

Public Debt as Private Liquidity: Optimal Policy*

George-Marios Angeletos
MIT

Fabrice Collard
University of Bern

Harris Dellas
University of Bern

September 25, 2016

Abstract

We study the Ramsey policy problem in an economy in which public debt contributes to the supply of assets that private agents can use as buffer stock and collateral, or as a vehicle of liquidity. Issuing more debt eases the underlying financial friction. This raises welfare by improving the allocation of resources; but it also tightens the government budget by raising the interest rate on public debt. In contrast to the literature on the Friedman rule, the government's supply of liquidity becomes intertwined with its debt policy. In contrast to the standard Ramsey paradigm, a departure from tax smoothing becomes desirable. Novel insights emerge about the optimal long-run quantity of public debt; the optimal policy response to shocks; and the sense in which a financial crisis presents the government with an opportunity for cheap borrowing.

*We are grateful to Behzad Diba for his collaboration on an earlier project that paved the way to the present paper; to Pedro Teles for illuminating discussions; and to various seminar participants for their feedback. Angeletos also thanks the University of Bern, Study Center Gerzensee, and the Swiss Finance Institute for their hospitality.

1 Introduction

In the presence of financial frictions, public debt can influence market efficiency and asset prices by contributing to the supply of assets that can be used as buffer stock or collateral. This insight, formulated in various related forms by Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998), seems particularly relevant in the aftermath of the Great Recession.

For instance, this insight suggests that the government may be able to mitigate a financial crisis not only by changing the composition of its balance sheet—a popular proposition in the quantitative-easing literature¹—but also by raising its net indebtedness. As another example, consider the argument made by Eggertsson and Krugman (2012) and Guerrieri and Lorenzoni (2011) that the recent financial crisis pushed the economy towards the zero lower bound (ZLB) by reducing the natural rate of interest. In this context, the issuance of public debt could, not only raise welfare by easing the underlying borrowing constraints, but also provide breathing space to monetary policy by raising the natural rate of interest and relaxing the ZLB constraint.

In spite of these considerations, the literature has yet to develop a benchmark for how optimal fiscal policy is determined in economies in which public debt is non-neutral, because it influences the virulence of financial frictions, but does not generate a free lunch for the government because taxation is distortionary. The contribution of this paper is to provide such a benchmark.

We accomplish this in three steps. First, we develop a stylized micro-founded model that accommodates the aforementioned function of public debt within the Ramsey policy paradigm (Barro, 1979, Lucas and Stokey, 1983, and Aiyagari et al., 2002). Second, we obtain a convenient reduced-form representation of the planner’s problem. Finally, we characterize analytically the solution to a class of reduced-form problems that nests the one obtained from our model.

This three-step approach illuminates the key policy trade offs. It also clarifies how the policy problem we are interested in relates to, but also differentiates from, those studied in the literature on the Friedman rule (Chari, Christiano and Kehoe, 1996, Correia and Teles, 1999). Last but not least, it reveals a formal link between our problem and the non-convex optimal-control problems studied in Skiba (1978), Brock and Dechert (1983), and Buera (2008), a link that is instrumental for the characterization of the optimal policy.

Optimal policy is shown to balance three objectives: smoothing taxes; easing the underlying financial friction; and lowering the government’s cost of borrowing. In our preferred scenario, the economy converges to a steady state in which the planner preserves, or even exacerbates, the friction in order to depress the interest rate on public debt. The optimal dynamics around such a steady state feature mean-reversion, contrasting the unit-root property exhibited in Barro (1979) and Aiyagari et al. (2002). Finally, the optimal response to a financial crisis may be driven primarily by fiscal considerations rather than by the need to alleviate the financial friction.

¹See, inter alia., Cúrdia and Woodford (2011), Gertler and Karadi (2011) and Gertler and Kiyotaki (2010)

Microfoundations. Our framework shares the same neoclassical backbone as the models in Barro and AMSS (henceforth shortcuts for Barro (1979) and Aiyagari et al. (2002), respectively). Agents are infinitely lived; markets are competitive; prices are flexible; and the government can collect taxes only by distorting the supply of labor. Unlike these models, however, ours includes a certain friction in how markets allocate resources for any given fiscal policy.

This friction is similar—in its essence although not in its exact form—to those found in Kiyotaki and Moore (1997) and Holmström and Tirole (1998). Private agents are subject to uninsurable idiosyncratic shocks, which are interpreted as liquidity shocks. These shocks can be met by reallocating a certain good from one agent to another. Reallocation requires borrowing and borrowing requires collateral. The risk-free debt issued by one private agent can serve as collateral, or buffer stock, for another agent. Nonetheless, collateral is in short supply at the aggregate level, because private agents can pledge only a fraction of their future income. Finally, public debt eases this shortage, in effect, by letting the government’s commitment to raise more taxes in the future substitute for the private agents’ inability to pledge part of their future income in private markets.²

Reduced Form. In the presence of uninsurable idiosyncratic risk, tractability becomes an issue, especially if one wishes to maximize welfare over arbitrary policy paths, as we do here. We cut the Gordian knot in a similar manner as Lagos and Wright (2005) and Guerrieri and Lorenzoni (2009): we make assumptions that prevent the cross-sectional distribution of net worth from becoming a relevant state variable in the planner’s problem. Although this carries costs, it leads to a convenient, and insightful, reduced-form representation of the planner’s problem as an optimal-control problem over the stock of public debt and the level of taxation.

Inspection of this reduced form reveals three key properties, which are likely to be present in richer frameworks as well. First, public debt emerges as a key policy instrument that can be employed to regulate the severity of the financial distortion. Second, this regulation involves two conflicting objectives: issuing more public debt helps attain a more efficient allocation of resources, but it also lowers the premium the market is willing to pay for collateral, or liquidity, thereby raising the government’s cost of borrowing relative to the planner’s discount rate. Third, the planner’s problem is generally non-convex due to the dependence of the aforementioned premium on the problem’s state variable, namely the stock of public debt. The results obtained in this paper depend on these three elementary properties, not on the details of the underlying micro-foundations.

Public Debt as Money. The function of public debt in our framework resembles that of “money” in the literature on the Friedman rule: the welfare gain from easing the financial distortion is akin to the utility of money, while the fiscal gain from paying a lower interest rate on public debt is akin to seigniorage. There is, however, a fundamental difference.

²The same mechanism explains the non-neutrality of public debt in Woodford (1990) and Aiyagari and McGrattan (1998) as well.

In the Friedman-rule literature, the government trades two types of assets: one, called “money”, conveys utility, while the other, called “bonds”, does not. By reshuffling its portfolio of money and bonds, the government can therefore regulate private liquidity without changing its net debt position. This feature helps disentangle the optimal supply of liquidity from optimal fiscal policy. A similar disentangling is also present in the recent literature on quantitative easing (Cúrdia and Woodford, 2011, Gertler and Karadi, 2011, Gertler and Kiyotaki, 2010), even though the form of liquidity provision is different.

By contrast, our framework lets the option provision of liquidity to be intertwined with optimal fiscal policy: at least at the margin, the government can enhance private liquidity only by raising debt burden. This property is key for understanding the nature of our results, their novelty vis-a-vis the pertinent literature, and their likely robustness.³

Optimal policy. When the financial friction is absent, our baseline model reduces to a deterministic version of Barro and AMSS. The planner’s problem is then convex; tax smoothing is optimal; and the long-run level of public debt is indeterminate, in the sense that any sustainable level is consistent with steady state. All these properties are upset once the friction is introduced.

The planner’s problem is non-convex due to the dependence of interest rates on public debt. As a result, the “first-order approach” to solving the problem fails: there generally exist multiple paths that satisfy the relevant first-order and transversality conditions. We nevertheless obtain a sharp characterization of the optimal policy by adapting the optimal-control methods of Skiba (1978), Brock and Dechert (1983), and Buera (2008) to a continuous-time formulation of our problem.

Depending on primitives, either the economy converges to the same steady state from all initial conditions; or there exist multiple steady states, each one associated with a separate basin of attraction.⁴ Either way, the steady state can be understood as a long-run target for the government’s supply of liquidity and the associated financial distortion in the economy. Reaching the desired long-run target necessitates a departure from tax smoothing, the welfare cost of which acts as an adjustment cost. This in turn helps explain, not only why the transition to steady state takes time, but also why the steady state may depend on initial conditions.

Two scenarios are possible in the long run. In the one, the financial distortion vanishes as the economy converges to steady state; in this sense, the analogue of the Friedman rule applies in the long run (asymptotically), although not in the short run (at any finite time). In the other scenario, the analogue of the Friedman rule does not apply in either the short or the long run: the financial distortion remains binding in steady state.

³In our framework, the aforementioned property comes together with an irrelevance of the composition of the government’s portfolio. This is because both properties follow from our assumption that all assets provide identical liquidity services. The latter property, however, is a simplification; only the property mentioned above is essential.

⁴The description of our results in this Introduction presumes that the financial friction is binding in the initial position of the economy, an assumption that is relaxed in the main text.

The first scenario reinforces an important lesson from the Friedman-rule literature: satiating the economy with liquidity *may* be optimal even if this requires a grave increase in the debt burden. We nevertheless find the second scenario of greater interest because it allows us to analyze the optimal policy response to shocks around a non-trivial steady state, that is, a steady state in which the two distortions coexist and the relevant trade offs remain active.

Wars, Recessions, and Crises. Consider a “war”, that is, a transitory positive shock to government spending. In Barro/AMSS, the optimal policy response involves running a deficit during the life of the shock in order to spread the additional tax burden over time. In comparison to this benchmark, we find that the government should run a smaller deficit, because doing so helps attenuate the increase in its cost of borrowing *relative* to the underlying social discount rate. The optimal tax and debt responses therefore exhibit, not only mean reversion, but also smaller volatility.

A similar result applies to a “traditional recession,” by which we mean a recession that does not affect either the social or the private value of liquidity. Things are different for a “financial recession”, which we identify with a shock that aggravates the underlying financial friction.

Interestingly, such a shock can have an ambiguous effect on the planner’s incentives to supply liquidity: while the increase in the social value of liquidity has the obvious effect of encouraging the issuance of public debt, the increase in the private value of liquidity pulls in the opposite direction, because it raises the amount of seigniorage that the planner can extract by squeezing liquidity. In a benchmark scenario in which these two effects are of the same size, perhaps because of the absence of fire-sale or other pecuniary externalities, the planner’s optimal supply of liquidity is invariant to the severity of the financial shock for any given shadow cost of taxation. It follows that the optimal policy response is driven by fiscal considerations and not, as one may have expected, by the apparent need to ease the aggravated financial friction.

What are these considerations? By reducing aggregate income and the tax base, the shock causes an increase in the tax burden. To the extent that the shock is transitory, it also makes it desirable to run a deficit for tax-smoothing reasons. These effects are present in financial and traditional recessions alike. What is different, though, in a financial recessions is that the interest rate on public debt falls *relative* to the planner’s discount rate. This effect, which is akin to an exogenous increase in seigniorage, counters the aforementioned increase in the tax burden and justifies a higher deficit during a financial than during a traditional recession of comparable magnitude.

This result provides a formal basis for the argument, made by commentators such as Paul Krugman and Brad DeLong, that the recent recession called for high deficits, not only because of the need to stimulate aggregate demand, but also because the financial crisis made it “cheap” for the US government to borrow. This argument appears plausible; but it is important to note that it has no place in the standard Ramsey framework because of the absence of a wedge between the interest rate and the planner’s discount rate in that framework.

