Public Debt as Private Liquidity: Optimal Policy[*]

George-Marios Angeletos
MIT

Fabrice Collard
Toulouse School of Economics

Harris Dellas
University of Bern

May 3, 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 in the private sector; but this depresses the premium that the private sector is willing to pay for this service, which in turn can reduce fiscal space. This trade off shapes the optimal long-run quantity of public debt. It justifies a departure from tax smoothing in the short run. It helps clarify the precise circumstances under which a low risk-free rate represents an opportunity of "cheap borrowing" by the government. It seems robust to different micro-foundations of the "services" provided by public debt. And it may subsume additional considerations, such as whether public debt crowds out capital.

[*]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

Various commentators have argued that the low interest rates on US public debt present an opportunity for cheap government borrowing. While tempting, this argument finds no place in a basic, complete-markets, representative-agent setting. As illustrated most recently by Barro (2021), such a setting can rationalize the observed rates on public debt, provided one accommodates appropriate aggregate risk ("disaster risk"). But in such a setting, tax smoothing remains optimal regardless of how low or cyclical these rates might be—there is never an actual opportunity for cheap government borrowing.

In this paper we revisit this logic in the presence of financial frictions. Such frictions can not only rationalize a low risk-free rate but also bestow public debt with a special function: public debt may ease frictions and raise welfare by contributing to the supply of assets that private agents can use as collateral or buffer stock, (Woodford, 1990; Aiyagari and McGrattan, 1998; Holmström and Tirole, 1998). Corroborating this idea, Krishnamurthy and Vissing-Jorgensen (2012) document that the spread between government and private bonds, an empirical proxy of the market's valuation of the "services" provided by public debt, decreases with the supply of public debt.

We show how this mechanism not only allows for government borrowing to be cheaper—in a proper sense—when the relevant market friction is tighter, but also introduces a novel policy trade off: the flip side of easing the friction and improving market allocations is an unwelcome increase in the government's cost of borrowing. Our main contribution 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 aims at distilling the essence of the issue at hand. This problem is basically the same as that in Barro (1979), except that public debt is non-neutral for two reasons in addition to distortionary taxation. First, the interest rate paid on public debt depends on its quantity, reflecting the *private* value of the liquidity, collateral, or other "convenience" services provided by it. Second, welfare also depends on the quantity of 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 an economy 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), as well as another in which public debt serves as collateral (Holmström and Tirole, 1998). Although highly stylized, these examples illustrate how the lessons provided below may transcend various applications. At the very least, they help fix ideas: 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 create assets or "inside" liquidity; and the private value of this service may differ from it 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, other hand, translates to a Laffer curve: the relevant rents are maximized at a quantity of public debt strictly less than the bliss point.

Because taxes are smoothed in a steady state, one may expect the long run to be shaped solely by the other two forces. This intuition has to 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 is basically right: 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 have emphasized: 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 a spike in idiosyncratic risk, a shortage of collateral, or a flight to safety.

Additional points. The trade off we study in this paper finds support in a literature that documents that the interest-rate spread between government bonds and high-grade corporate bonds is both cyclical and sensitive to the supply of public debt (Krishnamurthy and Vissing-Jorgensen, 2012; Greenwood and Vayanos, 2014). At the same time, our analysis shifts the spotlight from that spread to the wedge between the interest rate on government bonds and the underlying social discount rate (equivalently, the risk-free rate that would have obtained in a frictionless, representative-agent, counterfactual). The optimal policy hinges on the properties of this wedge, not of the aforementioned spread.

This and a few additional lessons are made possible by nesting a few 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 change—and so do a few other "details," such as the substitutability between corporate and government bonds or the possibility that public debt crowds out (or, in) capital accumulation. But the reduced-form problem remains the same, and so do our results.

This helps illustrate how our insights may transcend various applications, including the literature spurred by Blanchard (2019). Some of the issue that take center stage in this literature remain outside our scope, such as debt sustainability and debt bubbles (Brunnermeier, Merkel and Sannikov, 2021; Mehrotra and Sergeyev, 2020; Reis, 2021), inequality (Aguiar, Amador and Arellano, 2021), capital taxation (Bassetto and Cui, 2021), and the ZLB constraint on monetary policy (Mian, Straub and Sufi, 2022). Still, because much of this literature shares with our paper the key idea that public debt serves as a buffer stock or collateral, it naturally inherits the policy trade off we have highlighted here.

Turning to a technical aspect of our paper, we observe that the policy problem of interest may be non-convex and, hence, the standard first-order approach may fail: there can exist multiple paths that satisfy the planner's Euler and transversality conditions, and the challenge is to find out which one of them is truly optimal. We address this challenge by adapting the optimal control methods of Skiba (1978) and clarify how our main lessons are robust to the existence of multiple steady states. Finally, we 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.

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 flaw 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

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

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, or 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. What is endogenous is the time path of $s(t)$ and $b(t)$. To ease the notation, we henceforth suppress their dependance on t and write the planner's problem as follows.

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

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

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

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* $b_{bliss} \in (0, \bar{b})$ *such that* $V'(b), \pi(b) > 0$ *if* $b < b_{bliss}$, *and* $V'(b), \pi(b) \leq 0$ *otherwise.*

A0 and A1 require little explanation: the former is technical; the latter stylizes the distortionary effects of taxation and follows directly from Barro (1979). What is new here is A2. This assumption stylizes 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). We corroborate these interpretations, and shed additional light on the economics *behind* our reduced-form problem, with the help of two micro-founded examples.

3 A Detour: Stylized Micro-foundations

In this section we show how to derive our reduced-form problem from two stylized but micro-founded economies: an OLG economy with uninsurable idiosyncratic risk; and an infinite-horizon variant with a certain kind of collateral constraints. 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.

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

[2] To 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 $b_{\text{bliss}}$from below.

3.1 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 7 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 as in Aiyagari and McGrattan (1998).

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

$$\mathscr{U}_{it} = c^y_{it} - v(h_{it}) + \delta \mathbb{E}_{it}\left[u(c^o_{it+1})\right],$$

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

$$c^y_{it} + q_t a_{it} = (1-\tau_t) w_t h_{it} \qquad \text{and} \qquad c^o_{it+1} = e_{it+1} + a_{it}.$$

The notation is as follows: c^y_{it} and c^o_{it+1} 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 fixed scalars $k>0$,$\epsilon \in (0,k)$, 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); $\lim_{c \to 0} u'(c) = \lim_{h\to\infty} v'(h) = \infty$ and $\lim_{c\to\infty} u'(c) = \lim_{h \to 0} v'(h) = 0$ (the familiar Inada conditions are satisfied).

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. Finally, the government's 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 stock of debt inherited in period t, g is the exogenous level of government spending, $q_t b_t$ are the proceeds from new debt issuance, and $s_t \equiv \tau_t w_t h_t$ is tax revenue. And the planner's objective is to maximize the following welfare measure:

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

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

The demand for safe assets. Since young households are identical, they all make the same choices in equilibrium: 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:

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

Since market clearing in the bond market imposes $a_t = b_t$, we infer that

$$q_t = Q(b_t; \epsilon, k), \tag{5}$$

where the function Q is defined by the right-hand side of (4), replacing a_t with b_t. This function can be read interchangeably as the equilibrium price function for government debt and as the inverse of the aggregate demand for assets. Because $u'' < 0 < u'''$, the following properties are true:

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, the present proxy for inside liquidity* ($Q_k < 0$).

These properties translate to our reduced form by letting $\pi(b) \equiv Q(b) - \beta$ and help stylize the demand for 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 at large: that the equilibrium risk-free rate falls with more idiosyncratic risk (Bewley, 1980; Aiyagari et al., 2002) 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 pledgeable income a la Holmström and Tirole (1998), a point that will become clear in our second example below and also in the extension discussed in Section 7.

Let us now turn attention from the *private* value of government debt to the corresponding *social* value. By martket clearing in the goods market, $c_t^y + \int c_{it}^o di = Ah_t + k$. By aggregating the budgets of the old and market clearing in the asset market, $\int c_{it}^o di = b_{t-1} + k$. Combining, we infer that $c_t^y = Ah_t - b_{t-1}$. Next, using these properties to compute welfare, we conclude that

$$\mathscr{W} \equiv \sum_{t=0}^{\infty} \beta^t \int \mathscr{U}_{it} di = \sum_{t=0}^{\infty} \beta^t \left\{Ah_t - v(h_t) + V(b_t; \epsilon, k)\right\} + \zeta$$

where $\zeta \equiv \frac{1}{1-\beta}k + b_{-1}$ is a constant (invariant to policy) and

$$V(b) \equiv \varphi\delta u(k-\epsilon+b) + (1-\varphi)\delta u(k+\epsilon+b) - \beta^{-1}b.$$

defines the (total) social value of public debt. It is easy to verify that $V(b) = \pi(b) \equiv Q(b;\epsilon,k) - \beta$, that is, the (marginal) private and social values coincide. This is *not* true in our second example below, because of pecuniary externalities; but our reduced form accommodates both scenarios.

Reduced form. By the Inada conditions on u, $\lim_{b\to\infty} Q(b) = 0$ necessarily and $Q(0) > \beta$ for sufficient risk (ϵ close to k). It follows that 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 readily 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 by verifying that U satisfies Assumption A1.

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

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

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

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

The following remark completes the picture. So far, we have emphasized the function of government debt as a buffer stock against idiosyncratic risk, because this connects to the literature we are primarily concerned with. But in the present example, government debt plays 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—as collateral, buffer stock, or vehicle for life-cycle savings—may be important for certain questions, it is not essential for our purposes.

3.2 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, which becomes ripe in the "morning." 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. The afternoon good maps to the standard consumption good in the textbook Ramsey paradigm; the morning good helps introduce a tractable form of the friction of interest.

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

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

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

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

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. The modeling role of these shocks is to introduce a value for net trades of the morning good across the agents. For simplicity, there are only two idiosyncratic states: either the household has a high taste for the morning go 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 respectively probabilities φ and $1-\varphi$, for some $\varphi, \epsilon \in (0,1)$.[3]

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:

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

where y_t^{def} denotes the income received in the (off-equilibrium) event of default.[4] The presence of a_{it} 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.[5]

Finally, the production side of the economy is the same as in the previous example, and so is the government's budget constraint. What changes—naturally—is the government's objective, which is now given by utility behind the veil of ignorance, namely by (8) 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 scare, 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 $b < b_{bliss}$, 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 $b > b_{bliss}$, 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 $b < b_{bliss}$ increase with ϵ, the amount of idiosyncratic risk, and decrease with ξ, the pledgeable income.

[3]This 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.

