\infty}} \quad & \sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t) \\ \text{s.t.} \quad & Q(\tau_{t+1}, b_{t+1}) b_{t+1} = b_t - S(\tau_t, b_t) \end{aligned}$$

Q.E.D. ■

C Additional Results

In this Online Appendix we discuss how our insights generalize in the presence of state-contingent debt, we elaborate on the connection to the literature on the Friedman rule, and finally we illustrate the determinants of the optimal debt issuance within Example 2.

C.1 Allowing for state-contingent debt

We now discuss how our analysis qualifies the insights of Lucas and Stokey (1983). Relative to Barro (1979) and AMSS (Aiyagari et al., 2002), 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.

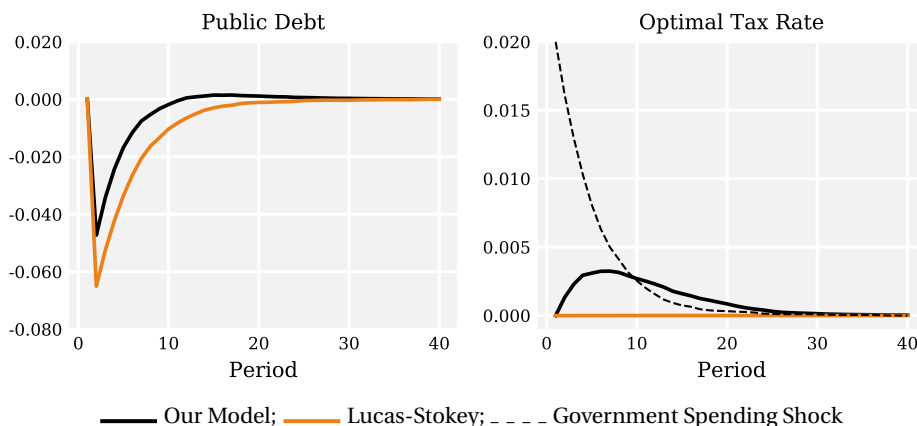
The answer differs in our setting. When state-contingent debt is available, our planner maintains the option to equate the shadow cost of taxation across different histories of shocks, exactly as in Lucas and Stokey (1983). But unlike that environment, the planner no longer finds it optimal to do so. 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 7 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.

Figure 7: Response to a war shock, with state-contingent debt



To sum up, once public debt is non-neutral for the reasons accommodated in this paper, the difference between Barro/AMSS and Lucas-Stokey is attenuated, and the qualitative responses of the optimal tax and the optimal allocation are the same whether debt is state-contingent or not.

C.2 On the Friedman Rule

Our paper departs from 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 (e.g., Krishnamurthy and Vissing-Jorgensen, 2012). To elaborate on the role played by it, revisit Example 2 in Section 3 and suppose that the government enacts a law that prohibits the use of corporate bonds as collateral in morning transactions. This means that, although both kinds of bonds can be used as stores of value, only government bonds convey money-like services. This in turn pegs the question of whether the government can not only issue the money-like asset (here, government bonds) but also save in the non-money asset (here, corporate bonds). If such saving is impossible, our analysis remains intact. At the other extreme, if such saving is not only possible but also completely unrestricted, the government’s net borrowing is completely disentangled from liquidity provision and, by the same token, the interest rate the government must pay for any additional borrowing is ρ . The trade off we have emphasized therefore ceases to apply, and optimal policy is determined in exactly the same fashion as in Barro (1979).

To see this more clearly, write the government’s net liabilities as $b = m - n$, 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

$$\dot{b} = \rho b - \pi(m)m + g - s, \quad (39)$$

where $\pi(m)m$ are the rents from issuing the money-like asset and s is tax revenue. The following properties are then evident: government borrowing, b , is disconnected from liquidity provision, m ; and the cost of borrowing is ρ , not $\rho - \pi$. Therefore, when the government varies b , it does not any more face the trade off we emphasized in our paper. 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 (39) over time to obtain the familiar intertemporal budget constraint:

$$b_0 + G = \int_0^{+\infty} e^{-\rho t} [\pi(m)m + S(\tau)] dt. \quad (40)$$

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 (40). Let λ denote the Lagrange multiplier on the intertemporal budget. It is then immediate that the optimal supply of liquidity is given by

$$m^* = \arg \max_m \Omega(m, \lambda), \quad (41)$$

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, debt management is disentangled from *both* liquidity provision and rent extraction; the trade off emphasized in our paper ceases to apply; tax smoothing reigns supreme; and the optimal path for τ and b is determined in exactly the same fashion as in Barro (1979).

Finally, as mentioned in the main text, another difference between our paper and the literature on the Friedman rule is the answer to the question of whether satiation obtains in the long run. That literature has traditionally favored a positive answer to this question, but in general the answer is ambiguous: it depends on assumptions about preferences, technologies, and tax instruments (see Chari, Christiano and Kehoe, 1996; Correia and Teles, 1999, and references therein). As already explained, a similar ambiguity is present in our framework for arbitrary π and V , although not for those obtained in the examples of Section 3. How this translates in richer, more realistic, settings is an open question. But we have hard time thinking of satiation as the empirically relevant scenario for the context of this paper.

C.3 Relation to Aiyagari and McGrattan (1998)

In this Online 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,

$$\hat{b} = \arg \max_b \{V(b) - \hat{\lambda} R(b)b\}, \quad (42)$$

which illustrates how Aiyagari and McGrattan (1998) treat the entire interest payments on public debt, $R(b)b$, as a cost. But as first highlighted in Section 4.4, the component ρb of these interest-rate payments is *not* a cost. That is, the *truly* optimal steady state satisfies

$$b^* = \arg \max_b \{V(b) + \lambda^* \pi(b)b\}, \quad (43)$$

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 followed in 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) accommodate two important effects that our main analysis abstracts from: the first is the role of inequality; the second is the possibility that public debt may crowd out capital accumulation by offering a substitute form of buffer stock. Although our modeling choices prevent us from studying the first issue, this limitation does not extend to the second issue: in Section 6 we discuss how our insights are robust to the possibility that public debt crowds out private capital, as well as to the exact opposite scenario.