Robustness. In our baseline model, the financial friction impedes the allocation of a good across consumers; public debt influences welfare and interest rates without affecting aggregate quantities; and private liquidity is tightly connected to public debt because there is no capital or other physical assets. In a variant model we allow the friction to impede the allocation of resources across firms, thus also allowing public debt to influence aggregate productivity and income. Furthermore, we relax the tight relation between public debt and private liquidity by allowing for capital accumulation. It then becomes possible that public debt crowds out capital (Aiyagari and McGrattan, 1998) and, in this sense, that “outside money” crowds out “inside money” (Stein, 2012, Brunnermeier and Sannikov, 2016). Last but not least, a wedge can be present between the private and the social value of liquidity, as in models with fire-sale or other externalities operating through collateral constraints (Shleifer and Vishny, 1992, Lorenzoni, 2008, Dávila, 2015). The modeling strategy we follow in this paper helps clarify that the obtained policy lessons are robust to all these “details”.

Layout. Section 2 introduces the baseline model. Section 3 characterizes its equilibrium for given policy. Section 4 obtains the desired reduced-form representation of the planner’s problem. Section 5 contains the core results. Section 6 provides additional insights and elaborates on certain connections to the literature. Section 7 studies the optimal response to shocks. Section 8 concludes with suggestions for future research. Appendix A presents the variant model, which helps capture financial frictions on the production side. Appendices B and C contain auxiliary results and proofs.

2 The baseline model

Our baseline model is designed with three goals in mind: to accommodate the function of public debt that we are interested in; to keep the analysis tractable; and to build a bridge between the literature on financial frictions and the literature on the Friedman rule.

There is a unit-mass of ex-ante identical households, indexed by $i \in [0, 1]$, and a representative firm. Time is discrete, indexed by $t = 0, 1, 2, \dots$, and each period is split into a “morning” and an “afternoon”. There are two edible goods. The one is the (exogenous) fruit of a tree, which becomes ripe in the morning of each period. The other is the (endogenous) output of the representative firm, which is produced in the afternoon with the labor of the households. Each good has to be consumed in the sub-period in which it is produced, or else it perishes. We refer to the first good as the “morning good” and to the second one, which is also our numeraire, as “the afternoon good”. Utility is assumed to be linear in the afternoon good, helping prevent the wealth distribution from becoming a relevant state variable in the planner’s problem. Finally, liquidity risk takes the form of taste shocks and a financial friction impedes the efficient allocation of the morning good.

The representative firm. The representative firm is competitive and produces the afternoon good using labor. Aggregate output is given by $y_t = Ah_t$, where h_t denotes the labor input and A

denotes the exogenous aggregate productivity (assumed to be time-invariant for the time being). It follows that, in equilibrium, the pre-tax wage is given by $w_t = A$ and all income goes to labor.

The government. The government’s objective is to maximize social welfare (ex-ante utility). Its budget constraint is given by

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

where b_t is the stock of public debt inherited from period $t-1$, g is the exogenous level of government spending, $w_t h_t$ is labor income, and τ_t is the tax rate.

The households. Preferences are given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(c_{it}, x_{it}, h_{it}; \theta_{it}) \right] \quad (2)$$

where $\beta \in (0, 1)$ is the subjective discount factor, $x_{it} \in \mathbb{R}_+$ and $c_{it} \in \mathbb{R}$ denote period t consumption in the morning and in the afternoon respectively, $h_{it} \in \mathbb{R}_+$ denotes labor supply, and θ_{it} is an idiosyncratic taste shock.

$$U(c, x, h; \theta) \equiv c + \theta u(x) - \nu(h), \quad (3)$$

$u(\cdot)$ is strictly increasing and strictly concave, and $\nu(\cdot)$ is strictly increasing and strictly convex. The modeling role of the taste shock is to introduce a benefit from reallocating the endowment of the morning good. It is i.i.d. across households and follows a continuous Markov process, with transition density $\varphi(\theta'|\theta)$ and unconditional density $\varphi(\theta)$.

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

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

$$c_{it} + p_t x_{it} + q_t a_{it+1} = a_{it} + (1 - \tau_t) w_t h_{it} + p_t \bar{e} \quad (4)$$

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

Let $z_{it} = p_t(x_{it} - \bar{e})$ denote the value, in terms of the numeraire, of the household’s net trade of the morning good. We interpret z_{it} as *intra-period* borrowing and lending. Whenever $z_{it} > 0$, 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 whenever $z_{it} < 0$, the household is a “lender” in the sense that it has a net positive position in the market for such IOUs.

Once the afternoon arrives, a borrower may be tempted to renege on her promise to pay back. The only action available to a lender in the case of such an event is to confiscate a fraction $\xi \in (0, 1)$ of the borrower’s afternoon labor income as well as all of his assets. For default to be averted in equilibrium, the following constraint must therefore hold:

$$z_{it} \leq \xi w_t h_{it}^{def} + a_{it} \quad (5)$$

where h_{it}^{def} denotes labor supply in the (off-equilibrium) event of default.

Finally, applying the same logic to inter-period borrowing, we have the following bound on the amount of debt the household can issue at the end of each period:

$$-a_{it+1} \leq \xi w_{t+1} h_{it+1}^{def}. \quad (6)$$

Remark 1. The intra-day IOUs can be interpreted as short-term credit lines that help cover “liquidity needs” and that are collateralized by the holdings of the risk-free asset. Ruling out trades of this asset in the morning may appear to impose that this asset is less liquid than the IOUs. This is not true. Because we have allowed the entire market value of an agent’s asset holdings to be posted as collateral, the equilibrium allocations remain unaffected if we let the risk-free asset be traded in the morning alongside, or in place of, the IOUs. Accordingly, we can think of the risk-free asset either as collateral or, more simply, as buffer stock. One way or another, the key and only friction is that constraints (5) and (6) restrict the consumption bundle that household i can attain within any given period as function of its pledgeable income and its net financial worth.

Remark 2. Some of the modeling choices we make for the sake of tractability complicate the mapping between the theory and the data. For instance, it’s not clear how to interpret the morning and the afternoon in the data. Furthermore, by forcing the heterogeneity in private net worth to have no effect on macroeconomic quantities and asset prices, we abstract from mechanisms that are at the core of the recent macroeconomic literature on financial friction. Nonetheless, as it will become clear in the sequel, the essence of the friction in our model is the same as in that literature in the following two critical aspects. First, insofar as pledgeable income is sufficiently limited at the aggregate level, the economy is prevented from attaining an efficient allocation of resources in equilibrium.⁵ Second, because an agent can ease her own constraints by accumulating more assets, a precautionary motive is present, and a related premium emerges in equilibrium asset prices.

Remark 3. In line with Holmström and Tirole (1998), we have assumed that privately-issued and government-issued assets are equally good either as stores of value or as forms of collateral. This clarifies that our results do not require the government to be the sole supplier of liquidity: there can be both “outside money” in the form of T-bills and “inside money” in the form of, say, corporate

⁵In the present model, the relevant allocation is that of the morning good across consumers with different taste shocks. In the variant of Appendix A, it is that of capital across entrepreneurs with different productivity shocks.

bonds. Notwithstanding this fact, the assumption that these assets are *perfect* substitutes to one another is only for simplicity. As we elaborate in Section 6, what is key is, instead, the absence of a hypothetical asset that serves no liquidity function within the private sector but can be used as an investment vehicle by the government.

3 Equilibrium

We start the characterization of the equilibrium by studying the collateral constraint. Default does not take place in equilibrium. However, to determine the pledgeable income of the household, we need to determine its labor supply in any off-equilibrium path following default.

This raises the question of what tax rate the household would face in the event of bankruptcy. Since this is an *off*-equilibrium object, it can differ from the *on*-equilibrium tax rate τ_t . Without serious loss of generality, we assume that it is zero.⁶ The labor supply in the event of bankruptcy then solves $\nu'(h_{it}^{def}) = (1 - \xi)w_t$. Using $w_t = A$, we conclude that the financial constraints (5) and (6) can be restated as, respectively,

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

where $\phi = \phi(A, \xi) \equiv \xi A(\nu')^{-1}((1 - \xi)A)$. In Section 7, we discuss how the the dependence of ϕ on A and ξ shapes the optimal policy response to shocks. For the time being, however, we treat A , ξ , and ϕ as constants and omit the dependence of all the endogenous objects on them.⁷

We now proceed to study how the financial friction impedes the equilibrium trades of the two goods. Because of the assumed preferences, the optimal consumption of the morning good maximizes $\theta u(x) - px$ subject to the collateral constraint. Let

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

denote the resulting utility net of the cost of any purchases of the good, and let

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

be the discounted, previous-period expectation of the aforementioned object. If it were not for the collateral constraint, the optimal consumption of the morning good would have been invariant to a . The dependence of $\mathcal{U}(a, \theta, p)$ on a therefore reflects exclusively the role of assets in relaxing future collateral constraints. In this sense, $\mathcal{U}(a, \theta, p)$ captures the *private value of liquidity* for type θ .

⁶This abstracts from the possibility that taxation can “destroy” some of the collateral in the economy and isolates public debt as the only instrument through which the planner can regulate the bite of the financial friction.

⁷Note that ϕ is strictly increasing in A , whereas it is strictly concave in ξ with a maximum obtained at an interior $\bar{\xi} \in (0, 1)$. Section 7 presumes that $\xi < \bar{\xi}$, so that the friction is aggravated with either a lower A or a lower ξ .

Letting $\tilde{c}_{it} \equiv c_{it} + z_{it}$ denote the household's total expenditure on the two goods, we arrive at the following characterization of its intertemporal consumption-and-saving problem.

Proposition 1. *For any given sequence $\{p_t, q_t\}_{t=0}^\infty$, the optimal plan for $\{\tilde{c}_{it}, h_{it}, a_{it+1}\}_{t=0}^\infty$ maximizes the following objective:*

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

subject to the following constraints:

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

In a nutshell, it is *as if* individual assets entered the utility function. The only caveat to such a reinterpretation of our model is that the (indirect) utility of assets depends on the expected price of the morning good, an issue that we address in the following section. Notwithstanding this qualification, Proposition 1 has two key implications. First, all households choose the same supply of labor. Second, optimal saving depends on a household's current shock and the anticipated price of the morning good, because these two determine the probability that the collateral constraint will bind in the next period; but it is invariant to the history of earlier shocks or asset holdings.⁸

To simplify the exposition, we henceforth assume differentiability of \mathcal{U} . Optimal asset holdings then satisfy the following Euler condition:

$$\mathcal{U}_a(a_{it+1}, \theta_{it}, p_{t+1}) \leq \pi_t \equiv q_t - \beta, \tag{7}$$

with equality whenever $a_{it+1} > -\phi$. That is, either the marginal value of liquidity, as measured by \mathcal{U}_a , is equated with the premium π_t at which the asset is priced over the discount factor; or, the value of liquidity is so low that the corner solution $a_{it+1} = -\phi$ obtains.⁹

If the idiosyncratic shock is persistent, higher types will demand more liquidity (i.e., will save more) than lower types for any given premium. Sufficiently high types can thus be on the demand side of the market for the risk-free asset ($a_{it} > 0$), whereas sufficiently low types can be on the supply side ($a_{it} < 0$). In this sense, our model allows for a non-trivial private supply of collateral. However, this supply is bounded, because the pledgeable income of all types is bounded. Furthermore, insofar as the risk is sufficiently large, the demand by the high types exceeds the supply by the low types when the interest rate is equal to the discount rate, implying that the equilibrium premium must be positive. Conversely, when the equilibrium premium is zero, it must be that all agents are unconstrained. In this sense, $q_t > \beta$ signals that aggregate collateral is in short supply, whereas $q_t = \beta$ signals that the economy's demand for liquidity has been "satiated".