[4]For 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 $h_t^{\text{def}} \equiv (v')^{-1}(w_t)$.

[5]Applying the same logic to inter-period borrowing, we infer that the following constraint must also hold: $-a_{it} \leq \xi y_t^{\text{def}}$. But this constrain never binds in equilibrium, because consumers choose the same savings and, by market clearing, these saving equal $b_t \geq 0$. We can thus ignore this constraint and focus on (10). An earlier version of our paper introduced additional heterogeneity and let some agents issue private debt in the afternoon, which in turn served as additional "inside collateral" for the next morning. This allowed the constraint $-a_{it} \leq \xi y_t^{\text{def}}$ to bind for some agents, 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 that in our first example, where $\pi(b)$ was *negative* for $b > b_{\text{bliss}}$, reflecting an over-abundance of consumption at old age. 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.

From consumer frictions to production frictions. In Appendix A we study a third example, in which a collateral constraint impedes the efficient allocation of capital and labor across firms. This further increases the proximity to Holmström and Tirole (1998) 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 Liquidity provision versus interest-rate suppression

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

$$\max_{s_1, b_1, s_2} \left\{ U(s_1) + V(b_1) + \beta U(s_2) \right\}$$

$$\text{s.t.} \quad Q(b_1)b_1 + s_1 = g_1 \quad \text{and} \quad s_2 = g_2 + b_1$$

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

$$b^* = \arg\max_{b_1} \left\{ \lambda_1 Q(b_1) b_1 + V(b_!) - \beta\lambda_2 b_1 \right\}. \tag{11}$$

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 $Q(b) = \beta + \pi(b)$ and let $\lambda_1 = \lambda_2 = \lambda$, which amounts to presuming perfect tax smoothing. Condition (11) then reduces to

$$b^* = \arg\max_b \Omega(b,\lambda), \qquad \text{where} \qquad \Omega(b,\lambda) \equiv V(b) + \lambda\pi(b)b. \tag{12}$$

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 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 is "cheap" when 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 $b_{\text{seig}} \equiv \arg\max_b \pi(b)b$ 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 $b = b_{\text{bliss}} \left(\equiv \arg\max_b V(b)\right)$. For any $\lambda \in (0,\infty)$, the optimal debt issuance strikes a balance between these two goals, i.e., $b_{\text{seig}} < b^* < b_{\text{bliss}}$. The stronger the fiscal preoccupation, as measured by λ, the closer b^* is to b_{seig}.

Clearly, such fiscal considerations are present only when $\lambda > 0$, i.e., when taxation is distortionary. But distortionary taxation *alone* does not suffice for satiation not to be optimal. Suppose in particular that $\pi(b)$ was either 0 or positive but invariant to b. Then, b^* would have equaled b_{bliss}, no matter how large λ might have been. This underscores the following point:

Observation. *On the margin, the cost of more liquidity provision is not higher taxes but rather higher interest rates: preserving the shortage of aggregate collateral makes sense because and only because it helps suppress the government's cost of borrowing.*

This puts the spotlight on the trade off between liquidity provision and interest-rate suppression. But it abstracts 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 fill in these holes below.

4.2 Balancing the three objectives

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

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

where $\Omega(b,\lambda)$ has the same meaning and definition as in our two-period heuristic above. But whereas there we treated λ as exogenous and forced it to be constant over time, here we recognize that the path of

λ is endogenous and is the mirror image of the optimal path of taxes. Indeed, since s oughts to maximize the Hamiltonian, we have that $\lambda = -U'(s)$, which means that λ measures the tax distortion. Furthermore, the optimal path for λ must satisfy the planner's Euler condition, which can be written as

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

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 heuristic used earlier. We will 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 time trend 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, 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. This is where the methods of Skiba (1978) come in handy. The general case is studied in Appendix B. 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 translates to a sufficiently large shadow value for depressing the government's cost of borrowing. Together, these assumptions suffice for the following result:

Theorem 1. *There exists 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 below satiation and $[0, b_{\text{skiba}})$ as its basin of attraction.[6] A detailed proof is provided in Appendix B. Here, we sketch out the main ideas with the help of the phase diagram in Figure 1.

Figure 1: Phase Diagram and the Optimal Path.

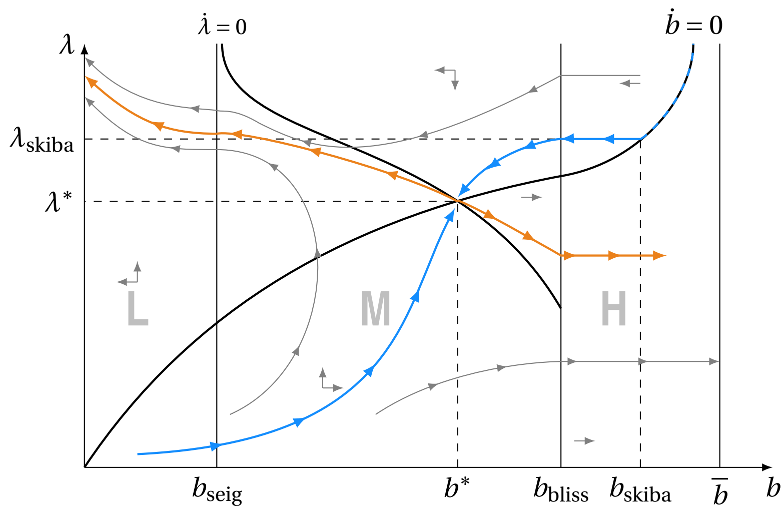

To start with, consider the $\dot{b} = 0$ locus. This corresponds to balanced budget, is given by $\lambda = -U'(s)$ with $s = g + \big(\rho - \pi(b)\big)\, b$, and is illustrated in the figure by the curve labeled "$\dot{b} = 0$." This curve is upward sloping because a higher level of debt requires a higher rate of taxation for the budget to be balanced.[7]

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

In region L, which is defined by $b < b_{\text{seig}}$ and corresponds to the upward-sloping portion of the debt Laffer curve, increasing b raises both $V(b)$ and $\pi(b)b$, so there is no trade off between "liquidity provision" and "interest rate manipulation." It follows that, for any $\lambda \geq 0$, the marginal value of raising debt is positive, $\Omega_b(b, \lambda) > 0$, and therefore also $\dot{\lambda} > 0$. That is, the locus $\dot{\lambda} = 0$ does not exist in region L.[8] By direct implication, there is also no steady state within this region.

[6] The threshold b_{skiba} is an example of the "Skiba points" that emerge in non-convex, optimal-control problems.

[7] To be precise, this curve is upwards slopping if and only if the interest-rate bill, $R(b)b = (\rho - \pi(b))b$, is increasing in b. This is necessarily the case in regions M and H but may fail in region L. Nonetheless, this possibility does not affect the result, because, as explained below, it is always optimal to leave region L and enter region M.

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

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

The graph of $\gamma(b)$ gives the curve labeled "$\dot{\lambda} = 0$" in Figure 1. This curve is decreasing, reflecting the fact that a higher λ shifts the balance in favor of interest-rate suppression. The balanced-budget line, on the other hand, is increasing, reflecting the higher tax distortion implied by higher level of debt. It follows that the two lines intersect at a unique point $(b,\lambda) = (b^*, \lambda^*)$, which identifies the unique steady state within regions L and M. To the left of this point, debt and taxes increase over time; and to the right of it, debt and taxes fall over time.

Finally, consider region H, which corresponds to levels of debt above satiation. In this region, we have that $V(b)$ 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 the *entirety* 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 exist 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.

Along this plan, debt falls gradually, crossing b_{bliss} within finite time and converging asymptotically to b^*. Compared to the Barro-like alternative, this plan necessitates 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 interest-rate suppression, once debt has fallen below b_{bliss}. Provided that this happens fast enough, which is the case if $b_0 \in (b_{\text{bliss}}, b_{\text{skiba}})$, the sacrifice in terms of tax smoothing is justified. The converse is true if $b_0 > b_{\text{skiba}}$. And this completes the characterization of the optimal plan.

4.4 The optimal long-run quantity of public debt

Suppose that $b_0 < b_{\text{bliss}}$, or at least $b_0 < b_{\text{skiba}}$. This rules the uninteresting scenario in which the financial friction is *never* binding. The following result zeroes in on the determination of the resulting steady state.

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

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

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

This connects the steady state to our two-period heuristic: b^* balances the desire to ease the financial friction in the private sector with the value of suppressing the government's cost of borrowing. The only subtlety is that the government budget imposes a 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 characterized the optimal policy path starting from arbitrary initial condition b_0, we established that this path converges to a point b^* that is itself invariant to b_0, and we characterized this point. By contrast, Aiyagari and McGrattan (1998) maximize welfare over steady-state policies, which herein translates to imposing budget balance in all periods, and hence also $b_t = b_0$ for all t, and maximizing welfare over b_0.[9] As explained in Appendix C.3, this over-estimates the opportunity cost of liquidity provision and, therefore, under-estimates the optimal long-run quantity of public debt.

4.5 Relaxing Auxiliary Assumptions B0-B3

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

Start with B3, which required that g (and the need for fiscal space) be high enough. 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. While logically possible, this scenario is the least interesting for our purposes.

Consider next B0, which let 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 $b > b_{\text{bliss}}$, capturing 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 scenario with $b^* < b_{\text{bliss}}$.

[9] Aiyagari and McGrattan (1998) were forced to do this exercise because of the complexity of their incomplete-markets environment. In particular, the key difficulty in that paper, which instead we assume away, is that the wealth distribution is a relevant state variable for the aggregate dynamics and hence also for the planner's problem.

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. 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, the steady-state level of debt drops from b^*_{old} to b^*_{new} and the optimal dynamic response is as follows: taxes initially increase more than the increase in g, in order to allow for debt and interest rates to fall; but thanks to this, the government can eventually afford a lower increase in taxes than the increase in g.

Figure 2: Permanent Increase in Government Spending

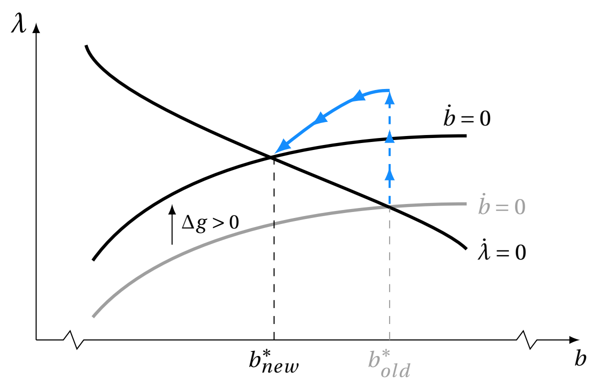

Proposition 5. *An unanticipated permanent increase in g calls for an increase in taxes by more than one-to-one in the short run and by less than one-to-one in the long run.*

[10] Throughout this section, we let public debt be risk-free, as in Barro (1979) and Aiyagari et al. (2002). The opposite scenario, which allows public debt to be fully state-contingent, is considered in Appendix C.1. As in Lucas and Stokey (1983), this scenario allows the government to insure its budget against shocks; but now the optimal state-contingency balances such insurance with the objectives of providing liquidity and suppressing interest rates.