Finally, note that discussion here presumes, like the analysis in Section 4.3, that the Auxiliary Assumptions B0-B3 hold. As explained in Online Appendix D.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 an essentially static optimization problem.

C.4 Optimal debt issuance in Example 2

Consider the problem introduced in Section 4.1, which as shown in Section 4.4 also informs the optimal steady state. Recall that this problem zeroed in on the trade off between liquidity provision and interest-rate suppression (or rent extraction). We now illustrate how exactly this trade off works in Example 2 from Section 3.

We start by verifying that the debt Laffer curve is single-peaked, the function Ω is also single-peaked, and the former's peak is below the latter's one.

Lemma 2. *Consider Example 2 from Section 3. The functions $\pi(b)b$ and $\Omega(b, \lambda)$ are strictly concave in $b \in [0, b_{\text{bliss}}]$ and their respective maxima, $b_{\text{seign}} \equiv \operatorname{argmax}_b \pi(b)b$ and $b^\diamond \equiv \operatorname{argmax}_b \Omega(b, \lambda)$, satisfy $0 < b_{\text{seign}} < b^\diamond < b_{\text{bliss}}$.*

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

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

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

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

which establishes the concavity of $\Pi(b) \equiv \pi(b)b$. Next, note that $\Pi'(0) = \beta\varphi \frac{1+\epsilon-\xi}{\xi} > 0$ and $\Pi'(b_{\text{bliss}}) = \pi'(b_{\text{bliss}})b_{\text{bliss}} = -\beta\varphi \frac{1+\epsilon-\xi}{1+\epsilon} < 0$. It follows that b_{seign} is the unique solution to $\Pi'(b) = 0$ 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^\diamond is the unique solution to $\Omega_b(b, \lambda) = 0$. Furthermore, because $g'(b_{\text{seign}}) = 0$, $g'(b_{\text{bliss}}) < 0$, $V'(b_{\text{seign}}) > 0$, and $V'(b_{\text{bliss}}) = 0$, we have that $\Omega_b(b, \lambda) > 0$ at $b = b_{\text{seign}}$ and $\Omega_b(b, \lambda) < 0$ at $b = b_{\text{bliss}}$, and therefore that b^\diamond is strictly between b_{seign} and b_{bliss} . \square

This verifies, in effect, that Example 2 verifies all the Auxiliary Assumptions invoked in Section 4. What is more, because we now have a simple closed-form characterization of Ω , we can go a step further to study the comparative statics of b^\diamond with respect to the underlying primitives. Those are reported in the next proposition

Proposition 9. *Consider Example 2, fix λ , and let $b^\diamond = \operatorname{argmax}_b \Omega(b, \lambda)$. The following comparative statics are true: b^* increases with the size of the liquidity risk (ϵ), decreases with the value of fiscal space (λ), and is generally non-monotonic in the amount of inside collateral (ξ).*

Proof. Using our closed-form solution for π and V along with the fact $b_{\text{bliss}} = 1 + \epsilon - \xi$, we can show that

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

Furthermore,

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

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

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

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

$$\frac{\partial^2 \Omega}{\partial b \partial \xi} = \frac{\partial^2 V}{\partial b \partial \xi} + \lambda \frac{\partial^2 \Pi}{\partial b \partial \xi}, \quad \frac{\partial^2 V}{\partial b \partial \xi} = V''(b) < 0, \quad \text{and} \quad \frac{\partial^2 \Pi(b)}{\partial b \partial \xi} = \frac{\beta \varphi \theta (b - \xi)}{(\beta + \xi)^2}.$$

Because $\frac{\partial^2 \Pi(b)}{\partial b \partial \xi}$ sign with the position of b relative to ξ , the effect of ξ on b^\diamond is generally ambiguous. In particular, we have found numerically that b^\diamond is inversely U-shaped with respect to ξ . \square

Although b^\diamond 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, $b^\diamond + \xi$.¹⁹ It then also follows that, at the optimal quantity of public debt, more private collateral depresses the liquidity premium ($\frac{\partial \pi(b^\diamond)}{\partial \xi} < 0$), whereas the converse is true with an aggravation of liquidity needs ($\frac{\partial \pi(b^\diamond)}{\partial \epsilon} > 0$).

To conclude, these findings complement the intuitions developed in Sections 4.4 and 5. 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 .

¹⁹To see this, let $z \equiv b + \xi$ and re-express V , π , and Ω as functions of z instead of b :

$$\pi(b) = \bar{\pi}(z) \equiv \beta \varphi \frac{1 + \epsilon - z}{z}$$

$$V(b) = \bar{V}(z) \equiv \beta \{ \varphi(1 + \epsilon) \log(z) - (1 + \varphi \epsilon) \log(\varphi z + 1 - \varphi) \}$$

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

Because \bar{V} and $\bar{\pi}$ are invariant to ξ , it is immediate that $\frac{\partial^2 \bar{\Omega}}{\partial z \partial \xi} = -\lambda \bar{\pi}'(z) > 0$, which via the IFT implies that $z^\diamond \equiv \arg \max_z \bar{\Omega}(z, \lambda) = b^\diamond + \xi$ 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^\diamond increases with ξ .