⁸Clearly, these properties hinge on the assumed preference specification.

⁹Throughout, we presume that $q_t \geq \beta$ ($\pi_t \geq 0$) for all t . The alternative scenario is ruled out in equilibrium: if the interest rate were higher than the discount rate, the demand for assets would have been infinite.

4 The reduced-form Ramsey problem

Instead of solving for the fixed point of the equilibrium, we choose to construct two functions that serve as “sufficient statistics” for the effects of the financial friction on welfare and on the government budget. This leads not only to a convenient reduced-form representation of the planner’s problem, but also to a clear understanding of the key policy trade offs.

The first function is denoted by $V(b)$ and is defined by the value of the following problem:

$$\max_{(p,q) \in \mathbb{R}_+^2 \ \& \ (x,a): [\underline{\theta}, \bar{\theta}] \rightarrow \mathbb{R}_+ \times [-\phi, +\infty)} \int \theta u(x(\theta)) \varphi(\theta) d\theta \quad (8)$$

$$\text{subject to} \quad \int x(\theta) \varphi(\theta) d\theta = \bar{e} \quad (9)$$

$$\int a(\theta_-) \varphi(\theta_-) d\theta_- = b \quad (10)$$

$$\phi + a(\theta_-) - p(x(\theta) - \bar{e}) \geq 0 \quad \forall (\theta, \theta_-) \quad (11)$$

$$\theta u'(x(\theta)) \geq p \quad \forall \theta \quad (12)$$

$$[\theta u'(x(\theta)) - p] [\phi + a(\theta_-) - p(x(\theta) - \bar{e})] = 0 \quad \forall (\theta, \theta_-) \quad (13)$$

$$a(\theta_-) + \phi \geq 0 \quad \forall \theta_- \quad (14)$$

$$\beta + \mathcal{U}_a(a(\theta_-), \theta_-, p) \leq q \quad \forall \theta_- \quad (15)$$

$$[\mathcal{U}_a(a(\theta_-), \theta_-, p) - \pi] [a(\theta_-) + \phi] = 0 \quad \forall \theta_- \quad (16)$$

The second function is denoted by $Q(b)$ and is defined by the value of q that obtains in the solution of the above problem. We also let $\pi(b) \equiv Q(b) - \beta$.

To interpret these objects, note that (8) is the ex-ante utility of the morning good; (9) and (10) are, respectively, the resource constraint for the morning good and the clearing condition for the asset market; and finally (11)-(16) are the household’s optimality conditions for morning consumption and asset holdings, together with the associated collateral constraints and complementary slackness conditions. It follows that the solution to (8)-(16) identifies the best cross-sectional allocation of asset holdings and morning-good consumption among those that can be implemented as an equilibrium for a given quantity b of public debt. $V(b)$ is the value of this allocation; $Q(b)$ identifies the interest rate that supports this allocation; and $\pi(b)$ is the associated liquidity premium.

Next, note that welfare (ex-ante utility) is given by $\mathcal{W} \equiv \mathbb{E} [\sum \beta^t (c_{it} + \theta_{it} u(x_{it}) - \nu(h_{it}))]$; that $\mathbb{E}[c_{it}]$ is aggregate consumption, c_t ; that all agents supply the same amount of labor; and that $\mathbb{E}[\theta_{it} u(x_{it})]$ equals $V(b_t)$ in the best implementable allocation. It follows that welfare reduces to

$$\mathcal{W} = \sum_{t=0}^{\infty} \beta^t [c_t - \nu(h_t) + V(b_t)] \quad (17)$$

The budget constraint is

$$Q(b_{t+1})b_{t+1} = b_t + g - \tau_t A h_t \quad (18)$$

Finally, note that the only implementability constraint that has not been incorporated in $V(b_t)$ is the one for the supply of labor:

$$\nu'(h_t) = (1 - \tau_t)A \quad (19)$$

The planner's problem reduces to finding the sequence $\{c_t, h_t, \tau_t, b_{t+1}\}_{t=0}^{\infty}$ that maximizes (17) subject to (18), (19), and the resource constraint.

This problem is equivalent to that of a representative-agent economy in which public debt generates a welfare flow of $V(b)$ and is priced at $Q(b) = \beta + \pi(b)$. In a model in which “assets enter the utility function”, we would further have the restriction $\pi(b) = V'(b)$, reflecting a coincidence of the private and social value of assets. Here, instead, it turns out that $\pi(b) > V'(b)$ because of a pecuniary externality that operates through the collateral constraint.¹⁰ Nonetheless, what is essential for our subsequent results is only the dependence of V and Q on b , not the aforementioned wedge. This dependence epitomizes the dual role of public debt, not only in our paper, but also in Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998): by regulating the bite of the financial friction, public debt can both ease the allocation of resources (the effect captured by V) and manipulate interest rates (the effect captured by Q).

How this dual role shapes the optimal policy will be explored in Section 5. For now, we opt to further simplify the planner's problem by solving (19) and the resource constraint for consumption and labor supply as functions of the tax rate. More specifically, let $H(\tau) \equiv (\nu')^{-1}(\frac{1-\tau}{A})$ and $S(\tau) \equiv \tau AH(\tau)$ denote the equilibrium values of, respectively, labor supply and tax revenue, as functions of the tax rate. It is straightforward to check that S is single-peaked—i.e., there is a Laffer curve—and attains its maximum value, \bar{s} , at $\tau = \bar{\tau}$ for some $\bar{\tau} \in (0, 1)$. For any $s \leq \bar{s}$, we thus have that, whenever the planner wishes to collect tax revenue equal to s , the tax rate that implements this goal is given by $\tau = T(s) \equiv \min\{\tau : S(\tau) = s\}$. Next, let $U(s) \equiv AH(T(s)) - \nu(H(T(s)))$ measure the resulting utility from consumption and leisure, as a function of tax revenue, and note that $U(s)$ is decreasing in s , reflecting the welfare cost of taxation.

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

$$\max_{\{s_t, b_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t [U(s_t) + V(b_t)] \quad (20)$$

$$\text{subject to } Q(b_{t+1})b_{t+1} = b_t + g - s_t \quad (21)$$

This result represents the planner's problem as a discrete-time optimal-control problem in which the state variable is the level of public debt and the control variable is the tax revenue (or, equivalently, the tax rate). The fact that *only* these variables show up in this problem hinges on the

¹⁰See Appendix 8 for a proof. The intuition is that, when an agent decides how much collateral to hold, she does not internalize the effect that her enhanced ability to buy the morning good will have on the price of that good and thereby on the tightness of the collateral constraint of other agents.

Lagos-Wright specification of our model; without it, the entire cross-sectional distribution of private wealth would have emerged as a state variable in place of b_t . Nevertheless, the policy trade offs that are encapsulated in the above representation are not driven by this feature of our model.

Recall that the economy exhibits two distortions relative to the first best: a financial friction and a tax distortion. The above representation reveals that, while both distortions are necessarily bad for efficiency and welfare, they can both help the government budget.

These insights seem quite general. As an example, in Appendix A we show that a similar reduced-form representation of the planner's problem obtain from alternative micro-foundations. The financial friction then impedes the efficient allocation of capital across firms and thereby also distorts investment decisions, resembling the effects found, inter alia, in Kiyotaki and Moore (1997) and Holmström and Tirole (1998). As in the latter paper, the issuance of public debt may raise aggregate TFP; but it may also crowd out capital, as in Aiyagari and McGrattan (1998). These effects represent new channels through which the financial friction can affect, not only welfare, but also the tax basis and thereby the government budget. Despite the potential richness of these new channels, the essence of the policy problem remains the same: regulating the financial distortion has a dual effect on welfare and the government budget, similarly to the present model. As it will become clear in the sequel, the key lessons of our analysis depend on this dual role of the financial distortion, not on the particulars of the micro-foundations.

With these points in mind, the analysis in the next section sidesteps the micro-foundations and studies directly a class of reduced-form policy problems like the one obtained in Proposition 2 above. To facilitate this transition, we close this section by identifying a certain threshold for the level public debt, above which the financial friction ceases to bind. This threshold is analogous to a satiation level of real-money balances in monetary models.

Let V_{bliss} denote the value of $\mathbb{E}[\theta u(x)]$ obtained at the first-best allocation of the morning good. Clearly, $V(b) \leq V_{bliss}$ for all b . Insofar as θ is bounded from above, V_{bliss} is attained for sufficiently high b ; and insofar as ϕ is sufficiently small, V_{bliss} is not attained when b is zero. There is then a threshold $b_{bliss} > 0$ such that $V(b) < V_{bliss}$ if $b < b_{bliss}$, and $V(b) = V_{bliss}$ if $b > b_{bliss}$. Finally, the agents are willing to pay a positive premium if and only if they expect their collateral constraints to bind, which in turn happens if and only if the first-best allocation is not attained. That is, $\pi(b) > 0$ if and only if $V(b) < V_{bliss}$. We conclude that b_{bliss} identifies the minimal level of public debt above which the friction ceases to bind and the private demand for collateral is “satiated”.

Depending on parameters, b_{bliss} may exceed the maximal sustainable level of public debt. In the subsequent analysis, we assume that the opposite is true. Although this is not strictly needed, it makes the analysis more interesting by permitting us to show that the planner may opt not to preserve the financial distortion even when it is feasible to eliminate it.

5 Optimal Policy: Meet Skiba

We now sidestep the micro-foundations and characterize the solution to a class of reduced-form policy problems like the one of Proposition 2. As mentioned in the Introduction, this cannot be accomplished by the usual first-order approach: because these problems are non-convex, there can exist multiple paths that satisfy the relevant first-order and transversality conditions. This is where Skiba (1978) and Brock and Dechert (1983) come to the rescue.

Skiba (1978) led the way of studying non-convex optimal-control problems in economics by allowing for a non-convex technology in the Neoclassical growth model. Brock and Dechert (1983) extended some of Skiba's results, including a variant of Arrow's maximum principle, to a richer class of settings. See also Buera (2008) for an elegant application within the context of entrepreneurship.

To facilitate the application of the relevant methods to our setting, as well as to develop a useful phase-diagram representation of the solution, we switch the analysis from discrete to continuous time. In particular, we henceforth consider the following class of problems, indexed by different configurations of the functions (U, V, R) .

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

$$\max \int_0^{+\infty} e^{-\rho t} [U(s) + V(b)] dt \quad (22)$$

$$\text{subject to } \dot{b} = R(b)b + g - s \quad \forall t \quad (23)$$

$$b(0) = b_0 \quad (24)$$

This is the continuous-time analogue of the problem described in Proposition 2. $U(s)$ and $V(b)$ capture the welfare effects of, respectively, the tax and the financial distortion; $R(b)$ is the interest rate on public debt; $\rho > 0$ is the underlying discount rate; \bar{s} is the upper bound on tax revenue; $\bar{b} \equiv \frac{\bar{s}-g}{\rho} > 0$ is the maximal sustainable level of debt; and $\underline{b} \leq 0$ is a lower bound.¹¹ Finally, we let $\pi(b) \equiv \rho - R(b)$; we put aside the details of the underlying micro-foundations; and we impose only the following minimal requirements on the objects that appear in the planner's problem.

Main Assumptions. *We consider economies in which the following properties hold:*

A1. *U, V , and π are continuously differentiable.*¹²

A2. *U is concave in s , with maximum attained at $s = 0$.*

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

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

¹¹The negative of \underline{b} can not exceed the aggregate pledgeable income of the private sector as a whole; we also assume that \underline{b} no less the highest value of b that solves $b = -g/R(b)$, which identifies the value of assets that suffices for the government to finance its spending with zero taxation.