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 transitory fiscal shocks, what is often referred to in the literature as "wars". We illustrate this in Figure 3, using a stochastic example in which government spending follows a symmetric two-state Markov process, with the probability of staying in the same state equal to 0.9. 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 permits the planner to moderate the increase in interest rates, which would have further tightened the budget.[11]

Figure 3: Optimal Response to a War

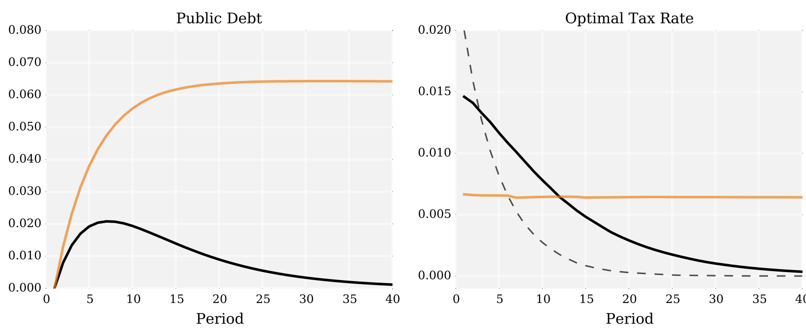

Debt and Taxes in our Model; Debt and Taxes in Barro/AMSS; Government Spending.

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

Because this raises the social value of liquidity and the profit from interest-rate suppression in proportion to each other, it leaves the $\dot{\lambda} = 0$ locus unaffected. If the $\dot{b} = 0$ locus had also been unaffected, the optimal policy response would have been to stay put. But $\dot{b} = 0$ actually shifts down, because the shock reduces the interest-rate costs on public debt, thus also reducing the taxes needed for balanced budget. In a nutshell, the private sector's flight to quality brings a bonanza for the government.

[11] If 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.

Proposition 6. *A flight to quality, modeled as a joint increase in the social and private value of liquidity, is equivalent to a positive income shock in the government budget. This justifies more government provision of liquidity, because and only because it reduces the shadow cost of such provision.*

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

Figure 4: IRFs to a Financial vs Traditional Recession

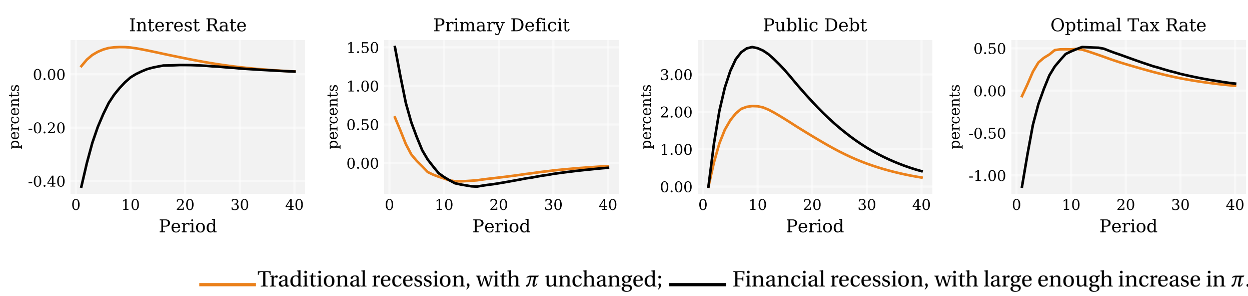

This provides a formal basis for the argument made by Paul Krugman, Brad DeLong and others that

[12] Formally, 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 transitory negative shock to ω_t.

the reduction in the government's cost of borrowing 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 variation in the observed interest rate *per se*, but rather the extent to which this represents variation in the wedge between that rate and the counterfactual rate that would have obtained in the absence of a financial friction. Were ρ to drop along side the interest rate, leaving π the same, public debt would *not* be cheaper.

6 Discussion

In this section we circle back to the micro-foundations behind our reduced form so as to shed additional light on the following issues: the possibility that public debt crows out (or, in) capital; the importance of the assumption that debt is non-neutral; and the precise "premium" that matters for our analysis.

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, modulo the following change in V:

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

At the same time, because the equilibrium quantity of capital coincides with the arg max of this problem, the following is true: when the government issues more debt, it crowds out capital, similarly to Diamond (1965) and Aiyagari and McGrattan (1998). Summing up:

Proposition 7. *In the extension described above, public debt issuance 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 applies if we add physical capital in our second example and let it serve as collateral in morning transactions.[14] Together, these extensions illustrate how our insight can be robust to the possibility that public debt crowds out "inside liquidity," or discourages the "private creation of collateral." Under our prism, such crowding out is *not* an additional, separate element of the costs and benefits of debt issuance; it is subsumed in the trade off we have already analyzed.

Finally, one can reverse the crowding-out property, again without affecting our results. We offer such an example in Appendix A. Compared to Aiyagari and McGrattan (1998) and the examples discussed so

[13] This 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, because this is outside the scope of our paper.

[14] This was shown in an earlier draft.

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 and as money. In the micro-foundations that underly our reduced form, public debt is non-neutral because and only because private borrowing capacity is not sufficiently sensitive to future tax obligations. To see this more clearly, modify our second example 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 issuance is perfectly offset by a commensurate reduction in pledgeable income. The same point applies to Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998): if borrowing constraints adjusted to future tax obligations, public debt would be neutral in all those papers, too. This clarifies the precise reason why public debt issuance can ease the financial friction and can thereby have a causal effect on both welfare and interest rates.

Put differently, the above discussion explains what it takes for public debt to have money-like attributes. At the same time, it is worth highlighting that our analysis departs from the Friedman-rule literature in that it ties private liquidity to the entire net liability of the government, as opposed to one particular component of its portfolio. To elaborate on the role played by this assumption, consider again the second example from Section 3 and suppose that the government enacts a law that prohibits the use of corporate bonds as collateral in morning transactions. This 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).

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 liquidity premium measured in these papers does not map to π (nor to V') in our setting. To see this, let r_{priv} and $r_{\text{gov}}(\equiv r)$ denote the interest rates on, respectively, private and government bonds. In the examples of Section 3, private and government bonds were perfect substitutes, so that $r_{\text{priv}} = r_{\text{gov}}$. The object measured by the aforementioned works was therefore identically zero, but the object that was relevant for our results, $\pi = \rho - r_{\text{gov}}$, was of course positive (and also sensitive to b). Furthermore, our second example could have accommodated $r_{\text{gov}} < r_{\text{priv}}$ by letting government bonds be a superior form of collateral in morning trades; but no matter the difference between

r_{gov} and r_{priv}, the relevant "sufficient statistic" for the policy problem would still be $\pi = \rho - r_{\text{gov}}$. [15]

A similar point applies to the difference between the expected return to capital, r_K, and the interest rate on government debt, which is a focal point of Blanchard (2019): if this difference reflects only a compensation for aggregate risk, as it is, e.g., in the case in Barro (2021), it need not matter for optimal policy. This is precise the case in a variant of the example of Section 6 that lets 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.

Needless to say, these insights are not universal: in richer models, $\rho - r_{\text{gov}}$ may not be a sufficient statistic, plus the relevant social discount rate may be more complicated than the fixed parameter ρ here. Still, we hope that this discussion helps redirect future work on the object that emerges as the most relevant one under the prism of our analysis.

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 the amount of collateral or liquidity, as in Woodford (1990), Aiyagari and McGrattan (1998), Holmström and Tirole (1998) and a recent literature spurred by Blanchard (2019). Issuing more debt raises welfare by easing the underlying financial friction. But it also tightens the government budget by raising interest rates relative to the social discount rate.

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

These insights were delivered with the help of a reduced-form Ramsey problem, whose main difference from the textbook counterpart (Barro, 1979; Lucas and Stokey, 1983) was to let 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 variety of 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 private and the social value of liquidity, due to pecuniary externalities operating via collat-

[15] In an earlier version of that example, we had made sure that some private agents find it strictly optimal to issue bonds in the afternoon of each period by introducing persistence in the θ shock. Clearly, the same is true in the present version if we introduce heterogeneity in discount rates. In both cases, one can go on to rationalize $q_{\text{gov}} > q_{\text{priv}}$, or equivalently $r_{\text{gov}} < r_{\text{priv}}$, by letting privately-issued bonds enter the collateral constraint (10) with a haircut relative to government-issued bonds. Importantly, none of these modifications upset the reduced-form representation of the policy problem.

eral constraints; the degree of substitutability between government bonds and private bonds; and the possibility that public debt crowds out, or perhaps crowds in, physical capital.

Needless to say, we do not mean to suggest that one can always abstract from these issues or the relevant micro-foundations: the applicability of our reduced form relied on strong simplifying assumptions, which are unlikely to hold in practice. Still, by providing a benchmark where our reduced form was without any loss, we hope to have accomplished the following goals as effectively as possible. First, we shined the spotlight on a trade off that may transcend a wide variety of micro-foundations, namely the trade off between increasing the aggregate supply of liquidity or collateral, on the one hand, and keeping government borrowing cheap, on the other hand. Second, we clarified when such borrowing is truly cheap. And third, we alerted that some of the issues that the literature is often focused on could be of secondary importance relative to the trade off we have emphasized in this paper.

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, Albert Marcet, Thomas Sargent, and Jesus Seppälä. (2002) "Optimal Taxation without State-Contingent Debt." *Journal of Political Economy*, 110(6): 1220–1254.

Aiyagari, S. Rao, and Ellen McGrattan. (1998) "The Optimum Quantity of Debt." *Journal of Monetary Economics*, 42(3): 447–469.

Barro, Robert J. (1979) "On the Determination of the Public Debt." *Journal of Political Economy*, 87(5): 940–971.

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 miméo.

Brunnermeier, Markus K., Sebastian Merkel, and Yuliy Sannikov. (2021) "Debt As Safe Asset." *Princeton University mimeo.*

Dechert, W.D., and K. Nishimura. (1981) "A Note on Optimal Growth with a Convex-Concave Production Function." State University of New York at Buffalo Department of Economics Discussion Papers.

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.