D Characterization of Optimal Plan

In this Online Appendix we offer a complete, self-contained, characterization of the solution to problem (1)-(2). 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 finally explain how the results generalize away from that benchmark, that is, when we relax Auxiliary Assumptions B0-B3

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

D.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 Online 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.²⁰

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:

$$\dot{b} = \Psi(b, \lambda) \equiv g + (\rho - \pi(b))b - s(\lambda), \quad (44)$$

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 in λ : a higher λ means higher taxes today, which in turn means lower debt tomorrow.²¹ By the Implicit

²⁰Had 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.

²¹Note 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.

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 3. *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,

$$\dot{\lambda} = v(b) - \lambda\pi(b)(\sigma(b) - 1). \quad (45)$$

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 (45) 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 (45) is non-zero, optimality requires a non-zero drift in λ , that is, a deviation from tax smoothing.

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

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

We can then restate the Euler condition (45) as follows:

$$\dot{\lambda} = \begin{cases} v(b) \left[1 - \frac{\lambda}{\gamma(b)} \right] & \text{if } b \in \Delta \\ 0 & \text{if } b \notin \Delta \end{cases} \quad (46)$$

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 $b < b_{\text{bliss}}$). 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

$$\frac{d[\pi(b)b]}{db} = \pi'(b)b + \pi(b) = -(\sigma(b) - 1)\pi(b). \quad (47)$$

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 $[0, b_{\text{bliss}}]$, 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.

D.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 assumptions

- B0.** For $b > b_{\text{bliss}}$, $V'(b) = \pi(b) = 0$.
- B1.** the ratio v/π is constant;
- B2.** the elasticity σ is increasing in $b \in (0, b_{\text{bliss}})$.

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 $\pi(b_{\text{seign}}) + \pi'(b_{\text{seign}})b_{\text{seign}} = \pi(b_{\text{seign}})(1 - \sigma(b_{\text{seign}})) = 0$. An implication of **B2** is then that $\sigma(b) < 1$ for $b < b_{\text{seign}}$ and $\sigma(b) > 1$ for $b > b_{\text{seign}}$. Dispensing from **B3** will allow use to obtain a more general characterization of the cases implied by **B0–B2**.

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

Proposition 10. *Let Assumptions **B0–B2** hold. There exists a unique $b^* \in (\underline{b}, b_{\text{bliss}}]$ such that, for any initial point $b_0 < b_{\text{bliss}}$, the optimal level of public debt converges monotonically to b^* . Furthermore, $b^* < b_{\text{bliss}}$ if $g > \hat{g}$ and $b^* = b_{\text{bliss}}$ 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 $b_0 < b_{\text{bliss}}$. 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 10 in multiple steps, developing additional insights on the way. We start by noting that Property **B0–B2** imply the following structure for the function γ , which is instrumental for the subsequent analysis.

Lemma 4. *Let Assumptions **B0–B2** hold. The domain of γ is $\Delta = [\underline{b}, b_{\text{seign}}) \cup (b_{\text{seign}}, b_{\text{bliss}})$, where $b_{\text{seign}} \equiv \arg \max_b \pi(b)b$. For $b \in [\underline{b}, b_{\text{seign}})$, γ is negatively valued and decreasing. For $b \in (b_{\text{seign}}, b_{\text{bliss}})$, γ is positively valued and decreasing. Finally, $\gamma(b) \rightarrow -\infty$ as $b \rightarrow b_{\text{seign}}$ from below and $\gamma(b) \rightarrow +\infty$ as $b \rightarrow b_{\text{seign}}$ from above.*

Proof. Recall that $b_{\text{seign}} = \arg \max_b \pi(b)b$, 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 $b \leq b_{\text{seign}}$. From the definition of γ and the assumption $V'(b) \propto \pi(b)$, we have

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

The latter result together with the definition of b_{seign} implies that $\lim_{b \uparrow b_{\text{seign}}} \gamma(b) = -\infty$ and $\lim_{b \downarrow b_{\text{seign}}} \gamma(b) = \infty$. 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 $[\underline{b}, b_{\text{bliss}})$. \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 3, ψ is positively valued and strictly increasing. Together with Lemma 4, this means that γ and ψ can intersect at most once. In particular, letting $\gamma_{\text{bliss}} \equiv \lim_{b \uparrow b_{\text{bliss}}} \gamma(b)$ and $\psi_{\text{bliss}} \equiv \psi(b_{\text{bliss}})$,²² we have the following property.

Lemma 5. *Let Assumptions **B0–B2** hold. If $\gamma_{\text{bliss}} > \psi_{\text{bliss}}$, then γ and ψ never intersect. If instead $\gamma_{\text{bliss}} < \psi_{\text{bliss}}$, then γ and ψ intersect exactly once, and this intersection occurs at $b = b^*$, for some $b^* \in (b_{\text{seign}}, b_{\text{bliss}})$.*

Proof. From Lemma 4, we know that $\psi(b)$ and $\gamma(b)$ can only intersect in $(b_{\text{seign}}, b_{\text{bliss}})$. Given (i) the monotonicity of $\sigma(b)$ and hence $\gamma(b)$, (ii) the fact that $\psi(b)$ is increasing and (iii) $\lim_{b \downarrow b_{\text{seign}}} \gamma(b) = \infty$,