¹²To be precise, we allow V and π to be non-differentiable at $b = b_{bliss}$.

A1 is technical. A2 means that the welfare cost of taxation is convex, a property that helps guarantee the monotonicity of the optimal debt dynamics. A3 captures the dual role of the financial friction on welfare and interest rates; it also imposes that the level of public debt that satiates the economy's demand for collateral is sustainable, a property that, as mentioned before, is not strictly needed but makes the analysis more interesting. Finally, A4 restricts the interest rate to be non-negative, an assumption that is not strictly needed but simplifies the exposition.

Beyond these restrictions, we do not need to make any further assumptions such as that the wedge between the private and the social value of liquidity has a particular sign: π can be either higher or lower than V' . In this regard, the lessons derived below appear to be robust to the kind of pecuniary externalities and the related inefficiencies that have been emphasized by, inter alia, Shleifer and Vishny (1992), Lorenzoni (2008), and Dávila (2015). We also do not restrict V to be concave. But even if we were to do so, this would not eliminate the non-convexity of the planner's problem: the non-convexity emerges naturally from the property that the liquidity premium depends on the problem's state variable (the level of public debt), regardless of the curvature of V .

The Hamiltonian and the Euler condition. Denote the costate variable with λ and define the Hamiltonian as follows:

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

The economic interpretation of λ is familiar: it measures the shadow value of tax revenue. By the same token, λ also measures the welfare cost of the interest payments on public debt.

For given values of the state and the costate, the optimal control maximizes the Hamiltonian. In our setting, the control is the tax revenue s (or, equivalently, the tax rate τ that attains s). Whenever $\lambda > 0$, the optimal tax equates the welfare cost of taxation with the shadow value of tax revenue: $U'(s) + \lambda = 0$. When instead $\lambda < 0$, taxes are set to zero.¹³ For future reference, and with some abuse of notation, we let $s(\lambda)$ denote the optimal tax as a function of the costate.

For any pair of values of the state, b , and the costate, λ , let

$$\mathcal{H}(b, \lambda) \equiv \max_s H(s, b, \lambda) \tag{25}$$

be the value of the Hamiltonian attained at the optimal value for the control. A path for (b, λ) is optimal *only if* it satisfies the planner's Euler condition at all t along with the transversality condition at infinity, namely:

$$\dot{\lambda} = \rho\lambda + \mathcal{H}_b(b, \lambda) \quad \forall t \quad \text{and} \quad \lim_{t \rightarrow \infty} e^{-\rho t} \lambda(t) b(t) = 0.$$

If the problem were convex, the above conditions together with the budget constraint and the initial condition $b(0) = b_0$ would not only be necessary but also sufficient for optimality. However,

¹³If we had allowed the tax rate to be negative, the condition $U'(s) + \lambda = 0$ would apply even when $\lambda < 0$. This, however, would not affect the rest of the analysis.

as already explained, the problem under consideration is not convex. As a result, it is possible that there exist multiple paths that satisfy all these conditions. The challenge is then to find which of the many candidate paths is the optimal one. Before addressing this issue, however, let us first elaborate on the policy trade offs encapsulated in the planner's Euler condition.

The Key Trade Offs. Using the definition of \mathcal{H} in (25), the planner's Euler condition can be restated as follows:

$$\dot{\lambda} = v(b) - \lambda\pi(b)(\sigma(b) - 1). \quad (26)$$

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

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

is the elasticity of the premium with respect to the quantity of public debt. Consider, as a reference point, the case in which public debt has no liquidity value, so that $v(b) = \pi(b) = 0$ for all b . Condition (26) 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 (26) is non-zero, optimality requires a non-zero drift in λ , that is, a deviation from tax smoothing.

To interpret this deviation and develop intuition, we momentarily consider an analogy between our dynamic economy and a certain static problem. In Section 6, variants of this analogy will also help relate our work to the literature.

Consider a one-period economy, in which consumers value a good that can be produced and supplied only by the government. Let the consumption of this good be denoted by b and its *social* value by $V(b)$, for some increasing function V . Suppose further that the *private* valuation of this good at the margin is $\pi(b)$, which means that the government can charge the market a price $\pi(b)$ for its provision. Finally, let λ be the government's shadow value of the profit earned from selling this good to the consumers. We can then express the optimal supply of this good as

$$b^* = \arg \max_b \Omega(b, \lambda), \quad (27)$$

where

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

is the sum of the social value of the good and the shadow value of the profit that the government makes from selling this good.

In our setting, the relevant good is the government's supply of the risk-free asset, or equivalently the quantity of public debt, and $V(b)$ measures its social value. As for the profit that the government

can make from selling public debt, this is tied to the wedge between the market interest rate, $R(b)$, and the social discount rate, ρ . In the standard Ramsey paradigm, this wedge is zero, implying that the planner sees neither a profit nor a cost to issuing an extra unit of debt. In our setting, by contrast, the financial friction depresses the interest rate below the social discount rate. This in turn explains the precise sense in which the government can profit from the provision of liquidity to the private sector. This profit is given by $\pi(b)b$. It is akin to seigniorage in monetary models and can be thought of as a fee that the planner charges the market for the service provided.

Next, suppose that $\pi(b)b$ is single-peaked, attaining a global maximum at $b = b_{seig}$. Clearly, $b_{seig} \in (0, b_{bliss})$, since $\pi(b)b$ is zero at the boundaries of this interval, strictly positive in its interior, and non-positive outside of it. By contrast, $V(b)$ is increasing in b and maximized at $b = b_{bliss}$. It follows that raising b increases both $V(b)$ and $\pi(b)b$ as long as $b < b_{seig}$; once $b > b_{seig}$, however, raising b increases $V(b)$ only at the expense of reducing $\pi(b)b$. This captures the trade off between collateral creation, or “liquidity provision”, and interest-rate manipulation, or “seigniorage”, that was mentioned in the introduction.

By construction of Ω , the net value of raising b at the margin is given by Ω_b . Clearly, an interior solution to (27) requires $\Omega_b = 0$: the solution of the static problem requires that the marginal loss in terms of seigniorage just offsets the marginal gain in terms of extra welfare.

What about the solution to our dynamic problem? Its Euler condition can be restated as $\dot{\lambda} = \Omega_b$. In a steady state, this reduces to $\Omega_b = 0$, which is the same as the FOC of the static problem mentioned above. At first glance, this suggests that the steady state of our problem is analogous to the solution of the aforementioned static problem. We will later show that this is true only if we impose additional restrictions on V and π . Even then, the analogy has to be qualified by the fact that the steady-state value of λ is endogenous to the steady-state value of b in our setting, whereas the aforementioned analogy treated λ as exogenous. Last but not least, our dynamic problem may (generically) admit multiple steady states, whereas the static problem (generically) admits a unique solution. These observations underscore the limitations of the proposed analogy.

That said, the analogy helps understand how optimal policy in our setting balances the aforementioned two objectives, that of enhancing liquidity and that of manipulating the interest rate, with a third objective, that of smoothing the tax distortion. To see this more clearly, note that the Euler condition equates λ , the drift in the shadow value of the government budget, with Ω_b , the marginal effect of b on Ω . The former captures the welfare cost of departing from tax smoothing; the latter captures the dual effect of the financial friction on welfare and the government budget.

These considerations then suggest the following analogy. The steady state can be interpreted as a long-run target for the level of public debt and the associated tax and financial distortions. As long as the initial level of debt is away from this target, reaching the target requires a departure from tax smoothing. The welfare cost of this departure acts as an adjustment cost: the larger this

cost, the larger the sacrifice that the planner is willing to suffer in terms of accepting a deviation from the target at any point in time. The analysis in the remaining of this section sheds light on how the aforementioned three policy objectives interact into shaping, not only the transitional dynamics, but also the target itself: even though the tax-smoothing objective drops out once the economy has reached a steady state it still plays a crucial role in determining which is the relevant steady state in the first place.

The phase diagram. In order to provide a complete characterization of the optimal policy, we study the phase diagram of the ODE system for b and λ implied by the budget constraint and the planner's Euler condition. Consider first the budget constraint. This can be expressed as follows:

$$\dot{b} = \Psi(b, \lambda) \equiv g + (\rho - \pi(b))b - s(\lambda), \quad (28)$$

where $s(\lambda)$ denotes the optimal tax revenue. It is straightforward to check that $s(\lambda)$ is increasing in λ and therefore that $\Psi(b, \lambda)$ is decreasing in λ : a higher λ means higher taxes today, which in turn means lower debt tomorrow.¹⁴ By the Implicit Function Theorem, we then have that there exists a function $\psi : [\underline{b}, \bar{b}] \rightarrow \mathbb{R}_+$ such that $\Psi(b, \psi(b)) = 0$ for all b ; equivalently,

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

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

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

The proof as well as the interpretation of this lemma is straightforward: $\psi(b)$ 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 planner's Euler condition. Let $\Delta \equiv \{b \in [\underline{b}, b_{bliss}) : \sigma(b) \neq 1\}$ and define the function $\gamma : \Delta \rightarrow \mathbb{R}$ as follows:

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

¹⁴Note 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.

where, recall, $v(b) \equiv V'(b)$ and $\sigma(b) \equiv -\pi'(b)b/\pi(b)$. We can then restate the planner's Euler condition, namely condition (26), as follows:

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

By implication,

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

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

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

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

It follows that the points at which $\sigma(b) = 1$ correspond to the critical points of the function $\pi(b)b$, which, as explained before, can be interpreted 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.

A Useful Benchmark

To develop a better understanding of the general case, we find it useful to start with a certain benchmark. This benchmark is not meant to be empirically plausible; we don't even have much insight to offer on what micro-foundations may justify it. Its function is purely pedagogical.

Benchmark. (i) the ratio v/π is constant; (ii) the elasticity σ is monotone.

The first assumption imposes that the wedge between the social and the private value of collateral is invariant to b ; this nests as a special case the case in which the wedge is absent and $\pi = v \equiv V'$, similarly to monetary models. The second assumption guarantees that $\pi(b)b$ is single-peaked and also extends the aforementioned local monotonicity of σ to its entire domain. Together, these assumptions lead to following sharp characterization of the optimal debt dynamics.

Theorem 1. *Consider the benchmark defined above. There exists a unique $b^* \in (b, b_{bliss}]$ 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 *unique* steady-state level of public debt, to which the economy converges from *any* initial point $b_0 < b_{bliss}$; it also relates b^* to b_{bliss} , the satiation point. We therefore reach the following lessons about the optimal supply of liquidity and the associated financial distortions. As long as the financial friction is binding to start with (namely $b_0 < b_{bliss}$), the friction remains binding at any finite time ($b_t < b_{bliss}$ for all t). In this sense, the analogue of the Friedman rule does not apply in the short run. Whether it applies in the long run depends on the relation between b^* and b_{bliss} , which in turn depends on the tightness of the government's budget constraint: if the constraint is sufficiently tight, as captured above by the restriction $g > \hat{g}$, then the planner finds it optimal to preserve the financial distortion, not only along the transition to steady state, but also in steady state. In fact, when $g > \hat{g}$ and $b_0 \in (b^*, b_{bliss})$, the planner opts to aggravate the friction over time, squeezing liquidity out of the economy.

In the sequel, we prove the above result in multiple steps, developing additional insights on the way. We start by noting that the assumptions of our benchmark imply the following structure for the function γ , which is instrumental for the subsequent analysis.

Lemma 2. *Consider the benchmark defined above. 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, $\gamma(b) \rightarrow -\infty$ as $b \rightarrow b_{seign}$ from below and $\gamma(b) \rightarrow +\infty$ as $b \rightarrow b_{seign}$ from above.*

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 $\psi_{bliss} \equiv \psi(b_{bliss})$,¹⁵ we have the following property.