Kiyotaki, Nobuhiro, and John Moore. (1997) "Credit Cycles." *Journal of Political Economy*, 105(2): 211–248.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. (2012) "The Aggregate Demand for Treasury Debt." *Journal of Political Economy*, 120(2): 233–267.

Lagos, Ricardo, and Randall Wright. (2005) "A Unified Framework for Monetary Theory and Policy Analysis." *Journal of Political Economy*, 113(3): 463–484.

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.

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

Online Appendices

A Variant Micro-foundations with Capital

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

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

Let κ_{it} denote the amount of capital owned by household i in the morning of period t. It is given by $\kappa_{it} = (1-\delta)\kappa_{it-1} + \iota_{it-1}$, where δ denotes depreciation and ι_{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.

Relative to the second example 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, public debt can have an ambiguous effect on capital accumulation. On the one hand, more public debt can crowd *in* capital via the aforementioned channel, namely by raising aggregate TFP and thereby the mean return to investment. On the other hand, more public debt can crowd *out* capital by offering a substitute form of collateral or buffer stock, as in Aiyagari and McGrattan (1998).

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

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

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

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

Proof. See Appendix D. □

To relate this proposition to Proposition 2, note that W, Q, and S capture, respectively, the per-period welfare flow, the market price of public debt and the tax revenue.[16] As we move from the second example of Section 3 to the new model, the micro-foundations that underlie these objects change, and so do their functional forms. For instance, the two distortions now have non-separable effects on welfare, interest rates, and the tax base. Yet, the strategy for obtaining the desired representation remains the same: the key step is to define W as the welfare flow that obtains when the planner takes as given (τ_t, b_t) and optimizes over the set of the cross-sectional allocations of labor, capital, and asset holdings and the aggregate supplies of capital and labor; Q and S are then defined by, respectively, the interest rate that supports the best implementable allocation and the primary surplus induced by it. Importantly, the only reason why W, Q and S depend on b is that the latter controls the financial friction. The representation obtained therefore encapsulates, once again, the dual role of the financial distortion on welfare and the government budget. What is new relative to the second example of Section 3 is that the financial friction affects the budget, not only via interest rates, but also via the tax base: by interfering with the allocation of capital, it affects wages, income, and tax revenue for any given tax rate. However, neither this feature nor the details of the underlying micro-foundations need alter the properties of optimal policy.

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

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

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

$$b(0) = b_0 \tag{19}$$

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

[16] In the second example of Section 3, we could express the policy problem in terms of s rather than τ because there was a one-to-one mapping between them. In the current model, this is no more true and, accordingly, we keep the tax rate as a control variable.

$W_b(b,\tau) = 0$ if $b \geq b_{bliss}(\tau)$; this allows for the possibility that the "satiation point" beyond which the friction ceases to bind may depend on the tax rate. Similarly, let $b_{seign}(\tau) \equiv \arg\max\{\pi(\tau,b)b + S(\tau,b)\}$; this is the analogue to the level of debt that maximized seigniorage in 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. Adjusting the notion of "liquidity plus seigniorage" accordingly gives

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

We can express the planner's Euler condition as

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

which has exactly the same interpretation as its counterpart in Section 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 $T(\lambda) = \arg\max_{\tau}\{W(\tau,b) + \lambda[\pi(\tau,b)b + S(\tau,b)]\}$. We therefore obtain essentially the same ODE system as in the second example of Section 3; the underlying micro-foundations and some details are different but the essence remains the same.

Figure 5: Entrepreneurial Model

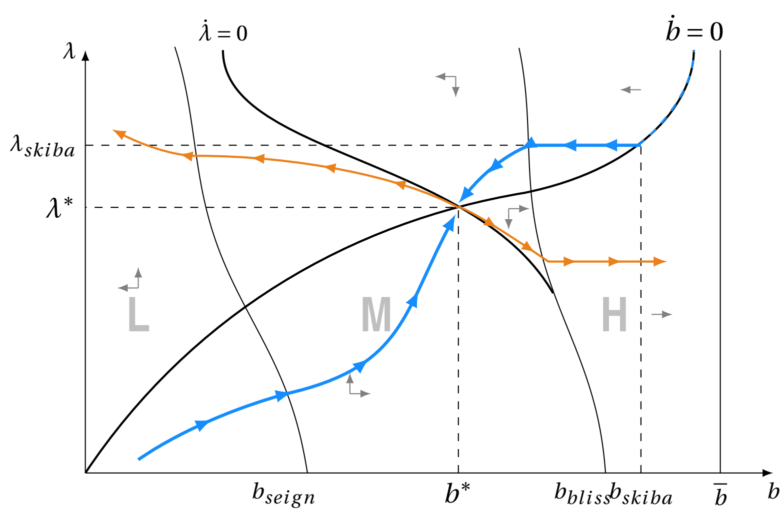

We illustrate this in Figure 5. For this example, we assume that $[\pi(\tau,b)b + S(\tau,b)]$, is single-peaked in b. This guarantees that the phase diagram can be split in three regions, similar to regions L, M and H in Figure 1. The boundaries of these regions are now curved, rather than vertical,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

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

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

Figure 6: Crowding in or out

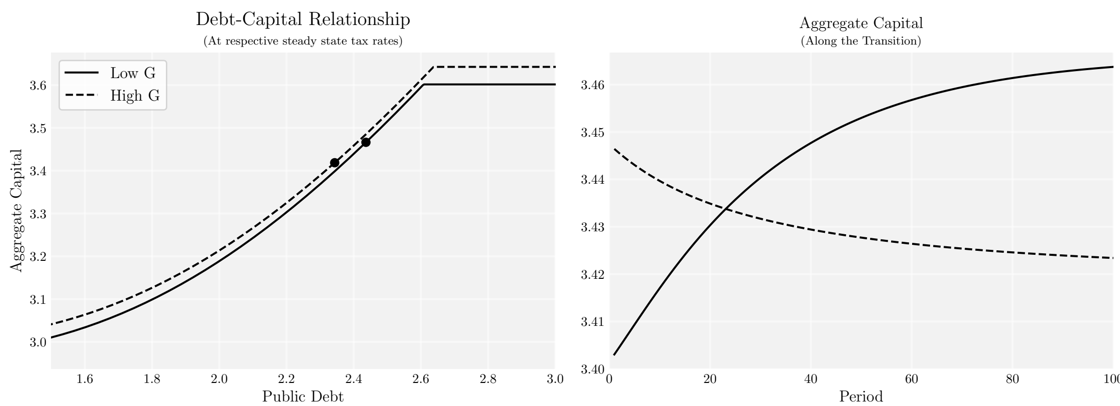

The left panel of Figure 6 considers the policy rule for aggregate capital. In particular, we consider two economies: one with a relatively low level of government spending (g =17% of steady-state output); and another with a relatively high level of government spending (g =27% of steady-state output)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

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

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

B Characterization of Optimal Plan

In this Appendix we offer a complete, self-contained, characterization of the solution to problem (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 show how Assumption B guarantees the existence of a unique steady state below satiation and prove Theorem 1.
- We show how, away from the aforementioned benchmark, it is possible to have multiple steady states below satiation, as well as no such steady state.
- We finally explain the precise sense in which the lessons obtained in the main text remain robust to the richer cases allowed here.

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

B.1 The ODE system

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

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

where $\Omega(b,\lambda) \equiv V(b) + \lambda\pi(b)b$ measures the social value of the liquidity services of public debt plus the profit made from providing these services, and λ measures the shadow value of tax revenue. Throughout this Appendix, we are ruling out both lump-sum taxes and lump-sum transfers. This allows the possibility that $\lambda < 0$, or equivalently $s < 0$ and $\tau < 0$, which means the planner may be using a distortionary subsidy in order to accumulate debt fast enough.[18]

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:

$$\dot{b} = \Psi(b,\lambda) \equiv g + \left(\rho - \pi(b)\right)b - s(\lambda), \tag{20}$$

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

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

in λ: a higher λ means higher taxes today, which in turn means lower debt tomorrow.[19] By the Implicit Function Theorem, there exists a function $\psi : [\underline{b}, \bar{b}) \to \mathbb{R}_+$ such that $\Psi(b, \psi(b)) = 0$ for all b; equivalently,

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

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

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

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. □

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,

$$\dot{\lambda} \quad = \quad v(b) - \lambda \pi(b) \left(\sigma(b) - 1\right). \tag{21}$$

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

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

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

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

We can then restate the Euler condition (21) as follows:

$$\dot{\lambda} = \begin{cases} v(b)\left[1-\frac{\lambda}{\gamma(b)}\right] & \text{if } b \in \Delta \\ 0 & \text{if } b \notin \Delta \end{cases} \tag{22}$$

By implication,

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

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

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

$$\frac{d[\pi(b)b]}{db} = \pi'(b)b + \pi(b) = -(\sigma(b) - 1)\pi(b). \tag{23}$$

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

B.2 The case studied in the main text

We now focus on a slightly more general case than the one studied in the main text—more specifically we dispense with Auxiliary Assumption **B3** and only maintain the following 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 $b \in (0, b_{bliss})$.

The first assumption imposes that the wedge between the social and the private value of collateral is invariant to b, the second guarantees that $\pi(b)b$ is single-peaked and also extends the aforementioned local monotonicity of σ to its entire domain. In the sequel, we will refer to the peak in $\pi(b)b$ as b_{seign}. This peak satisfies $\pi(b_{seign}) + \pi'(b_{seign})b_{seign} = \pi(b_{seign})(1 - \sigma(b_{seign})) = 0$. An implication of **B2** is then that $\sigma(b) < 1$ for $b < b_{seign}$ and $\sigma(b) > 1$ for $b > b_{seign}$. Dispensing from **B3** will allow use to obtain a more general characterization of the cases implied by **B0–B2**.

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

Proposition 9. *Let Assumptions **B0–B2** hold. There exists a unique $b^* \in (\underline{b}, b_{bliss}]$ such that, for any initial point $b_0 < b_{bliss}$, the optimal level of public debt converges monotonically to b^*. Furthermore, $b^* < b_{bliss}$ if $g > \hat{g}$ and $b^* = b_{bliss}$ if $g < \hat{g}$, for some $\hat{g}$.*

This result identifies b^* as the steady state to which the economy converges from *any* initial point $b_0 < b_{bliss}$. It also relates b^* to the satiation point b_{bliss}. In particular, it shows that b^* is strictly lower than b_{bliss} if and only if g is high enough. Theorem 1 in the main text then follows directly from noting that Property **B3** in the main text is the same as $g > \hat{g}$ here. The rest of the section is dedicated to proving Proposition 9 in multiple steps, developing additional insights on the way. We start by noting that Property **B0–B2** imply the following structure for the function γ, which is instrumental for the subsequent analysis.