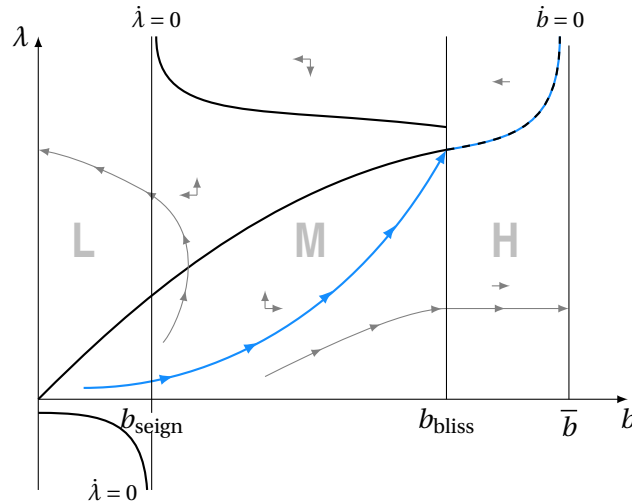
²²Recall that γ is defined to the left of the satiation point but not *at* it, which explains why we write $\gamma_{\text{bliss}} \equiv \lim_{b \uparrow b_{\text{bliss}}} \gamma(b)$ rather than $\gamma_{\text{bliss}} \equiv \gamma(b_{\text{bliss}})$. Also, the existence of the limit follows from the 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.

$\gamma(b)$ and $\psi(b)$ can intersect at most once. If $\gamma_{\text{bliss}} > \psi_{\text{bliss}}$, (i) and (iii) imply that $\gamma(b)$ lies above $\psi(b)$ everywhere in $(b_{\text{seign}}, b_{\text{bliss}}]$ and therefore they never intersect. In $\gamma_{\text{bliss}} < \psi_{\text{bliss}}$, (i)–(iii) imply that they intersect only once. \square

The two scenarios are illustrated in, respectively, Figures 8 and 9. The latter is the same as Figure 1 in the main text, reproduced here to ease the exposition.

Let us first consider Figure 8. The phase diagram is split in three regions: the region L, for $b < b_{\text{seign}}$; the region M, for $b \in (b_{\text{seign}}, b_{\text{bliss}})$; and the region H, for $b > b_{\text{bliss}}$. 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 9. What distinguishes the two figures is whether γ and ψ admit an intersection within region M. In Figure 8, they do not. This is because we have imposed $\gamma_{\text{bliss}} > \psi_{\text{bliss}}$, which together with the monotonicity of γ and ψ guarantees that γ lies above ψ throughout region M.

Figure 8: Benchmark, with $\psi_{\text{bliss}} < \gamma_{\text{bliss}}$.



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 8, the ODE system (44)-(46) admits no steady state to the left of the satiation point b_{bliss} (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_{bliss} 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 $b = b_{\text{bliss}}$ and $\lambda = \psi_{\text{bliss}} \equiv \psi(b_{\text{bliss}})$; the latter corresponds

to the level of taxes that balances the budget when the economy rests at the satiation point.

For any $b_0 < b_{\text{bliss}}$, 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 (44)-(46), then, and only then, the economy converges asymptotically to $(b_{\text{bliss}}, \lambda_{\text{bliss}})$. 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.²³

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 $b_0 < b_{\text{bliss}}$, the path that leads to satiation is the optimal path, and Proposition 10 applies with $b^* = b_{\text{bliss}}$. For any $b_0 \geq b_{\text{bliss}}$, 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 11. *Let Assumptions B0–B2 hold and suppose $\psi_{\text{bliss}} < \gamma_{\text{bliss}}$. If $b_0 < b_{\text{bliss}}$, debt converges monotonically to b_{bliss} and taxes exhibit a positive drift along the transition. If instead $b_0 \geq b_{\text{bliss}}$, debt stays constant at b_0 for ever, and tax smoothing applies.*

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

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

²³One 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 $\lambda < 0$.

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

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

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

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

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

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

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 $(b_{\text{bliss}}, \lambda_{\text{bliss}})$ therefore satisfy a saddle path property. It is also easy to show that the eigenvector associated to the stable eigenvalue is given by

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

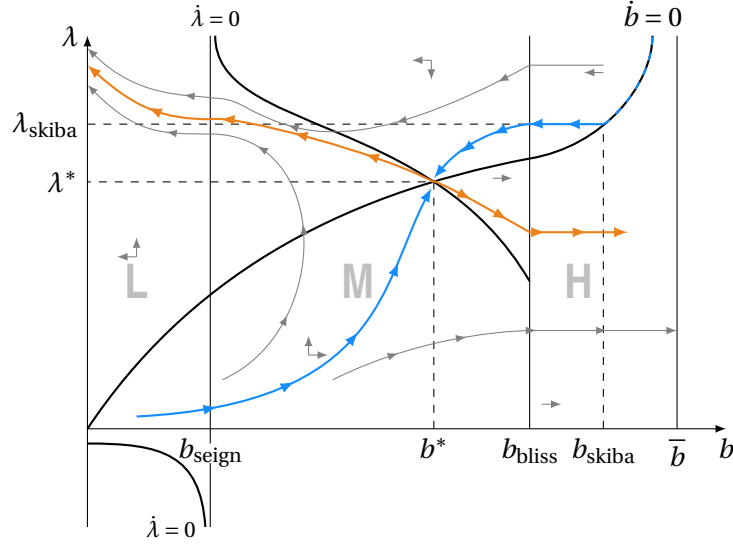
and is not degenerate as $\psi'(b) > 0$. In other words, starting from $b(0) \in \{b_{\text{bliss}} - \varepsilon; \varepsilon > 0\}$, there exists a unique path taking the economy to satiation. This establishes the first part of the proposition.