Lemma 3. *In the benchmark defined above, the following are true. 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 $b^\# \in (b_{seign}, b_{bliss})$.*

The two scenarios are illustrated in Figures 1 and 2. Let us first consider Figure 1. 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.

¹⁵Recall 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.

The above properties also hold true in Figure 2. What distinguishes the two figures is whether γ and ψ admit an intersection within region M. In Figure 1, 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 1: 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, the ODE system (28)-(29) admits no steady state to the left of the satiation point (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, we 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 (28)-(29), then, and only then, the economy converges asymptotically to $(b_{bliss}, \lambda_{bliss})$. In other words, there is a unique path that satisfies the planner's Euler condition and the budget constraint at all dates, and that eventually leads to satiation. This path is indicated with blue color in the figure.¹⁶

¹⁶One cannot rule out the possibility that λ_0 is negative for sufficiently low b_0 . When this is the case, the negative λ signals the high value that the planner attaches to raising the level of public debt. In fact, if it were feasible for

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 Theorem 1 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 3. *Consider the benchmark defined above 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.*

Let us now consider Figure 2. In this case, γ and ψ intersect exactly once, at $b = b^\# \in (b_{seig}, b_{bliss})$. Let $\lambda^\# \equiv \psi(b^\#)$ denote the shadow cost of taxation associated with balancing the budget when $b = b^\#$. By construction, the pair $(b^\#, \lambda^\#)$ identifies the unique steady state of the ODE system (28)-(29) to the left of the satiation point (i.e., within regions L and M). As is clear from the figure, this steady state is saddle-path stable.¹⁷ In particular, for any $b_0 < b_{bliss}$, we can find a continuous path that satisfies conditions (28)-(29) and that asymptotically converges to $(b^\#, \lambda^\#)$. Exactly the same arguments as in Figure 1 guarantee that this path is the unique candidate for optimality, and hence also the optimal path, as long as $b_0 < b_{bliss}$.


A crucial difference from the case in Figure 1 is that the economy now converges to a steady state characterized by a debt level that is strictly lower than the satiation level: Theorem 1 applies with $b^* = b^\# < b_{bliss}$. 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 $b_0 > b_{bliss}$.

debt to jump upwards, the planner would choose to do so up to the point where λ turns non-negative, and only thereafter to 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 $\lambda < 0$.

¹⁷The technical appendix establishes this result formally.

Figure 2: Benchmark, with $\psi_{bliss} > \gamma_{bliss}$.



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

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

For any b_0 , let $\mathcal{P}(b_0)$ be the set of all the paths for (b, λ) that start from b_0 , satisfy the ODE system in all t , and also satisfy the transversality condition at infinity. Since these conditions are

necessary for optimality, the optimal path is necessarily contained in $\mathcal{P}(b_0)$. More generally, we can reduce the planner's problem to that of choosing a path $\mathcal{P}(b_0)$. Next, note that any path in $\mathcal{P}(b_0)$ is associated with a different initial value for the costate and let $\Lambda(b_0)$ be the set of such initial values for the costate. Choosing a path in $\mathcal{P}(b_0)$ is therefore equivalent to choosing an initial value λ_0 in $\Lambda(b_0)$. The following result is helpful for evaluating the welfare associated with any candidate path.

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

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

Let us now go back to Figure 2. Pick any $b_0 \geq b_{bliss}$ and suppose there exists a continuous path that satisfies the ODEs and asymptotically converges to $b^\#$. As already noted, this path is a candidate for optimality. But so is the Barro-like plan that keeps b and λ constant for ever at, respectively, b_0 and $\psi(b_0)$. Note, next, that the first plan is associated with a higher λ_0 (i.e., higher taxes) than the second, because the first runs a surplus whereas the second balances the budget. Finally, note that, along any candidate path, $\mathcal{H}_\lambda(b, \lambda) = \dot{b}$. For the path that leads the economy to $b^\#$, we have that $\dot{b} < 0$ at $t = 0$, and hence $\mathcal{H}_\lambda(b_0, \lambda_0) < 0$. For the Barro-like plan, instead, $\dot{b} = 0$ and hence $\mathcal{H}_\lambda(b_0, \lambda_0) = 0$. Since \mathcal{H} is convex, this means that the Barro-like plan attains the minimum of \mathcal{H} over the set of candidate paths. It follows that, whenever the path that takes the economy to $b^\#$ exists, this path strictly dominates the Barro-line, 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 $\lim_{b \rightarrow \bar{b}} \psi(b) = \infty$; this verifies that b_{skiba} exists and is necessarily strictly between b_{bliss} and \bar{b} . It is then immediate that a continuous path that satisfies the ODEs and that converges to $b^\#$ exists if and only if $b_0 < b_{skiba}$, as illustrated in the figure.

We thus have the following complement to Proposition 3.

Proposition 4. *Consider the benchmark defined above and suppose $\psi_{bliss} > \gamma_{bliss}$. There exist unique points $b^\# \in (b_{seig}, b_{bliss})$ and $b_{skiba} \in (b_{bliss}, \bar{b})$ such as the optimal level of debt converges monotonically to $b^\#$ if $b_0 < b_{skiba}$, whereas it stays constant at b_0 for ever if $b_0 \geq b_{skiba}$. Taxes exhibit a positive drift as long as $b \in (b_{seig}, b^\#)$, a negative drift as long as $b \in (b^\#, b_{bliss})$, and are smoothed as long as $b > b_{bliss}$.*

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 3 and 4 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, as well as how it relates to the literature on the Friedman rule. 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. *$\psi_{bliss} > \gamma_{bliss}$ 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 Theorem 1 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.

Do the lessons obtained above apply outside the benchmark under consideration? Before we address this question, we complete the characterization of the benchmark by studying the comparative statics of the steady-state level of debt.

Proposition 5. *As long as $\psi_{bliss} > \gamma_{bliss}$, the steady-state debt level $b^\#$ increases with v , decreases with g , and can either increase or decrease with π .*

This result underscores how the steady-state level of debt resolves the trade-off between collateral creation and interest-rate manipulation, or between liquidity and seigniorage. The steady-state level $b^\#$ is always between b_{seig} and b_{bliss} , the levels that maximize, respectively, the seigniorage, $\pi(b)b$, and the social value of liquidity, $V(b)$.

If we shift the function v upwards holding the function π constant, we increase the marginal social value of liquidity while keeping constant both marginal and total seigniorage. This means that not only is the social value of easing the friction increased but also that the tightness of the budget constraint is maintained. The first effect is manifested in an upward shift in the graph of γ ; the latter in a constant position of ψ . Consequently, $b^\#$ increases unambiguously.

By contrast, if we raise g , we leave the relative value of liquidity unchanged but tighten the budget (γ stays constant but ψ increases), which in turn explains why $b^\#$ decreases with g and indeed gets closer to b_{seig} as g approaches its upper bound.¹⁸ A higher g means that it is more expensive to provide liquidity (equivalently, more worthwhile to extract seigniorage).

Next, if we keep v and σ constant but raise π , we increase both marginal and total seigniorage. The first effect shifts γ down, reflecting the reduction in the relative value of providing liquidity; the second effect shifts ψ down, reflecting the relaxation of the budget. In other words, the increase in π , holding v constant, makes the provision of liquidity cheaper but also relatively less valuable in the eyes of the planner, which explains why the overall effect on $b^\#$ is ambiguous.

These comparative statics anticipate the analysis of shocks conducted in Section 7. They also help highlight the following elementary point: as long as pecuniary externalities like those studied in Shleifer and Vishny (1992), Lorenzoni (2008), and Dávila (2015) drive a wedge between the private and the social private value of collateral, this wedge naturally impacts the planner's long-run target for the level of public debt, but does not alter the qualitative nature of the optimal policy dynamics.

Beyond the Benchmark

The benchmark studied above has two key properties: $\pi(b)b$ is singled-peaked, so that the phase diagram can be organized in the three regions described above; and γ is decreasing over the region M , so that it can intersect at most once with ψ . If we modified the benchmark by allowing either for

¹⁸The latter is given by $g_{max} \equiv T_{max} + [\pi(b_{seig}) - \rho]b_{seig}$, where T_{max} is the maximal feasible tax revenue.

a non-monotone σ or for $V' \neq \pi$ but maintained the aforementioned properties, then the preceding arguments go through and Propositions 3 and 4 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 .¹⁹ 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 . Figure 3 here along with Figures 10 and 11 in Appendix B 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.

Consider Figure 3. In order to simplify the exposition, we truncate region L , where $b < b_{seig}$, γ is negatively valued, and there can be no steady state; we thus focus on region M , where $b \in (b_{seig}, b_{bliss})$ 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_{seig} , 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 (which is proved in the Appendix), 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 6. *Consider any $(b^\#, \lambda^\#)$ such that $\lambda^\# = \gamma(b^\#) = \psi(b^\#)$, that is any steady-state point of the ODE system in the region to the left of the satiation point. There exists a finite scalar $\chi > 0$ such that the local dynamics around that steady-state point are*


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

This result restricts the *local* dynamics of the ODE system in the neighborhood of any steady state point, that is, around the intersections of γ and ψ . Consistent with this result, Figure 3 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.

¹⁹Recall that γ is necessarily decreasing in a neighborhood to the right of b_{seig} , 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_{seig}, b_{bliss}) .

Notwithstanding these restrictions on the local dynamics, there remain three distinct possibilities with regard to the *global* dynamics. Figure 3 considers one of these possibilities; to economize on space, the other two are delegated to Appendix 8.

Figure 3: Beyond the Benchmark: Rich Dynamics and Multiple Steady States



In Figure 3, 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 point gets closer to the graph of ψ , meaning that value of \dot{b} gets closer to 0. In the knife-edge case in which $b_0 = \tilde{\tilde{b}}$, this path is associated with $\dot{b} = 0$. Conversely, the path that leads to the highest steady state is associated with $\dot{b} > 0$ for all $b_0 \in (\tilde{b}, \tilde{\tilde{b}}]$, and with $\dot{b} = 0$ in the reverse knife-edge case in which $b_0 = \tilde{b}$.

Combining these observations, we obtain the following properties. When $b_0 = \tilde{b}$, the path that leads to $b_L^\#$ features $\mathcal{H}_\lambda = \dot{b} < 0$, whereas the path that leads to $b_H^\#$ features $\mathcal{H}_\lambda = \dot{b} = 0$. By the convexity of \mathcal{H} , the latter path is dominated. Conversely, when $b_0 = \tilde{\tilde{b}}$, it is the former path that now features $\mathcal{H}_\lambda = \dot{b} = 0$ and that is therefore dominated. By continuity,²⁰ 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 3 for an illustration: the bold segments of the two stable arms indicate the optimal selection among the two candidate paths.²¹

So far, we focused on region M. In region H ($b_0 \geq b_{bliss}$), the analysis is similar to Figure 2. 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^L 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 ($b_0 < b_{seig}$), there is a unique candidate path, one leading to $b_L^\#$.

Optimal policy in Figure 3 has the following properties: whenever $b_0 < \hat{b}$, debt converges monotonically to $b_L^\#$; whenever $b_0 \in (\hat{b}, b_{skiba})$, debt converges monotonically to $b_H^\#$; and 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.

In Appendix 8, we consider two variants of Figure 3 that maintain the same 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 $b_M^\#$. As we explain in the Appendix, this path then emerges as the

²⁰Here, 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).

²¹In 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 path for *all* initial conditions: in the one variant (Figure 10), it is optimal to converge to $b_H^\#$ for all $b_0 < b_{skiba}$; in the other (Figure 11), it is optimal to converge to $b_L^\#$.