Lemma 2. *Let Assumptions **B0–B2** hold. The domain of γ is $\Delta = [\underline{b}, b_{seign}) \cup (b_{seign}, b_{bliss})$, where $b_{seign} \equiv \arg\max \pi(b)b$. For $b \in [\underline{b}, b_{seign})$, γ is negatively valued and decreasing. For $b \in (b_{seign}, b_{bliss})$, γ is positively valued and decreasing. Finally, $\gamma(b) \to -\infty$ as $b \to b_{seign}$ from below and $\gamma(b) \to +\infty$ as $b \to b_{seign}$ from above.*

Proof. Recall that $b_{seign} = \arg\max_b \pi(b)b$, 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: $\pi(b)(1-\sigma(b)) \gtrless 0$ for $b \lessgtr b_{seign}$. From the definition of γ and the assumption $V'(b) \propto \pi(b)$, we have

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

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

Recall that the graph of γ identifies the $\dot{\lambda} = 0$ locus in the region to the left of the satiation point, whereas the $\dot{b} = 0$ locus is given by the graph of ψ. By Lemma 1, ψ is positively valued and strictly increasing. Together with Lemma 2, this means that γ and ψ can intersect at most once. In particular, letting $\gamma_{bliss} \equiv \lim_{b\uparrow b_{bliss}} \gamma(b)$ and $\psi_{bliss} \equiv \psi(b_{bliss})$,[20] we have the following property.

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

Lemma 3. *Let Assumptions **B0–B2** hold. If $\gamma_{bliss} > \psi_{bliss}$, then γ and ψ never intersect. If instead $\gamma_{bliss} < \psi_{bliss}$, then γ and ψ intersect exactly once, and this intersection occurs at $b = b^*$, for some $b^* \in (b_{seign}, b_{bliss})$.*

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

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

Let us first consider Figure 7. The phase diagram is split in three regions: the region L, for $b < b_{seign}$; the region M, for $b \in (b_{seign}, b_{bliss})$; and the region H, for $b > b_{bliss}$. 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 8. What distinguishes the two figures is whether γ and ψ admit an intersection within region M. In Figure 7, they do not. This is because we have imposed $\gamma_{bliss} > \psi_{bliss}$, which together with the monotonicity of γ and ψ guarantees that γ lies above ψ throughout region M.

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

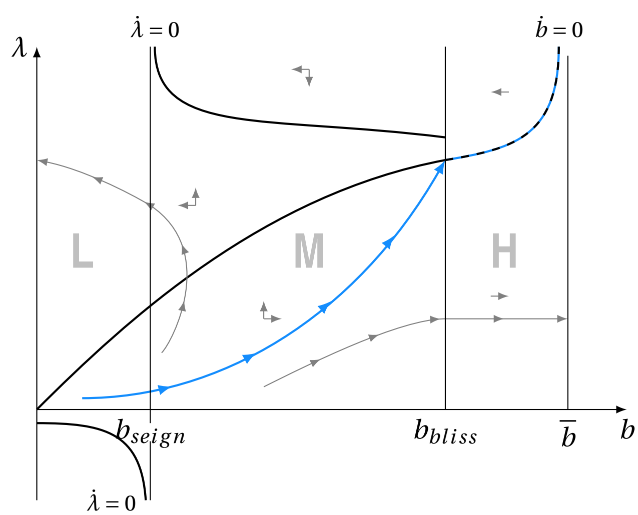

What do these properties imply for the solution to the planner's problem? Since γ and ψ never intersect for the case depicted in Figure 7, the ODE system (20)-(22) admits no steady state to the left of the satiation point b_{bliss} (regions L and M). By contrast, there is a continuum of such steady-state 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 $b = b_{bliss}$ and $\lambda = \psi_{bliss} \equiv \psi(b_{bliss})$; the latter corresponds to the level of taxes that balances the budget when the economy rests at the satiation point.

For any $b_0 < b_{bliss}$, there exists a unique value of the costate, $\lambda_0 < \psi(b_0)$, such as the following is true: if the economy starts from (b_0, λ_0) and thereafter follows the dynamics dictated by (20)-(22), then, and only then, the economy converges asymptotically to $(b_{bliss}, \lambda_{bliss})$. In other words, there is a unique path that satisfies the planner's Euler condition and the budget constraint at all dates, and that eventually leads to satiation. This path is indicated with blue color in the figure.[21]

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

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

Proposition 10. *Let Assumptions **B0 –B2** hold and suppose $\psi_{bliss} < \gamma_{bliss}$. If $b_0 < b_{bliss}$, debt converges monotonically to b_{bliss} and taxes exhibit a positive drift along the transition. If instead $b_0 \geq b_{bliss}$, debt stays constant at b_0 for ever, and tax smoothing applies.*

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

$$\dot{b} = \rho b - S(\lambda)$$
$$\dot{\lambda} = 0$$

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.

[21] One 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 $\lambda < 0$.

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

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

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

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

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

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

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

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

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

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

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

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

Let us now consider Figure 8. In this case, γ and ψ intersect exactly once, at $b = b^* \in (b_{seign}, b_{bliss})$. Let $\lambda^* \equiv \psi(b^*)$ denote the shadow cost of taxation associated with balancing the budget when $b = b^*$. By construction, the pair (b^*, λ^*) identifies the unique steady state of the ODE system (20)-(22) to the left of the satiation point (i.e., within regions L and M). As is clear from the figure, this steady state is saddle-path stable. In particular, for any $b_0 < b_{bliss}$, we can find a continuous path that satisfies conditions (20)-(22) and that asymptotically converges to (b^*, λ^*). Exactly the same arguments as in Figure 7 guarantee that this path is the unique candidate for optimality, and hence also the optimal path, as long as $b_0 < b_{bliss}$.

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

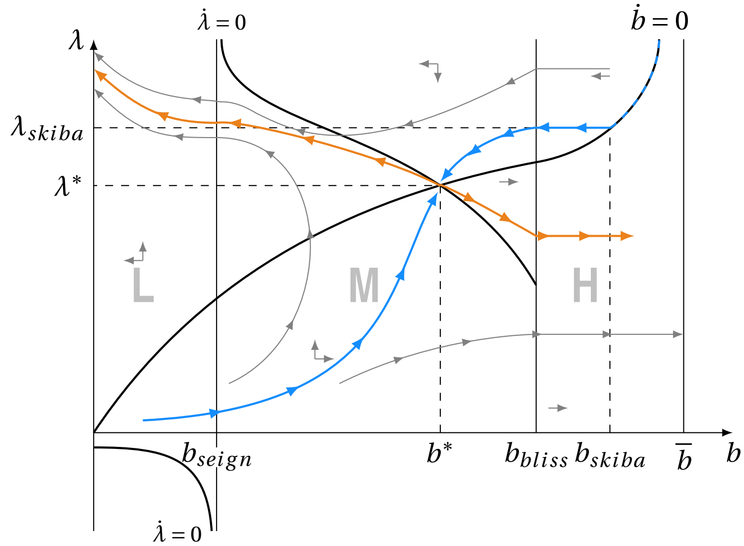

A crucial difference from the case in Figure 7 is that the economy now converges to a steady state characterized by a debt level that is strictly lower than the satiation level: Proposition 9 applies with $b^* < b_{bliss}$. Consequently, the sign of the drift in debt and taxes now depends on the initial position: if $b_0 < b^*$, then debt and taxes increase monotonically over time, whereas the converse is true if $b_0 \in (b^*, b_{bliss})$.

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 $b_0 > b_{bliss}$. This is no longer true. Instead, as it is evident in the figure, for any $b_0 \in [b_{bliss}, b_{skiba})$, there is an additional candidate for optimality: the path indicated with blue color in the figure.

This path lets b fall over time, crossing b_{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 (20)-(22) 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 $\mathscr{P}(b_0)$ 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 $\mathscr{P}(b_0)$. More generally, we can reduce the planner's problem to that of choosing a path $\mathscr{P}(b_0)$. Next, note that any path in $\mathscr{P}(b_0)$ 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 $\mathscr{P}(b_0)$ is therefore equivalent to choosing an initial value λ_0 in $\Lambda(b_0)$. The following result is helpful for evaluating the welfare associated with any candidate path.

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

Proof. See Brock and Dechert (1983). □

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

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

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

Consider $b_0 = b_{bliss}$. If we initiate the ODE system with a starting value $\lambda(0)$ slightly above $\psi_{bliss} = \psi(b_{bliss})$, which means that we run a sufficiently small enough surplus, then the resulting path for b never reaches b^*. By contrast, if we start with $\lambda(0)$ far above $\psi(b_{bliss})$, debt falls below b^* in finite time. Finally, note the path of b induced by the ODE system is continuous and monotonic in $\lambda(0)$. It follows that there exists a critical value $\lambda_{skiba} \in (\psi_{bliss}, \infty)$ such that, if we start with $\lambda(0) = \lambda_{skiba}$, then and only then the economy converges asymptotically to b^*.

By continuity, this kind of path also exists for b_0 above but close enough to b_{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 $b_{skiba} \in (b_{bliss}, \bar{b})$ as the level of debt that balances the budget when taxes are set at the level corresponding to λ_{skiba}; that is, $b_{skiba} \equiv \psi^{-1}(\lambda_{skiba})$. Note that ψ is continuous and monotone, $\lambda_{bliss} > \psi(b_{bliss})$, and $\lim_{b \to \bar{b}} \psi(b) = \infty$; 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 $b_0 < b_{skiba}$, as illustrated in the figure.

We thus have the following complement to Proposition 10.

Proposition 11. *Let Assumptions **B0–B2** hold and suppose $\psi_{bliss} > \gamma_{bliss}$. There exist unique points $b^* \in (b_{seign}, b_{bliss})$ and $b_{skiba} \in (b_{bliss}, \bar{b})$ such as the optimal debt level converges monotonically to b^* if $b_0 < b_{skiba}$, whereas it stays constant at b_0 for ever if $b_0 \geq b_{skiba}$. Optimal taxes exhibit a positive drift as long as $b \in (b_{seign}, b^*)$, a negative drift as long as $b \in (b^*, b_{bliss})$, and are smoothed as long as $b > b_{bliss}$.*

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

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

where $\varpi \equiv \pi(b)/V'(b)$ 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 $b^* < b_{bliss}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ and $V'(b^*) > 0$. Finally, from Lemma 2, we know that $\gamma'(b) < 0$ for $b \in (b_{seign}, b_{bliss}]$. 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. □

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

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

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

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

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

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

B.3 Beyond the Benchmark: 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 10 and 11 continue to hold.