Let us now consider a situation where $\gamma(b_{\text{bliss}}) < \psi(b_{\text{bliss}})$. 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. $(b_{\text{bliss}}, \lambda_{\text{bliss}})$ is not locally stable and starting from $b < b_{\text{bliss}}$, there exists no path leading the economy towards it. \square

Let us now consider Figure 9. In this case, γ and ψ intersect exactly once, at $b = b^* \in (b_{\text{seign}}, b_{\text{bliss}})$. 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 (44)-(46) 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 $b_0 < b_{\text{bliss}}$, we can find a continuous path that satisfies conditions (44)-(46) and that asymptotically converges to (b^*, λ^*) . Exactly the same arguments as in Figure 8 guarantee that this path is the unique candidate for optimality, and hence also the optimal path, as long as $b_0 < b_{\text{bliss}}$.

A crucial difference from the case in Figure 8 is that the economy now converges to a steady state characterized by a debt level that is strictly lower than the satiation level: Proposition 10 applies with $b^* < b_{\text{bliss}}$. 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 $b_0 \in (b^*, b_{\text{bliss}})$.

Figure 9: Benchmark, with $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$.



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 $b_0 > b_{\text{bliss}}$. This is no longer true. Instead, as it is evident in the figure, for any $b_0 \in [b_{\text{bliss}}, b_{\text{skiba}})$, 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_{bliss} 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_{bliss} , 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_{bliss} 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 (44)-(46) 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 6 (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 9. Pick any $b_0 \geq b_{\text{bliss}}$ 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 $b_0 = b_{\text{bliss}}$. If we initiate the ODE system with a starting value $\lambda(0)$ slightly above $\psi_{\text{bliss}} = \psi(b_{\text{bliss}})$, 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 $\psi(b_{\text{bliss}})$, 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 $\lambda_{\text{skiba}} \in (\psi_{\text{bliss}}, \infty)$ such that, if we start with $\lambda(0) = \lambda_{\text{skiba}}$, 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_{bliss} . 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_{bliss} . 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 $b_{\text{skiba}} \in (b_{\text{bliss}}, \bar{b})$ as the level of debt that balances the budget when taxes are set at the level corresponding to λ_{skiba} ; that is, $b_{\text{skiba}} \equiv \psi^{-1}(\lambda_{\text{skiba}})$. Note that ψ is continuous and monotone, $\lambda_{\text{bliss}} > \psi(b_{\text{bliss}})$, and $\lim_{b \rightarrow \bar{b}} \psi(b) = \infty$; this verifies that b_{skiba} exists and is necessarily strictly between b_{bliss} 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 $b_0 < b_{\text{skiba}}$, as illustrated in the figure.

We thus have the following complement to Proposition 11.

Proposition 12. *Let Assumptions B0–B2 hold and suppose $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$. There exist unique points $b^* \in (b_{\text{seign}}, b_{\text{bliss}})$ and $b_{\text{skiba}} \in (b_{\text{bliss}}, \bar{b})$ such as the optimal debt level converges monotonically to b^* if $b_0 < b_{\text{skiba}}$, whereas it stays constant at b_0 for ever if $b_0 \geq b_{\text{skiba}}$. Optimal taxes exhibit a positive drift as long as $b \in (b_{\text{seign}}, b^*)$, a negative drift as long as $b \in (b^*, b_{\text{bliss}})$, and are smoothed as long as $b > b_{\text{bliss}}$.*

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

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

where $\omega \equiv \pi(b)/V'(b)$ 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 $b^* < b_{\text{bliss}}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ and $V'(b^*) > 0$. Finally, from Lemma 2, we know that $\gamma'(b) < 0$ for $b \in (b_{\text{seign}}, b_{\text{bliss}}]$. 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 $b_0 < b_{\text{bliss}}$, so that the financial distortion is present in the initial period. Under this restriction, the combination of Propositions 11 and 12 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 $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$ 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 $\psi_{\text{bliss}} \equiv \psi(b_{\text{bliss}})$ and $\gamma_{\text{bliss}} \equiv \lim_{b \uparrow b_{\text{bliss}}} \gamma(b)$. It follows that $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$ if and only if $\Omega_b(b, \lambda) < 0$ for (b, λ) close enough to $(b_{\text{bliss}}, \psi(b_{\text{bliss}}))$, which leads to the following simple interpretation.

Fact 1. $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$ 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 10 is then completed by noting that $\psi_{\text{bliss}} > \gamma_{\text{bliss}}$ if and only if g is high enough, a property that holds even outside our benchmark and that is proved in Lemma 7 below.

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

D.3 Beyond the Benchmark: Relaxing Auxiliary Assumptions B1-B3

Thanks to Auxiliary Assumptions B1-B3, 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 11 and 12 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.²⁴ 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 10, 11 and 12 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 10. In order to simplify the exposition, we truncate region L, where $b < b_{\text{seign}}$, γ is negatively valued, and there can be no steady state; we thus focus on region M, where $b \in (b_{\text{seign}}, b_{\text{bliss}})$ 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 13. *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 $\gamma'(b^*) > \psi'(b^*) + \chi$.*

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

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

²⁴Recall 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 $b \downarrow b_{\text{bliss}}$. 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 $(b_{\text{seign}}, b_{\text{bliss}})$.