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

The number of possible scenarios would increase if we allowed γ and ψ to intersect more than three times. And yet an additional layer of complexity would emerge if the assumption that $\pi(b)b$ is single-valued were relaxed, as the tripartite structure of the phase diagram would be lost. As we illustrate in the Appendix, the phase diagram would then look like the outcome of patching together *multiple* pairs of L and M regions from our earlier examples.²² Furthermore, γ and ψ may admit an arbitrary number of intersections within *each* of the M regions, and the dynamics both within and across regions. Therefore, it is simply impossible to provide a complete taxonomy of all the possible scenarios once we departed from our benchmark.

Notwithstanding this complexity, we can establish the following result, which offers a qualified generalization of Proposition 4 in our benchmark.

Proposition 7. *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} .*

The proof is as follows. 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 $\pi(b)b$.²³ 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

²²As before, L designates regions in which $\pi(b)b$ is locally increasing, $\sigma(b) < 1$, and $\gamma(b) < 0$, whereas M designates regions in which the converse properties hold.

²³Because $\pi(b)b$ 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

the continuity and differentiability of γ and ψ , there exists at least one point $b^\# \in (b^\ddagger, b_{bliss})$ such that $\gamma(b^\#) = \psi(b^\#)$ and $\gamma'(b^\#) < \psi'(b^\#)$, that is, a steady-state point in which γ intersects ψ from above. If there are multiple such points, consider the highest one. By Proposition 6, we know that this steady state is saddle-path stable. Similarly to Figure 2, 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.²⁴

As already noted, Proposition 7 provides a generalization of Proposition 4 from our benchmark: $\gamma_{bliss} < \psi_{bliss}$ guarantees that it is optimal to lead the economy to a steady state in which the financial distortion binds in order to dampen the interest rate on public debt, not only for all initial levels of debt below b_{bliss} , but also over a range of initial levels above it. Unlike our benchmark, though, the optimal long-run quantity of public debt (and the associated levels of the financial distortion and the interest rate) need not be uniquely determined: as in the example of Figure 3, the economy may converge to different steady states starting from different initial conditions.

In our benchmark, $\gamma_{bliss} < \psi_{bliss}$ was both sufficient and necessary for the aforementioned long-run property to hold: sufficiency was established in Proposition 4, necessity in Proposition 3. Proposition 7 shows that sufficiency holds more generally, the same is not true for necessity. We prove this by example in the Appendix (see Figure 13 in particular).

$\sigma(b) < 1$ 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} .

²⁴This 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.

Complete Characterization

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

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

be the set of the points at which γ intersects ψ from above. As shown in Proposition 6, these points identify the saddle-path stable steady states of the ODE system.²⁵ Depending on primitives, $B^\#$ may be empty, or may contain an arbitrary number of elements.²⁶ Regardless of this, we have the following result.

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

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

The proof of this result follows from arguments similar to those underlying our earlier results and is therefore delegated to the Appendix. 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, then debt converges to the unique point in B^* for all $b_0 < b_{skiba}$. When instead B^* contains multiple points, then 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 2 and 3, B^* is a strict subset of $B^\#$ in Figures 10 and 11 in the Appendix. Finally, the result allows for $B^* = \emptyset$ and, equivalently, $b_{skiba} = \underline{b}$, thus nesting cases in which satiation obtains in the long run regardless of initial conditions. This in turn can be true either because $B^\#$ is itself empty, as in Figure 1, or because $B^\#$ is not empty but the path that leads to satiation happens to dominate any other candidate path.

²⁵In 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: an arbitrarily small perturbation of the primitives can break guarantee that $\gamma'(b) \neq \psi'(b)$ whenever $\gamma(b) = \psi(b)$. In any event, this detail does not affect the validity of the result stated below, because this result allows B^* to be a *strict* subset of $B^\#$.

²⁶We 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.

Theorem 2 characterizes the optimal policy for a *given* economy, that is, for a given specification of the reduced-form functions (U, V, π) or of the underlying micro-foundations. The next theorem offers a taxonomy of *all* the economies under consideration.

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 $b_{skiba} \geq b_{bliss}$.*

Furthermore, $\psi_{bliss} > \gamma_{bliss}$ is sufficient for an economy to belong to the last class.

To sum up, the result that the optimal path leads the economy to a steady state in which the financial distortion is preserved in order to depress the interest-rate cost of public debt, obtains *at least* as long as $\psi_{bliss} > \gamma_{bliss}$. As noted before, this condition has a simple interpretation: 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 result 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 $g > \hat{g}$.*

Local Dynamics and Local Comparative Statics. We conclude this section with two additional results that facilitate our later analysis of the response of the economy to shocks. 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 generalization of the comparative statics of Proposition 5.

Proposition 8. *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.*

Proposition 9. *For any $b^* \in B^*$, b^* increases continuously with a small enough increase in V or a small enough decrease in g , and is ambiguously affected by a small change in π .*

These two results together imply that, at least for small changes in either initial conditions or 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 in our benchmark. This observation forms the basis for Section 7, where we study the optimal policy response to financial and other shocks.

²⁷In 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.

6 Discussion and Additional Insights

In this section we offer some additional insights; we also link our contribution to the extant literature.

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

Suppose that the government enacts a law that prohibits the use of corporate bonds as collateral in morning transactions. This restriction adds a constraint to the planner's sub-problem defined in (8)-(16) and results in a change in the functions V and π . By shutting the private supply of collateral down the law may reduce V and increase π , so it has ambiguous welfare implications. But its effect on the price of corporate bonds is unambiguous: they are now priced at the discount factor, while government bonds command a premium over the discount factor. Our model is now directly comparable to those in the Friedman rule literature, with government bonds playing the role of money and corporate bonds the role of the non-money asset.

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

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

or equivalently

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

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

Therefore, when the government varies b , it does not face the key trade off present in our model. By the same token, the optimal supply of liquidity is disentangled from the optimal dynamics of debt and taxes.

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

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

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

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

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

$$m^* = \arg \max_m \Omega(m, \lambda^*), \quad (33)$$

where $\Omega(m, \lambda) \equiv V(m) + \lambda \pi(m)m$ measures the total value of “liquidity plus seigniorage”. Depending on primitives, m^* may or may not coincide with satiation. Regardless of this, however, tax smoothing obtains and the optimal fiscal policy is determined in exactly the same fashion as in the Barro/AMSS benchmark.

The exercise conducted above captures the essence of the literature on the Friedman rule. Saying that m^* may or may not coincide with satiation is equivalent to saying that, depending on primitives, the Friedman rule may or may not be optimal. The Friedman rule literature has identified, within certain types of models, conditions on preferences, technologies, and tax instruments that suffice for the optimality of the Friedman rule. The micro-foundations of the models studied in that literature are quite different from those in our model, making it difficult to identify the types of primitives that lead to satiation in our specification. Nonetheless, the exercise conducted above share with the literature on the Friedman rule the property that the supply of liquidity is disentangled from fiscal policy.

This is no longer the case when both assets contribute to liquidity. Section 5 established this under the simplifying assumption that the two assets are perfect substitutes to one another. The results obtained, however, are more general. To see this, consider the following extension of the two-asset exercise conducted above. Let $V(m, n)$ measure the joint welfare effect of the two assets, and let $\pi^m(m, n)$ and $\pi^n(m, n)$ denote their respective premia. Next, define $\hat{\Omega}(b, \lambda)$ as the maximum of $V(m, n) + \lambda[\pi^m(m, n)m + \pi^n(m, n)n]$ subject to $m + n = b$. Finally, let $\hat{\Pi}(b, \lambda)$ be the value of $\pi^m(m, n)m + \pi^n(m, n)n$ attained at the aforementioned maximum. The planner's Euler condition and the budget constraint then reduce to $\dot{\lambda} = \hat{\Omega}_b(b, \lambda)$ and $\dot{b} = Rb + g - s(\lambda) - \Pi(b, \lambda)$, respectively. This is essentially the same ODE system as the one in our baseline model. Unless $\hat{\Omega}_b$ is identically zero (the knife-edge case studied in the Friedman-rule literature), optimal policy departs from tax smoothing and gives rise to interesting dynamics in b . The exercise above then underscores the potential robustness of our findings to multi-asset settings that allow for compositional effects in the government's balance sheet.

To recap, the intertwining of fiscal policy with liquidity provision, and the resulting properties of the transitional, optimal dynamics of debt and taxes represent the key differences between our

results and those in the Friedman rule literature.²⁸ But they are not the only ones. In the exercise above, m^* attains the global maximum of $\Omega(m, \lambda)$. In our context, any steady state satisfies the first-order condition $\Omega_b(b^*, \lambda^*) = 0$. Yet, this does not necessarily mean that b^* maximizes $\Omega(b, \lambda^*)$. We elaborate on this point below.

On the optimal quantity of public debt. What is the optimal long-run quantity of public debt? The insights of the literature on the Friedman rule, as well as the static analogy we built in the beginning of Section 5, may lead one to conjecture that the optimal steady-state level of debt maximizes “liquidity plus seigniorage”, in the sense that $b^* = \arg \max \Omega(b^*, \lambda^*)$. We now address the validity of this conjecture.

Proposition 10. *(i) In our benchmark, whenever B^* is non-empty, it is given by a single point b^* such that $b^* = \arg \max_b \Omega(b, \lambda^*)$ and $\lambda^* = \psi(b^*)$.*

(ii) More generally, the following is true for any economy in which B^ is non-empty. Take any $b^* \in B^*$ and let $\lambda^* = \psi(b^*)$. If $\gamma'(b^*) < 0$, b^* attains a local maximum of $\Omega(b, \lambda^*)$. If instead $\gamma'(b^*) > 0$, b^* attains a local minimum of $\Omega(b, \lambda^*)$.*

(iii) There exist economies in which B^ is a singleton and, nevertheless, b^* attains a local minimum of $\Omega(b, \lambda^*)$.*

To prove this result, note first that

$$\Omega_b(b, \lambda^*) = (\sigma(b) - 1) \pi(b) [\gamma(b) - \lambda^*] \quad (34)$$


Next, 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 $b < b_{seig}$, whereas $\sigma(b) > 1$ and $\gamma(b) > 0$ for all $b > b_{seig}$. Finally, the monotonicity of γ guarantees that $\gamma(b) > \gamma(b^*)$ for $b \in (b_{seig}, b^*)$, whereas $\gamma(b) < \gamma(b^*) = \lambda^*$. Together with the fact that $\gamma(b^*) = \psi(b^*) = \lambda^* > 0$, this implies that $\Omega_b(b, \lambda^*) > 0$ for all $b \in [\underline{b}, b^*)$ and $\Omega_b(b, \lambda^*) < 0$ for all $b \in (b^*, \bar{b}]$, which proves part (i). $\pi(b)b$ need not be single-peaked and, most importantly, $\gamma(b)$ need not be monotone outside our benchmark. The argument above therefore breaks down: there can exist multiple steady states below b_{bliss} and, regardless of how many they are, there is no guarantee that *any* of them maximizes “liquidity plus seigniorage”. The only thing we know is that at any such point $\sigma(b^*) > 1$ and $\gamma(b^*) = \lambda^*$, which in turn proves part (ii): in a neighborhood of b^* , $\Omega_b(b, \lambda^*)$ crosses zero from above (meaning that b^* is a local maximum) if $\gamma'(b^*) < 0$, whereas it crosses zero from below (meaning that b^* is a local minimum) if $\gamma'(b^*) > 0$.