What if the aforementioned properties do not hold, as it may be the case for certain micro-foundations? There is a plethora of possibilities. To make progress, we will continue for a moment to assume that $\pi(b)b$ is single-peaked, which preserves the tripartite structure of the phase diagram, but will let $\gamma(b)$ be non-monotone over region M.[22] In this case, the graphs of γ and ψ may intersect multiple times. Clearly, any such intersection identifies a steady-state point of the ODE system. *What are the local dynamics around each of these points? Starting from a given initial* b_0*, how many paths are candidates for optimality? And what are the properties of the optimal path?*

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

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

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

(i) saddle-path stable if $\gamma'(b^*) < \psi'(b^*)$*;*

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

(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

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

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

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

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

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

Given that $b^* < b_{bliss}$, $\sigma(b^*) < 1$, $\gamma(b^*) > 0$ and $V'(b^*) > 0$. Therefore, the position of $\gamma'(b^*)/\psi'(b^*)$ with respect to 1 determines the sign of the determinant, and hence the position of the two eigenvalues around 0. Note that a steady state only exists in regions where $\sigma(b^*) > 1$ and hence $\gamma(b^*) > 0$. When $\gamma'(b^*) < \psi'(b^*)$, $\det(\mathbf{J}) < 0$ and hence the two eigenvalues are distributed around 0. Therefore, a saddle path exists (recall that $\mathrm{Tr}(\mathbf{J}) = \rho > 0$), 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:

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

The two roots are complex if the discriminant is negative

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

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

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

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

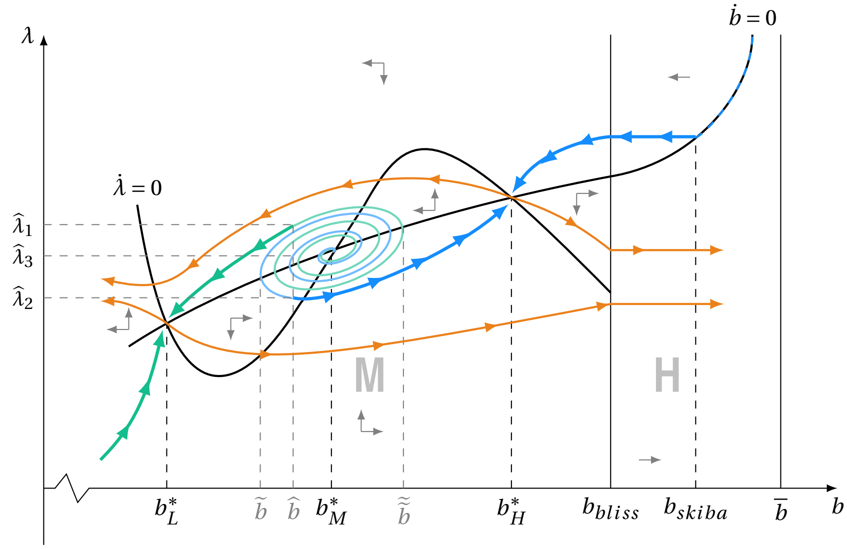

Notwithstanding these restrictions on the local dynamics, there remain three distinct possibilities with regard to the *global* dynamics. Figure 9 considers one of these possibilities.

In Figure 9, we have imposed the following property on the global dynamics: both the stable arm that leads to b_L^* from above and the one that leads to b_H^* from below cycle back to b_M^*. It follows that there exist values $\tilde{b}$ and $\tilde{\tilde{b}}$, as indicated in the figure, such that the following is true within region M. Whenever $b_0 < \tilde{b}$, $\Lambda(b_0)$ is a singleton and the unique candidate for optimality is the saddle path that leads to b_L^*. Whenever $b > \tilde{\tilde{b}}$, $\Lambda(b_0)$ is again a singleton, but now the unique candidate is the saddle path that leads to b_H^*. Finally, whenever $b_0 \in [\tilde{b}, \tilde{\tilde{b}}]$, there are multiple paths that are candidates for optimality. For instance, if we take $b_0 = \hat{b}$ as indicated in the figure, one candidate is obtained by setting $\lambda_0 = \hat{\lambda}_1 \equiv \max \Lambda(b_0)$ and letting debt decrease monotonically towards b_L^*; another candidate is obtained by setting $\lambda_0 = \hat{\lambda}_2 = \min \Lambda(b_0)$ and letting debt increase monotonically towards b_H^*; and yet another candidate is obtained by setting $\lambda_0 = \hat{\lambda}_3$ and letting debt to cycle twice around $\hat{b}$ before eventually converging to b_H^*. The closer b_0 is to b_M^*, the larger the number of candidates; when b_0 is exactly b_M^*, there is actually a countable infinity of candidates.

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

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

Combining these observations, we obtain the following properties. When $b_0 = \tilde{b}$, the path that leads to b_L^* features $\mathscr{H}_\lambda = \dot{b} < 0$, whereas the path that leads to b_H^* features $\mathscr{H}_\lambda = \dot{b} = 0$. By the convexity of $\mathscr{H}$, the latter path is dominated. Conversely, when $b_0 = \tilde{\tilde{b}}$, it is the former path that now features $\mathscr{H}_\lambda = \dot{b} = 0$ and that is therefore dominated. By continuity,[23] 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 9 for an illustration: the bold segments of the two stable arms indicate the optimal selection among the two candidate paths.[24]

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

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

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

[23] Here, 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).

[24] In 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.

optimal to converge to b_L^*.

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

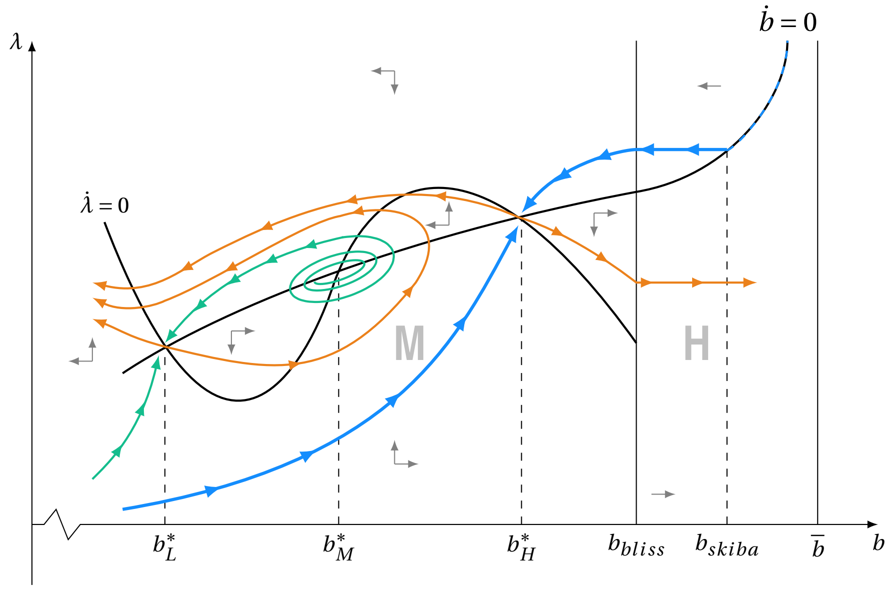

Let us fill in the details, starting with Figure 10. Unlike Figure 9, the stable arm corresponding to the highest steady state no longer cycles back to b_M^*; instead, it extends past b_L^*. This has the following important implication. If we consider $b_0 = b_L^*$, then there are two candidate optimal plans, namely the plan of staying put at b_L^* and the plan that leads to b_H^*. The former plan is dominated because it features $\mathscr{H}_\lambda = \dot{b} = 0$, whereas the latter features $\mathscr{H}_\lambda = \dot{b} > 0$. By continuity, the saddle path that leads to b_L^* is dominated also for any b_0 in an open neighborhood of b_L^*. But then the path leading to b_L^* can *never* be optimal: if the economy were to follow this path starting from any initial point b_0, the economy would enter the aforementioned neighborhood in finite time; at that point, switching paths would increase welfare, which contradicts the optimality of the original path. We conclude that, contrary to what happens in Figure 9, the path that leads to b_H^* in Figure 10 is now the optimal path for all $b_0 < b_{skiba}$.

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

These findings illustrate the following more general points and qualify some of the properties of the benchmark model. To the extent that the ODE system admits multiple steady states below b_{bliss}, any such point represents a point of indifference between the desire to depress the interest rate on public

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

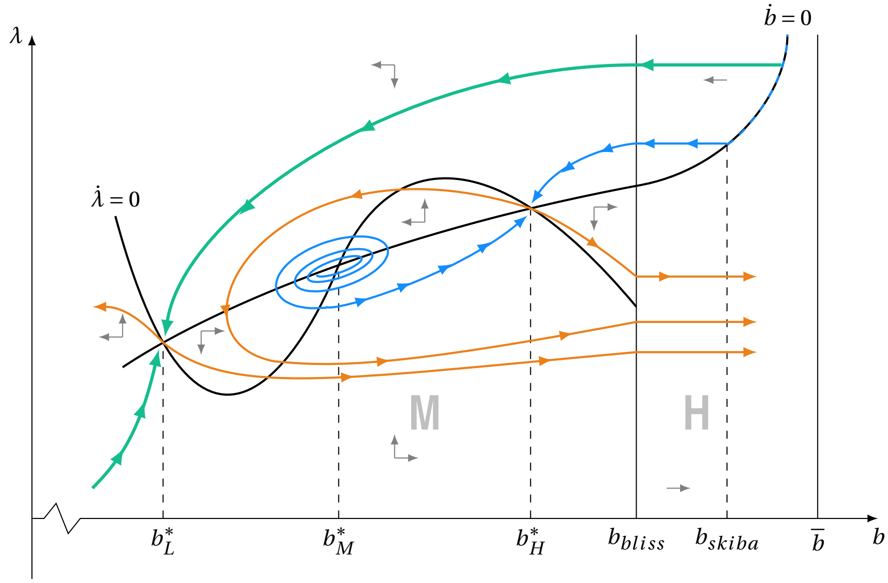

debt and the desire to improve liquidity and efficiency; this is our earlier observation that $\Omega_b = 0$ at any such point. Furthermore, to the extent that such a point is locally saddle-path stable, it is optimal to converge to it over time if the economy starts in a small enough neighborhood of 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 $\{b_1, b_2, b_3, \ldots, b_N\}$, with $\underline{b} < b_1 < b_2 < \ldots < b_N < b_{bliss}$, where N is an arbitrary finite number. First, note that $\sigma(b)$ crosses 1 whenever b crosses any of these points. Next, note that the last point, namely b_N, is necessarily a local maximum, because after that point $\pi(b)b$ falls to zero as b approaches b_{bliss}. It follows that $\sigma(b)$ is higher than 1 when $b \in (b_N, b_{bliss})$, lower than 1 when $b \in (b_{N-1}, b_N)$, higher than 1 when $b \in (b_{N-2}, b_{N-1})$, and so on. By the same token, γ is positively valued $b \in (b_N, b_{bliss})$, negatively valued than 1 when $b \in (b_{N-1}, b_N)$, positively valued when $b \in (b_{N-2}, b_{N-1})$, and so on.