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 $b^* < b_{\text{bliss}}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ 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

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

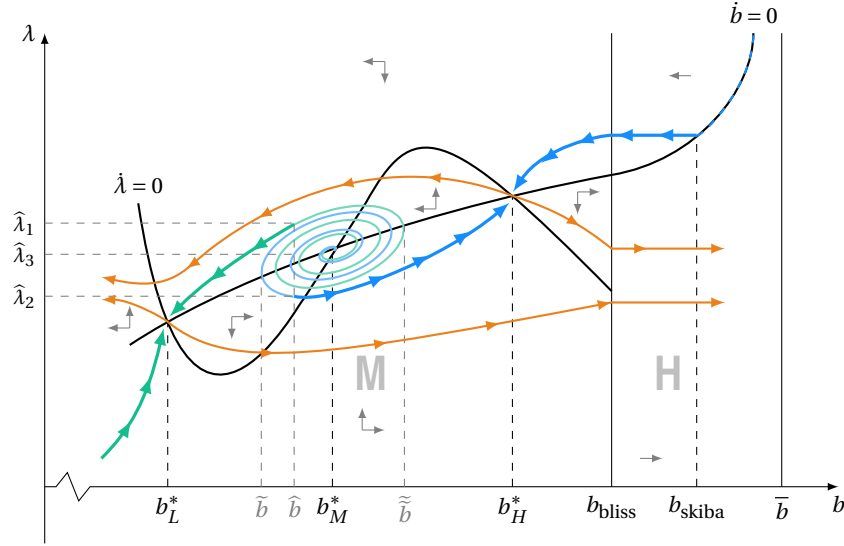
Therefore establishing the condition for the emergence of complex vs real explosive eigenvalues. \square

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 10 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 10 considers one of these possibilities.

In Figure 10, 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 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.

Figure 10: Beyond the Benchmark: Rich Dynamics and Multiple Steady States



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 6 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 6; 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 $b_0 = \tilde{b}$, 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 $b_0 = \tilde{\tilde{b}}$, it is the former path that now features $\mathcal{H}_\lambda = \dot{b} = 0$

and that is therefore dominated. By continuity,²⁵ the path that leads to b_L^* is therefore optimal for b_0 close enough to \tilde{b} , whereas the path that leads to b_H^* is optimal for b_0 close enough to $\tilde{\tilde{b}}$. 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 $\hat{b} \in (\tilde{b}, \tilde{\tilde{b}})$ 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 10 for an illustration: the bold segments of the two stable arms indicate the optimal selection among the two candidate paths.²⁶

So far, we focused on region M. In region H ($b_0 \geq b_{\text{bliss}}$), the analysis is similar to Figure 9. That is, there is a threshold $b_{\text{skiba}} \in (b_{\text{bliss}}, \bar{b})$ such that, as long as $b_0 \in (b_{\text{bliss}}, b_{\text{skiba}})$, 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 $b_0 \geq b_{\text{skiba}}$. Finally, in region L ($b_0 < b_{\text{seign}}$), there is a unique candidate path, one leading to b_L^* .

The kind of optimal policy illustrated in Figure 10 has the following properties: (i) whenever $b_0 < \hat{b}$, debt converges monotonically to b_L^* ; (ii) whenever $b_0 \in (\hat{b}, b_{\text{skiba}})$, debt converges monotonically to b_H^* ; and (iii) whenever $b_0 \geq b_{\text{skiba}}$, 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 $b_{\text{skiba}} > b_{\text{bliss}}$ 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 10. One of these variants is illustrated in Figure 11, the other in Figure 12. 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 11, it is optimal to converge to b_H^* for all $b_0 < b_{\text{skiba}}$; and in the case seen in Figure 12, it is optimal to converge to b_L^* .

Let us fill in the details, starting with Figure 11. Unlike Figure 10, 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 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

²⁵Here, 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).

²⁶In 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 11: Optimal to Converge to b_H^* for all $b_0 < b_{skiba}$

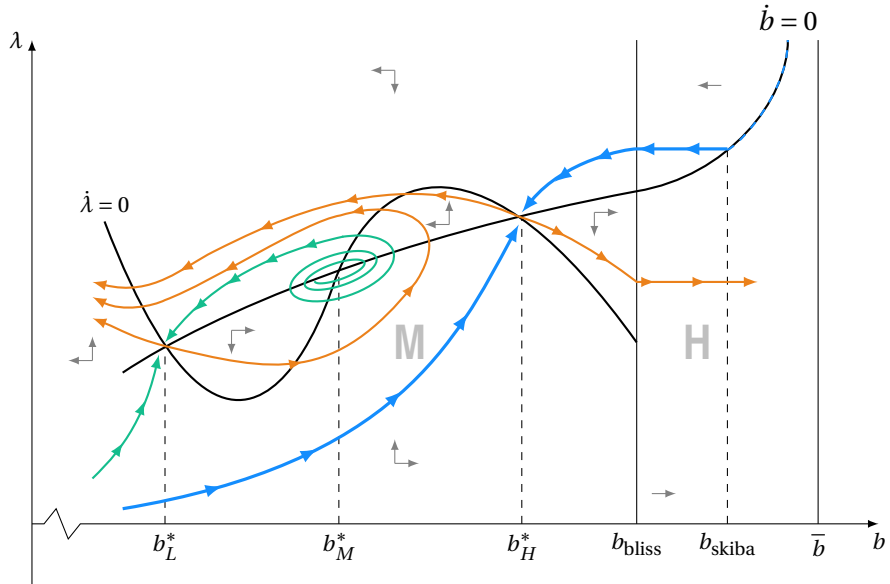
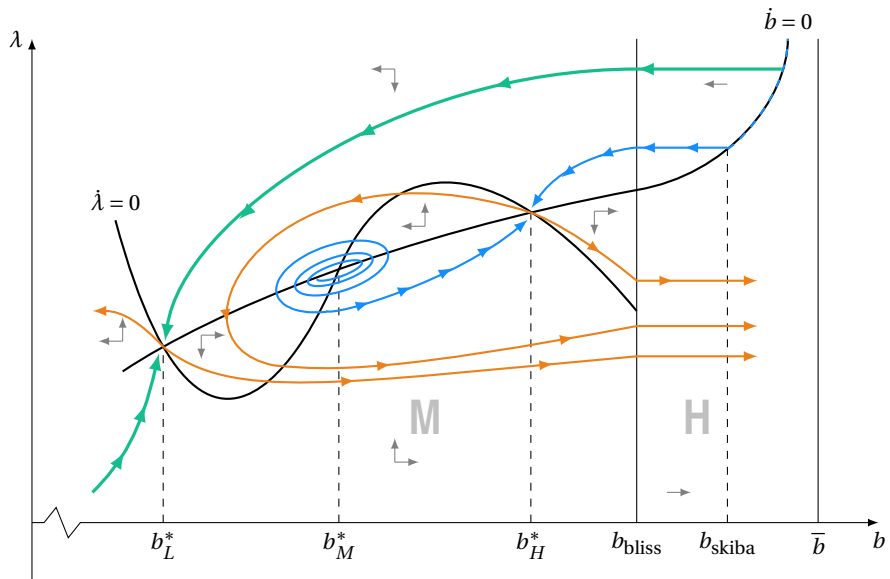