Figure 3 offers an example of the first case: there are two steady states, both satisfying $\gamma'(b^*) < 0$ and therefore both representing local maxima. The proof of part (iii) is completed by Figure 4, which offers an example of the second case.

²⁸Note that this difference is present even if we restrict attention to economies in which it is optimal to satiate the demand for collateral in the long run.

Let us elaborate. In the example of Figure 4, γ and ψ intersect only once, at the point (b^*, λ^*) . As in our benchmark, the economy converges to this point from any initial point $b_0 < b_{skiba}$, for some threshold $b_{skiba} > b_{bliss}$. Unlike our benchmark, however, γ is not globally decreasing. Instead, γ is “moderately” increasing after some point, so that $0 < \gamma'(b^*) < \psi'(b^*)$. Letting $\gamma'(b^*) < \psi'(b^*)$ guarantees that the dynamics are characterized by saddle-path stability, which is necessary for optimality. Letting $\gamma'(b^*) > 0$ guarantees that b^* attains a local minimum of $\Omega(b, \lambda^*)$. The global maximum of $\Omega(b, \lambda^*)$ is instead attained at \hat{b} .²⁹

Figure 4: Converging to a Local Minimum



How could it be optimal to stay at b^* instead of trying to reach \hat{b} ? The latter would indeed be optimal if λ were to stay constant at λ^* . Intuitively, when $\lambda = \lambda^* = \psi(b^*)$, the shadow cost of taxation is high enough that the value of the extra seignorage obtained by lowering b from b^* towards \hat{b} outweighs the reduction in liquidity. A “myopic” planner that treats λ as constant would thus do exactly this. However, if the economy were to transit to the lower debt level \hat{b} , then the shadow cost of taxation would also be lower. In fact, once $\lambda = \hat{\lambda} \equiv \psi(\hat{b})$, the same myopic planner would now find it optimal to raise b above \hat{b} , possibly up to b_{bliss} . This suggests that the myopic planner could engage in pointless cycles and underscores the limits of either the static analogy we built earlier or any attempt to extrapolate from the literature on the Friedman rule.

The bottom line is that, in general, the optimal long-run target for the level of public debt can not be understood in isolation of the departure from tax smoothing that is required for the

²⁹In the figure, \hat{b} is defined by the lowest intersection of the graph of γ with the horizontal line at λ^* . By construction, $\gamma(b) > \lambda^*$ for $b \in (b_{seig}, \hat{b})$, $\gamma(b) < \lambda^*$ for $b \in (\hat{b}, b^*)$, and $\gamma(b) > \lambda^*$ for $b \in (b^*, b_{bliss})$. Using (34) along with the fact that $\sigma(b) < 1$ and $\gamma(b) < 0$ for $b < b_{seign}$, we infer that Ω_b is positive for all b up to \hat{b} , negative for b between \hat{b} and b^* , positive for b between b^* and b_{bliss} , and zero thereafter. It follows that the maximum of $\Omega(b, \lambda^*)$ is necessarily attained either at \hat{b} or at b_{bliss} . Finally, we can guarantee that the maximum is attained at \hat{b} , rather than at b_{bliss} , by picking the locus of γ and the product $(\sigma(b) - 1) \pi(b)$ is such that the integral of Ω_b over (b^*, b_{bliss}) is lower, in absolute value, than the integral of Ω_b over (\hat{b}, b^*) .

transition to such a target: even though tax-smoothing considerations may appear to be dormant once the economy has reached its steady state, they can play a key role in determining what the steady state is in the first place.

More on the optimal quantity of public debt. We now relate our results to the results of Aiyagari and McGrattan (1998) on the optimal long-run quantity of public debt. This paper relies on a short-cut to escape the computational challenges of studying optimal policy in incomplete-market settings: instead of solving the problem of a Ramsey planner who chooses the dynamic path of taxes and debt levels so as to maximize ex-ante utility, it restricts taxes and debt to be constant over time, abstracts from transitional dynamics, and maximizes welfare in steady state.

In order to understand how this approach may fail to provide the right guidance to policy, we proceed to embed the Aiyagari and McGrattan exercise into our framework. Consider a planner who maximizes $U(s) + V(b)$ subject to $r(b)b = s$. Let \hat{b} denote the level of debt that solves this problem and let $\hat{\lambda}$ be the relevant Lagrange multiplier.³⁰ Clearly,

$$\hat{b} = \arg \max_b \left\{ V(b) - \hat{\lambda}[\rho - \pi(b)]b \right\} = \arg \max_b \left\{ \Omega(b, \hat{\lambda}) - \hat{\lambda}\rho b \right\}, \quad (35)$$

whose first-order condition gives $\Omega_b = \hat{\lambda}\rho > 0$. By contrast, the steady state solution in our planning problem satisfies $\Omega_b = 0$. This discrepancy reveals a simple fact: the Aiyagari and McGrattan exercise incorrectly treats the entire interest payments on public debt, $r(b)b$, as a cost while in reality the social planner should view debt issuance as a profit (seignorage) generating exercise to the tune of $\pi(b)b$.

In our benchmark, if we assume that $\psi_{bliss} > \gamma_{bliss}$ so that there is a unique steady state below satiation, then we also get $\hat{b} < b^*$. The Aiyagari-McGrattan exercise *underestimates* the optimal long-run quantity of public debt and also gives a \hat{b} that is necessarily below b_{bliss} in spite of the fact that in this case satiation is optimal in the long run when $\psi_{bliss} < \gamma_{bliss}$. Outside our benchmark, there can exist multiple steady states below b_{bliss} , some of which may be higher than b^{gold} and others lower than it. Furthermore, as shown in Proposition 10, even when the steady state below b_{bliss} is unique, this steady state may correspond to a local minimum of Ω . It follows that ignoring transitional dynamics may fail to give the right answer to the question of the optimal long term level of debt *even when* one modified the steady-state objective to take into account the fact that debt carries a seignorage-like benefit rather than a cost.³¹

The following analogy may help here. The original exercise conducted by Aiyagari and McGrattan is akin to studying the *golden rule* in the Neoclassical Growth Model. The point made here is that even the *modified* golden rule can be misleading.

³⁰To simplify, we assume that the solution is unique, which is true for generic (U, V, π) .

³¹The fact that public debt in Aiyagari and McGrattan also serves as an instrument for redistribution in addition to its role as private collateral does not matter for the essence of our criticism.

Crowding out, or crowding in, capital. A central aspect of Aiyagari and McGrattan (1998) that our analysis has so far abstracted from is the possibility that the issuance of public debt crowds out capital accumulation. This happens because, in that model, public debt is a substitute for physical capital (or for financial claims on such capital) as a buffer stock.

This channel is accommodated in the variant model we study in Appendix A. In that model, private agents can relax future financial constraints by saving, not only in government bonds, but also in physical capital. It follows that the issuance of public debt *can* crowd out capital accumulation by reducing the need to use physical capital as a form of collateral.

We emphasize “can” because our model also contains an additional channel, which pulls in the opposite direction. Because the financial constraint impedes the efficient allocation of capital across different productive opportunities, and because public debt serves as collateral and thus mitigates the financial constraint, an increase in the quantity of public debt can improve the allocation of capital, thereby also leading to an increase in the ex ante return to capital. This effect, which crowds *in* capital, is absent in Aiyagari and McGrattan (1998), because in their model the production side of the economy is first-best efficient; but it is at the core of Holmström and Tirole (1998) and, more generally, of the literature that emphasizes how financial frictions reduce aggregate TFP and discourage investment. One can think of this effect as a positive effect on the demand for investment, and the one emphasized by Aiyagari and McGrattan as a negative effect on the supply of saving. The overall effect then depends on which of these conflicting effects dominates.

Despite these ambiguities, the policy problem can still be understood in terms of the dual role that the financial friction has on welfare and the government budget. The crowding in or out of capital, and the endogeneity of aggregate TFP and income to the stock of public debt represent interesting questions on their own right, and will certainly matter for quantitative purposes, but are not essential for the qualitative insights we have produced in this paper.

When is it cheap for the government to borrow? In recent policy debates, low interest rates have been often interpreted as a signal that it is “cheap” for the government to borrow. Among others, Krugman and DeLong have argued that the US government should have run an expansionary fiscal policy during the Great Recession, not only for Keynesian-stimulus reasons, but also because interest rates were extraordinarily low.

This claim is in general misleading. Suppose we added deterministic variation in discount rates in the Barro-AMSS version of our model, or preference and technology shocks in the stochastic model of Lucas and Stokey (1983). We could then obtain arbitrary movements in interest rates, either across time or across states, without any change in the optimal mixture of taxes and debt. In such cases, low interest rates do *not* justify more borrowing. The reason is that the equilibrium interest rate captures not only the cost of borrowing but also what a society desires.

Nonetheless, our analysis helps identify situations in which the aforementioned claim may con-

tain the right policy recommendation. More borrowing may be optimal during a crisis, not simply because the interest rate on public debt is low, but rather because the low interest rate is a manifestation of the aggravated financial friction and of the associated increase in the liquidity premium that the government can extract from the market. We elaborate further on this point in the discussion of the optimal policy response to financial shocks in the next section.

On Ricardian Equivalence. In our paper, as in Woodford (1990), Aiyagari and McGrattan (1998), and Holmström and Tirole (1998), the financial friction comes together with a departure from Ricardian Equivalence. It is possible, however, to restore the neutrality of public debt without eliminating the financial friction altogether.

Consider a modification of our model that lets the private sector’s pledgeable income move one-to-one with future tax obligations. This modification preserves the financial friction but eliminates the effect of public debt on allocations, interest rates, and welfare: the solution to the sub-problem defined in (8)-(16) becomes invariant to the level of public debt, so that the functions V and π no longer depend on b . 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 these papers as well.

This, however, does not mean that the economy reduces to the Barro-AMSS benchmark: the interest rate is still depressed relative to the planner’s discount rate. Accordingly, the Euler condition gives $\dot{\lambda} = \pi\lambda > 0$, where $\pi > 0$ is the—now fixed—wedge between the interest rate and the planner’s discount rate. The following result then obtains.

Proposition 11. *Suppose that the friction is present but public debt is neutral. Then, optimal taxes and debt exhibit a positive drift. In the long run, debt converges to \bar{b} , the highest sustainable level.*

Insofar as π remains invariant to b , the above result continues to apply even if we let V increase with b (for some unmodeled reason). This clarifies how the lessons obtained in this paper depend on the causal effect of public debt on interest rates—a causal effect that is corroborated by the evidence in Krishnamurthy and Vissing-Jorgensen (2012). If it were not for this effect, the optimal quantity of public debt would increase without bound.

7 Wars, Recessions, and Financial Crises

We now turn to the characterization of the optimal policy response to various shocks.³² As in AMSS, the stochastic policy problem can only be solved numerically. Unlike AMSS, however, the local stability of any steady-state point below satiation guarantees that, as long as we limit ourselves


³²To this goal, we continue to assume that the only asset traded between the government and the private sector is the risk-free bond. This keeps the analysis directly comparable to Barro and AMSS. At the end of the section, we discuss an alternative that allows for state-contingent debt and builds a bridge to Lucas and Stokey (1983).

to the study of small shocks in the neighborhood of such a steady state, we can analyze the optimal policy response by studying the transitional dynamics of the deterministic model. What is more, the analysis of Section 5 guarantees that, at least as long as we restrict attention to small shocks, it makes no difference whether we conduct the analysis under the benchmark case in which there is a unique steady state below satiation or in the more general case in which there may exist multiple such steady states. In what follows, we therefore alternate between two complementary approaches: we build intuition using the phase diagram of the deterministic model; we numerically illustrate the impulse responses of the stochastic economy.