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

Figure 12: Multiple Regions

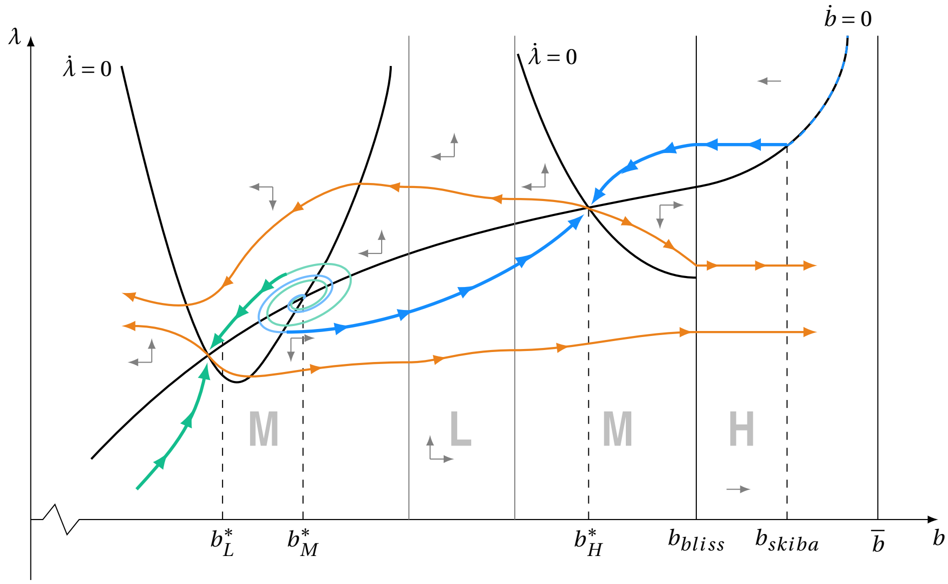

Notwithstanding all the complexity, we can thus establish the following result, which offers a qualified generalization of Proposition 11 in our benchmark.

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

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

Let $b^\ddagger \in (0, b_{bliss})$ be the *last* local maximum of $\pi(b)b$.[25] 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 $\gamma(b^*) = \psi(b^*)$

[25] Because $\pi(b)b$ 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}.

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

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.[26]

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

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

As already explained, the condition $\psi_{bliss} > \gamma_{bliss}$ has a simple interpretation: it means that, in the neighborhood of b_{bliss}, the shadow cost of taxation is sufficiently high so that the marginal value of depressing the interest rate on public debt outweighs the marginal cost of the financial distortion. Consistent with this interpretation, it is straightforward to show this case obtains when the level of government spending is sufficiently high.[27]

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

Proof. Note that ψ_{bliss} is continuous and increasing in g as long as $g < g_{max}$ and diverges to $+\infty$ as $g \to g_{max}$. 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, $\psi_{bliss} = 0$

[26] This 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.

[27] In 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.

if and only if $g = -\rho b_{bliss} < 0$. Finally, note that γ_{bliss} is *(i)* invariant to g; *(ii)* positive for the reasons offered above; and *(iii)* finite by assumption. It then follows that there exists a threshold $\widehat{g}$, necessarily less than g_{max} and possibly negative, such that $\psi_{bliss} > \gamma_{bliss}$ if and only if $g > \widehat{g}$. $\square$

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

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

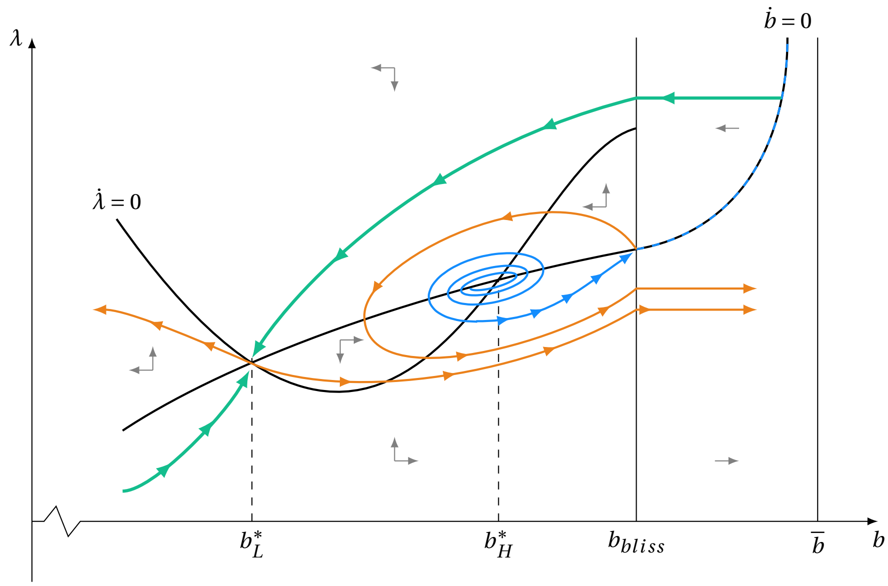

We illustrate this in Figure 13. As in our benchmark (see Figure 7 in particular), letting $\gamma_{bliss} > \psi_{bliss}$ guarantees the local existence of a candidate path that leads to satiation: for some $\epsilon > 0$ and all $b_0 \in (b_0 - \epsilon, b_{bliss})$, there exists a path that satisfies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

B.5 Beyond the Benchmark: Relaxing Auxiliary Assumption B0

In the preceeding 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 $b > b_{bliss}$. As mentioned in the main text, this amounts to changing the interpretation of region H (i.e., the region where $b > b_{bliss}$) from "harmless satiation" to "harmful excess."

In this case, the phase diagram takes the form illustrated in Figure 14. Relative to our benchmark, there are two changes: in region M, the $\dot{\lambda} = 0$ locus presents an asymptote at $b = b_{bliss}$; 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 $b_0 \in (0, \overline{b})$.

Figure 14: Relaxing Assumption **B0**

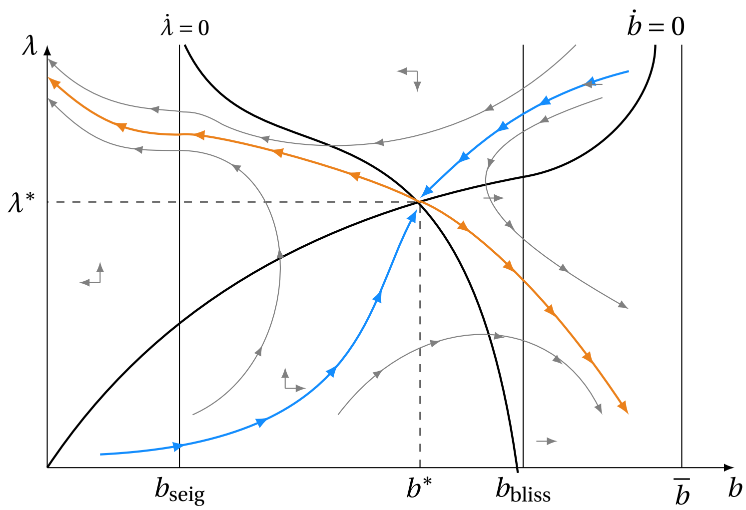

What if we relax *both* B0 and B1-B3? Clearly, this is a hybrid of the present and the previous scenarios: relaxing B1-B3 allows multiple intererior steady states below b_{bliss}, while relaxing B0 guarantees that $\dot{\lambda} < 0$ for $b > b_{bliss}$ and hence that the economy converges to a steady state below b_{bliss} for any intial 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.

B.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

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

be the set of the points at which γ intersects ψ from above. As shown in Proposition 12, these points identify the saddle-path stable steady states of the ODE system.[28] Depending on primitives, $B^{\#}$ may be empty, or may contain an arbitrary number of elements.[29] Regardless of this, we have the following result.

Theorem 2. *In every economy, there exists a threshold $b_{skiba} \in [\underline{b}, \overline{b}]$ 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 is a selection from the correspondence Λ (which was defined in the context of Lemma 4). To illustrate, consider Figure 9. In this example, the aforementioned correspondence is given by the combination of three segments: the thick green line on the left of $\hat{b}$, plus the solid blue line between $\hat{b}$ and b_{skiba}, plus the segment of the graph of the $\dot{b} = 0$ locus that rests on the right of b_{skiba}. As it is evident in this example, the correspondence λ^* is single-valued and continuous for all b_0 other than $\hat{b}$; the discontinuity at $\hat{b}$ reflects a switch in the optimal selection among different candidate paths. Moving beyond this specific example, the policy rule for the co-state can feature multiple such discontinuities. Any such discontinuity, however, has 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 $\lim_{b\uparrow\hat{b}} \Lambda^{opt}(b) > \lim_{b\downarrow\hat{b}} \Lambda^{opt}(b)$,[30] In other words, as we move from higher to lower levels of debt, the

[28] In 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 $B^{\#}$.

[29] We 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.

[30] At 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.

co-state can only jump upwards, which means that the rate of taxation and the level of government surpluses must also jump upwards. It then follows that lower initial conditions are necessarily associated with convergence to lower steady states, which in turn is the key to the result.

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

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

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

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

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 8 and 9, B^* is a strict subset of $B^{\#}$ in Figures 10 and 11.