Figure 12: Optimal to Converge to b_L^* for all $b_0 < b_{skiba}$



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 10, the path that leads to b_H^* in Figure 11 is now the optimal path for all $b_0 < b_{\text{skiba}}$.

Consider next Figure 12. This illustrates a diametrically opposite scenario from that shown in Figure 11: 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 10 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 $b_0 \leq b_{\text{bliss}}$.

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 $\Omega_b = 0$.

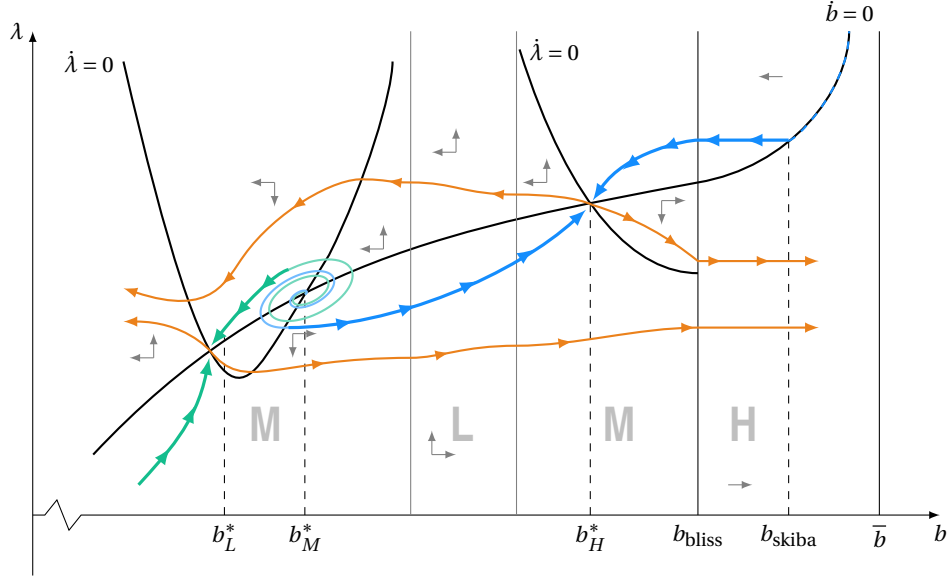
The number of possible scenarios would increase if we allowed γ and ψ to intersect more than three times. Yet an additional layer of complexity emerges if the assumption that $\pi(b)b$ is single-valued is relaxed. The tripartite structure of the phase diagram is then lost. Instead, the phase diagram 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 $\underline{b} < b_1 < b_2 < \dots < b_N < b_{\text{bliss}}$, 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 $b \in (b_N, b_{\text{bliss}})$, 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 $b \in (b_N, b_{\text{bliss}})$, 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 13. 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 13: Multiple Regions



Notwithstanding all the complexity, we can thus establish the following result, which offers a qualified generalization of Proposition 12 in our benchmark.

Proposition 14. *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 $\pi(b)b$.²⁷ 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 continuity and differentiability of γ and ψ , there exists at least one point $b^* \in (b^\ddagger, b_{bliss})$ 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

²⁷Because $\pi(b)b$ 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} .

such points, consider the highest one. By Proposition 13, we know that this steady state is saddle-path stable. Similarly to Figure 9, 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.²⁸

D.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.²⁹

Lemma 7. *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 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

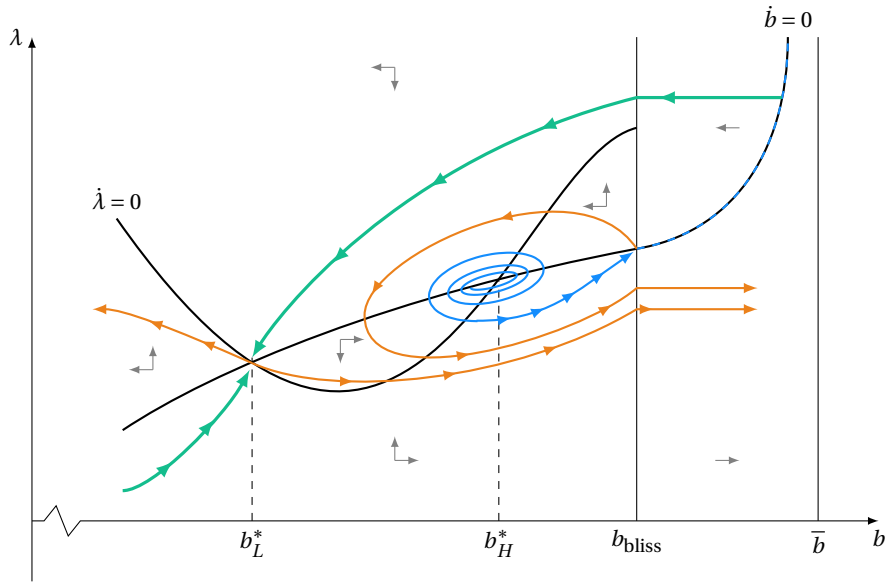
²⁸This argument mirrors Theorem 2 in Brock and Dechert (1983). Applied to our setting, this theorem states that, whenever the policy rule of the costate features a discontinuous jump, this jump is downward. By the same token, as we move from higher to lower levels of debt, the costate can only jump upwards, which means that lower levels of debt are necessarily associated with convergence to weakly lower steady states.

²⁹In 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.

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 12, necessity in Proposition 11. In the more general case allowed here, sufficiency remains valid by Proposition 14, but necessity may not apply.

Figure 14: No Satiation Despite $\psi_{bliss} < \gamma_{bliss}$ (or g low enough)



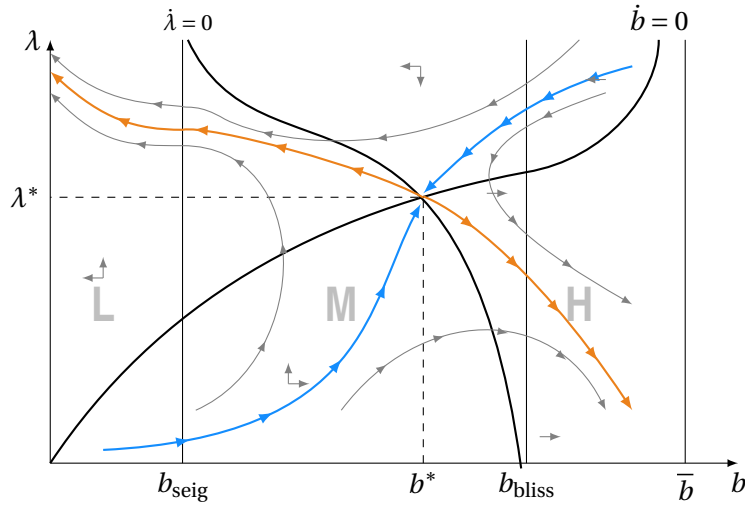
We illustrate this in Figure 14. As in our benchmark (see Figure 8 in particular), letting $\gamma_{bliss} > \psi_{bliss}$ guarantees the local existence of a candidate path that leads to satiation: for some $\epsilon > 0$ and all $b_0 \in (b_0 - \epsilon, b_{bliss})$, there exists a path that satisfies 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 6, 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.

D.5 Beyond the Benchmark: Relaxing Auxiliary Assumption B0

In the preceding analysis we relaxed Auxiliary Assumptions B1-B3 but maintained B0. We now do the converse, that we let $\pi(b)$ and $V'(b)$ turn negative for $b > b_{\text{bliss}}$. As mentioned in the main text, this amounts to changing the interpretation of region H (i.e., the region where $b > b_{\text{bliss}}$) from “harmless satiation” to “harmful excess.”

In this case, the phase diagram takes the form illustrated in Figure 15. Relative to our benchmark, there are two changes: in region M, the $\dot{\lambda} = 0$ locus presents an asymptote at $b = b_{\text{bliss}}$; and in region H, $\dot{\lambda}$ is now negative instead of 0. throughout region H. It follows that the Skiba point now coincides with \bar{b} , and the economy converges to b^* for *every* initial level of debt $b_0 \in (0, \bar{b})$.

Figure 15: Relaxing Assumption B0



What if we relax *both* B0 and B1-B3? Clearly, this is a hybrid of the present and the previous scenarios: relaxing B1-B3 allows multiple interior steady states below b_{bliss} , while relaxing B0 guarantees that $\dot{\lambda} < 0$ for $b > b_{\text{bliss}}$ and hence that the economy converges to a steady state below b_{bliss} for any initial position and any g . We conclude that relaxing B0 substitutes for the requirement that g is high enough and only reinforces the rationale for focusing on the scenario in which the economy converges to a steady state below satiation.

D.6 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

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

be the set of the points at which γ intersects ψ from above. As shown in Proposition 13, these points identify the saddle-path stable steady states of the ODE system.³⁰ Depending on primitives, $B^\#$ may be empty, or may contain an arbitrary number of elements.³¹ 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 6). To illustrate, consider Figure 10. 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 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 $\lim_{b \uparrow \hat{b}} \Lambda^{opt}(b) > \lim_{b \downarrow \hat{b}} \Lambda^{opt}(b)$,³² 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

³⁰In 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 $B^\#$.

³¹We 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.

³²At 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.

$\hat{b} = \underline{b}$. 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 14 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 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 9 and 10, B^* is a strict subset of $B^\#$ in Figures 11 and 12.

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 can be ruled out by assuming at least one of the following two conditions: that g is high enough, or $\psi_{bliss} > \gamma_{bliss}$; or that $V'(b)$ and $\pi(b)$ turn negative beyond the bliss point.

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 taxonomy of all the economies under consideration and of all possibilities that *do* obtain for some specification of (U, V, π, g) .

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, sufficient conditions for an economy to belongs to the last class are: that the need for fiscal space, as parameterized by g , is sufficiently high; and/or the social and private values of public debt turn negative at $b > b_{bliss}$.

Proof. That any economy must belong to one of these three classes follows from Theorem 2. That the three classes are non-empty follows from the examples we have already provided. Finally, the claimed sufficiency follows from the previous analysis as well. \square

D.7 Local Dynamics and Local Comparative Statics

We conclude this Online 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 15. *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 13, 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 13, that the eigenvalue associated with the stable arm is given by $\mu = \frac{\rho - \sqrt{\Delta}}{2}$ with $\Delta > 0$ (see proof of Proposition 13). 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 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

Proposition 16. *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 D.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.