Finally, it is generally possible that $B^* = \emptyset$, meaning that satiation obtains in the long run regardless of initial conditions. But as already explained, this scenario can be ruled out by assumping at least on 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, sufficent 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 nom-empty follows from the examples we have already provided. Finally, the claimed sufficiency follows from the previous analysis as well. □

B.7 Local Dynamics and Local Comparative Statics

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

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

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

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

we know from proposition 12, that the eigenvalue associated with the stable arm is given by $\mu = \frac{\rho-\sqrt{\Delta}}{2}$ with $\Delta > 0$ (see proof of Proposition 12). It is then straightforward to obtain the eigenvector $\mathbf{v} = (v_1, v2)$ 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. □

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

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

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

C Additional Results

C.1 Allowing for state-contingent debt

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

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

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

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

We illustrate this property in Figure 15 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 15: Response to a war shock, with state-contingent debt

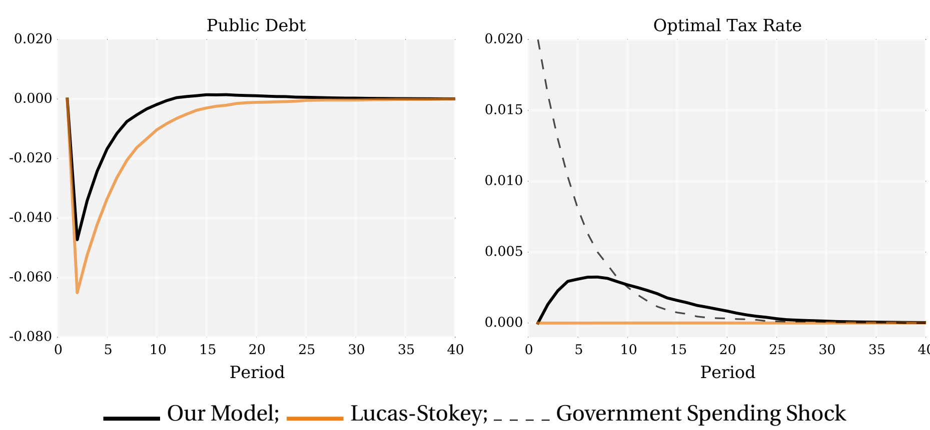

To sum up, once public debt is non-neutral for the reasons we have accommodated in this paper, the difference between Barro/AMSS and Lucas-Stokey is attenuated: the qualitative response of the optimal tax and the optimal allocation is the same whether the government has access to state-contingent debt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

C.2 On the Friedman Rule

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

Then, the government's total liabilities are given by $b = m - n$, where m is the stock of government bonds and n is the quantity of corporate bonds held by the government. The budget constraint is given by

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

or equivalently

$$\dot{b} = \rho b - \pi(m)m + g - s, \tag{24}$$

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 (24) over time to obtain the familiar intertemporal budget constraint:

$$b_0 + G = \int_0^{+\infty} e^{-\rho t}[\pi(m)m + S(\tau)]\mathrm{d}t. \tag{25}$$

where $G \equiv \int_0^{+\infty} e^{-\rho t} g \mathrm{d}t$ 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,

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

subject to the single integral constraint in (25). Let λ denote the Lagrange multiplier on the intertemporal budget. It is then immediate that the optimal supply of liquidity is given by

$$m^{\star} = \arg\max_{m} \Omega(m, \lambda), \tag{26}$$

where $\Omega(m,\lambda) \equiv V(m) + \lambda\pi(m)m$ measures "liquidity plus seigniorage". Depending on primitives, $m^{\star}$ 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 rains supreme; and the optimal fiscal policy is determined in exactly the same fashion as in Barro (1979).

C.3 Relation to Aiyagari and McGrattan (1998)

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

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

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

$$\hat{b} = \arg\max_{b} \left\{ V(b) - \hat{\lambda} R(b) b \right\}, \tag{27}$$

This underscores how the Aiyagari-McGrattan approach treats the interest payments on public debt, $R(b)b$, as a cost. But as first highlight in Section 4.4, the component ρb of these interest-rate payments is *not* a cost. Accordingly, the *truly* optimal steady state satisfies

$$b^{\star} = \arg\max_{b} \left\{ V(b) + \lambda^{\star}\pi(b) b \right\}, \tag{28}$$

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

In summary, the solution strategy Aiyagari and McGrattan (1998) not only abstracts from transitional dynamics (or, relatedly, the optimal response to shocks) but also offers a distorted perspective on optimal long-run quantity of public debt. At the same time, Aiyagari and McGrattan (1998) allow for an interesting economic effect that our main analysis abstract from: that public debt may crowds out capital accumulation by offering a substitute form of buffer stock. We explain why this possibility, or even the opposite one, does not fundamentally change the policy problem in Section 4.4 in the main text.

Finally, note that discussion here presumes, like the analysis in Section 4.3, that the Auxiliary Assumptions hold. As explained in Appendix B.3, the economy may feature multiple steady states when these assumptions do not hold. In these circumstances, the Aiyagari-McGrattan approach will never detect this multiplicity, for it is the (generically) unique solution to a static optimization problem.

D Main Text Proofs

Proof of Proposition 1. The results derive directly from the concavity of the utility function, $u(x)$, and the assumption $u''(x) < 0 < u'''(x)$.

Q.E.D. ■

Proof of Proposition 2. Follows directly from the main text.

Q.E.D. ■

Proof of Proposition 3. The household solves the problem

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[c_{i,t} + \theta_{i,t} \log(x_{i,t}) - v(h_{i,t}) \right] \tag{29}$$

subject to

$$c_{i,t} + p_{i,t} x_{i,t} + q_t b_{i,t+1} = b_{i,t} + (1-\tau_t) w_t h_{i,t} + p_t e_i \tag{30}$$

$$p_t (x_{i,t} - e_i) \leqslant \xi + b_{i,t} \tag{31}$$

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 $\mathbb{E}[\theta] = 1 + \varphi\vartheta$ 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 $e_L = 1/(1-\phi)$, such that the total endowment, $\overline{e} = \varphi e_H + (1-\varphi)e_L = 1$. The rest of the notation is identical to that used in Section 3.2.

In equilibrium, the borrowing constraint (31) 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:

$$\frac{1+\epsilon}{x_{Ht}} = p_t(1+\mu_t) \tag{32}$$

$$\frac{1}{x_{Lt}} = p_t \tag{33}$$

$$p_t x_{Ht} \leq \xi + b_t \tag{34}$$

$$\mu_t \geq 0 \tag{35}$$

$$\mu_t(\xi + b_t - p_t x_{Ht}) = 0 \tag{36}$$

$$\varphi x_{Ht} + (1-\varphi) x_{Lt} = 1 \tag{37}$$

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

$$\pi_t \equiv q_t - \beta = \beta\varphi\mu_{t+1} \geq 0 \tag{38}$$

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

$$\frac{1+\epsilon}{x_{Ht}^*} = \frac{1}{x_L^*} \tag{39}$$

$$\varphi x_H^* + (1-\varphi) x_L^* = 1 \tag{40}$$

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

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

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

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

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

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

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:

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

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

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

We therefore reach the following result:

Lemma 6. *Suppose $\xi < 1+\epsilon$, we have*

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

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

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

(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 $b < b_{bliss}$.*

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

Proof. The properties of π and V with respect to b, and hence *(i)*, follow directly from their closed-form characterization. Part *(ii)* follows from the fact that $b_{bliss} = 1+\epsilon-\xi$ along with the fact that, for any b_{bliss} and any $b < b_{bliss}$, 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 $b < b_{bliss}$. □

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

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

and can, after some algebraic manipulation, be rewritten as

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

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 $\pi(b) > V'(b)$.

Beside the results reported in the proposition we also consider the problem introduced in Section 4.1, which as shown in Section 4.4 also characterizes the optimal steady state. In particular, consider the following two objects:

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

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

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

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

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

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

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

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

which establishes the concavity of $g(b) \equiv \pi(b)b$. Next, note that $g'(0) = \beta\varphi\frac{1+\epsilon-\xi}{\xi} > 0$ and $g'(b_{bliss}) = \pi'(b_{bliss})b_{bliss} = -\beta\varphi\frac{1+\epsilon-\xi}{1+\epsilon} < 0$. It follows that b_{seign} is the unique solution to $g'(b) = 0$ and is strictly between 0 and b_{bliss}.

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

This result echoes the properties we establish in Section 4.1 for the steady state. But because we now have a simple closed-form characterization of Ω, we can go a step further to study the comparative statics of b^* with respect to the underlying primitives. Those are reported in the next proposition

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

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

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

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

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

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

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

Although b^* can be decreasing in ξ, which means that more private collateral can crowd out the government-provided collateral, there is no complete crowding out: an increase in ξ always increases total collateral, $b^*+\xi$.[31] It then also follows that, at the optimal quantity of public debt, more private

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

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

collateral depresses the liquidity premium ($\frac{\partial \pi(b^*)}{\partial \xi} < 0$), whereas the converse is true with an aggravation of liquidity needs ($\frac{\partial \pi(b^*)}{\partial \epsilon} > 0$).

To conclude, these findings complement the intuitions developed in Section 4.4. Strictly speaking, they do not apply to the steady state of the infinite-horizon model, because they treat λ as exogenous. But we can use the government budget evaluated at the steady state to obtain λ as an increasing function of b, an increasing function of g, and a decreasing function of π (and thereby a decreasing function of θ and an increasing function of ξ). We can then readily translate the result to the steady-state level of debt, modulo the replacement of λ with g. That is, the value of fiscal space is re-parameterized by g, but the comparative statics with respect to ϑ and ξ go through.

Q.E.D. ■

Proof of Theorem 1. Assumption **B3** corresponds to the case $g > \widehat{g}$ ofLemma 5, such that $\Psi_{bliss} > \gamma_{bliss}$.Then Proposition 11 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 $s^* = g + rb^* - \pi(b^*)b^*$. Note first that

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

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

Q.E.D. ■

Proof of Proposition 5 to 7. Follow directly from the discussions in the main text.

Q.E.D. ■

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

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

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

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

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

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

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

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

$$\mathbf{n}(x,p,w) = \int n(x,p,w;\theta)\varphi(\theta)\mathrm{d}\theta \tag{41}$$

$$\mathbf{k}(x,p,w) = \int k(x,p,w;\theta)\varphi(\theta)\mathrm{d}\theta \tag{42}$$

The problem of the household is

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

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 $\int a_{it}\mathrm{d}i = b_t$, let $\kappa_t \equiv \int \kappa_{it}\mathrm{d}i$ denote aggregate investment, and define

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

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

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

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

$$\nu'(h_t) = (1-\tau_t)w_t \tag{43}$$

$$q_t = \beta(1 + \Omega_x(x_{t+1},p_{t+1},w_{t+1})) \tag{44}$$

$$1 = \beta\left(1 + \Omega_x(x_{t+1},p_{t+1},w_{t+1})p_{t+1}\right) \tag{45}$$

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

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

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

$$h_t = H(b_t, \tau_t) \text{ and } k_t = K(b_t, \tau_t) \tag{46}$$

so that

$$w_t = W(b_t, \tau_t) \text{ and } p_t = P(b_t, \tau_t) \tag{47}$$

Likewise, using the resource constraint, we get

$$c_t = \theta f(k_t, n_t) + (1-\delta) k_t - k_{t+1} - g = \widetilde{C}(b_t, \tau_t) - k_{t+1} \tag{48}$$

Using (46) and (48) in the welfare function, we get

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

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 (44), 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

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

Q.E.D. ■