Fiscal shocks. Consider the comparative dynamics of an unexpected, once and for all, increase in g . This is illustrated in Figure 5. Prior to the change, the economy is assumed to be resting at the steady-state point b_{old}^* . The increase in g causes the $\dot{b} = 0$ locus to shift upwards, reflecting the increase in the taxes required for balanced budget. By contrast, the $\dot{\lambda} = 0$ locus does not move, because g does not enter the planner's Euler condition. It follows that the steady-state level of debt drops from b_{old}^* to b_{new}^* . The optimal dynamic response is therefore as follows: on impact, λ and the associated tax rate jump up from their old steady-state values to values that set the economy on a new saddle path; thereafter, debt and tax monotonically decrease towards the new steady state, so that the government runs a surplus along the transition.

To understand the logic of this policy response, compare it with that obtained in the Barro/AMSS benchmark, that is, in the absence of the financial friction. In that benchmark, the optimal response to an unexpected, once and for all, change in government spending is to restore budget balance and keep the level of debt constant for ever. Relative to this response, the optimal response in our model is to raise taxes even more in the short run, but to lower them in the long run. In effect, the policy maker deviates from tax smoothing in order to reduce liquidity, raise the premium, and thereby partially offset the higher tax burden.

Figure 5: Permanent Increase in Government Spending



The same logic underlies the optimal response to transitory fiscal shocks, what is often referred to in the literature as “wars”. We illustrate this in Figure 6, using the impulse response functions of a stochastic example (instead of the deterministic phase diagram). In this example, government spending is assumed to follow a two-state Markov process, with values g_p in the “peace” state and g_w in the “war” state ($g_w > g_p$), and with transition matrix given by

$$\begin{bmatrix} q_p & 1 - q_p \\ 1 - q_w & q_w \end{bmatrix}$$

where q_p and q_w are the probabilities of staying in the same state.

Figure 6 shows the impulse response of the debt level and the tax rate to the realization of a war, both in a numerical version of our baseline model (blue lines) and, for comparison, in its Barro/AMSS counterpart (red lines). In the top panel, the war is assumed to last only one period ($q_p = 1 - q_w = 0.9$, implying that a 1 period war occurs every 10 periods on average); in the bottom panel, the war is moderately persistent ($q_p = q_w = 0.9$). In both cases, Barro/AMSS predict that the war leaves a permanent mark on the level of debt and the rate of taxation, reflecting the unit-root property of that benchmark. Furthermore, the size of the tax response is simply the change in the annuity value of government spending. In our setting, by contrast, the debt level eventually reverts to its initial position, reflecting the determinacy of the long-run target level of debt. Furthermore, the accumulation of debt during the war is less pronounced than that in Barro/AMSS, because doing so permits the planner to moderate the reduction in the liquidity premium, which would have further tightened the budget.³³ By the same token, the planner raises taxes and runs smaller deficits during the war, but also enjoys lower taxes in the aftermath of the war.

Productivity shocks. Consider an adverse productivity shock. In the Barro/AMSS benchmark, the optimal policy response to such a shock is analogous to that of a “war”, because both shocks boil down to tightening the government budget. In our context, this analogy is preserved only as long as the productivity shock does not affect the severity of the financial friction. This, in turn, is true in our baseline micro-founded model only when $\xi = 0$, that is, when pledgeable income is zero. When, instead, $\xi > 0$, pledgeable income covaries with productivity, implying that the collateral constraint hardens during a recession. How this new effect reshapes the optimal policy response is discussed next.


Financial recessions. Consider an exogenous shock that causes income to fall, the tax basis to shrink, and both the private and the social value of liquidity to increase. As noted above, such a shock could be a negative productivity shock in our baseline model as long as $\xi > 0$. Alternatively, it could be a negative shock to ξ itself, that is, an exogenous tightening of the collateral constraint.³⁴

³³If the war is sufficiently persistent, this mechanism becomes so strong that the level of debt actually falls, as in the example with the permanent change discussed earlier.


³⁴To be precise, in our baseline model, a shock to ξ affects the private and social value of liquidity, but does not

Figure 6: Reponse to a War Shock (AMSS)

(a) i.i.d. Case



(b) Persistent Case



— Debt and Taxes in our Model; — Debt and Taxes in AMSS; - - - Government Spending.

More generally, we use the aforementioned shock to proxy for “financial recessions”, that is, for recessions featuring an aggravated financial friction and a higher-than-usual demand for collateral.³⁵

By raising both the private and the social value of liquidity, the shock can have an ambiguous effect on the ratio v/π . It may also have an ambiguous effect on σ . In the absence of any strong priors on the relative size of these effects, we focus on the scenario in which both the wedge between v and π and the elasticity σ remain constant. This scenario has two advantages: it keeps the analysis close to the literature on the Friedman rule, where v/π is *a fortiori* fixed at 1; and it leads to a sharp characterization of the optimal response.³⁶

Under this scenario, the $\dot{\lambda} = 0$ locus does not shift in response to the shock. This reflects the offsetting effects that the shock has on the social value of liquidity and on the planner’s opportunity cost of supplying liquidity: the aggravated friction presents the planner with an opportunity to extract more seigniorage at the margin, which in turn negates the apparent increase in the social value of easing the friction. It follows that the optimal policy response is dictated by the shift in the \dot{b} locus, that is, by the impact of the shock on the budget constraint.

By shrinking the tax basis, the shock causes the $\dot{b} = 0$ locus to shift upwards, reflecting an increase in the tax burden. If that were the only effect on the budget constraint, the optimal policy response would be analogous to the case of an adverse fiscal shock, which has already been analyzed. There is, however, an important offsetting effect. By raising the liquidity premium π , the shock also raises the implicit seigniorage that the government enjoys in terms of a lower cost of borrowing. This increase helps mitigate, and may even overturn, the reduction in tax revenue.

Depending on which effect dominates, the $\dot{b} = 0$ locus could therefore shift either way. If the interest-rate effect dominates, a financial recession looks, paradoxically, more like a traditional boom: the $\dot{b} = 0$ locus shifts down, meaning that the tax burden *falls*. But even if the tax-revenue effect dominates, as long as the interest rate falls, the tax burden increases *less* than otherwise. It follows that the optimal policy response to a financial recession features necessarily a smaller increase in the tax rate, and hence also a larger increase in the primary deficit, than a *comparable* traditional recession, by which we mean a recession that involves the same reduction in tax revenue but no exogenous movement in the liquidity premium and the interest rate.


affect income and the tax basis; this is because the financial friction affects only the distribution of consumption. In the variant of Appendix A, on the other hand, a negative shock to ξ affects, not only the value of liquidity, but also income and tax revenue; this is because the friction distorts the allocation of capital and reduces aggregate TFP.

³⁵To map this to our reduced-form policy problem, let z denote the shock; let $S(\tau, z)$, $\pi(b, z)$, and $v(b, z)$ denote the tax revenue, the liquidity premium, and the marginal social value of liquidity as functions of the shock; and finally let S increase with z , and π and v decrease with it, so that a “financial recession” corresponds to a low realization for z . As in our analysis of “wars”, we first develop intuition by treating z as a parameter in the phase diagram of the deterministic model, and then use a numerical example to illustrate the stochastic extension.

³⁶This scenario, however, abstracts from the realistic possibility that fire-sale externalities are exacerbated during a financial crisis; see footnote 38 for a discussion of how this possibility influences the optimal policy response.

Figure 7 uses a numerical example to illustrate how the above logic, which is based on the phase diagram of the deterministic model, shapes also the IRFs of the stochastic model. The orange lines correspond to the case in which the shock does not have a direct effect on the interest rate.³⁷ The black lines, instead, correspond to the knife-edge case in which the interest-rate effect negates the tax-revenue effect exactly. More specifically, our example imposes that the increase in $\pi(b)b$ equals the reduction in $S(\tau)$ when b and τ equal to their steady-state values, as well as that v increases in proportion to π . These assumptions guarantee that, following the shock, the planner experiences no change in either the budget constraint or the Euler condition, and therefore she finds it optimal to keep both τ and b constant at these values. It follows that the *entire* drop in tax revenue during the recession is debt-financed. And yet, the higher deficits do not translate to a higher debt burden over time, because the government is now able to “refinance” its original debt at lower interest rates.

Figure 7: IRF to a Financial Shock



What if the interest-rate effect is present but weaker than the tax-revenue effect? Not surprisingly, the optimal policy response is simply in-between the two cases in the figure. The optimal deficit is lower than in the knife-edge scenario studied above, but higher than in a comparable traditional recession. By the same token, both the tax rate and the debt burden increase during the recession, but less so than in the absence of the interest-rate effect.³⁸

These findings provide a formal basis for the argument—voiced by Krugman, DeLong, and

³⁷The shock is nevertheless allowed to have an *indirect* effect, insofar as the endogenous increase in the stock of public debt is accompanied by an endogenous increase in the interest rate.

³⁸As noted before, these findings presume that v/π is invariant to the shock. We now discuss a plausible relaxation of this assumption. Suppose v increases less than π during a financial recession, causing the aforementioned locus to shift down. This could reflect the exaggeration of fire-sale or other externalities that drive a wedge between the private and the social value of liquidity. Other things equal, this raises the planner’s incentive to extract seigniorage, which in turn tends to lower the optimal deficit level relative to the scenario in which v increases as much as π . No to misinterpret this result, that we kept the variation in π constant and lowered the variation in v . If, instead, we keep the variation in v constant and raise the variation in π , then we get a higher optimal deficit, because we have a stronger interest-rate effect.

others—that the reduction in the government’s cost of borrowing during a financial crisis is a signal of the optimality of running larger deficits. But it is important to emphasize the part of the statement that says “during a financial crisis”: as emphasized in the previous section, 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 shadow discount rate of the planner.

Allowing for state-contingent debt. We conclude this section by discussing 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.

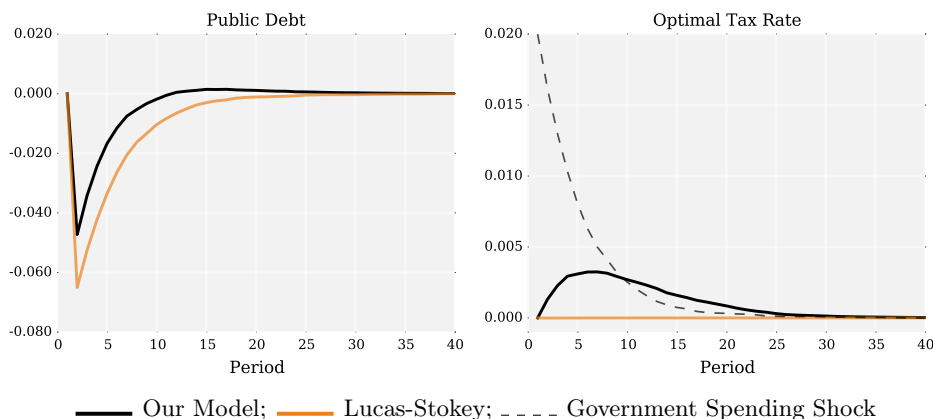
The answer differs in our setting. When state-contingent debt is available, our planner maintains the option to equate the shadow cost of taxation across different histories of shocks, exactly as in Lucas and Stokey (1983). But unlike that environment, the planner no longer finds it optimal to do so. 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 8 using a persistent war. This is exactly the same as in the bottom of Figure 6, 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 8: 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.

8 Concluding remarks

The Great Recession has rekindled interest on the macroeconomic effects and the policy implications of financial frictions. In the presence of such frictions, debt instruments issued by the government can serve, not only as store of value, but also as collateral or vehicles of private liquidity: they help private agents relax the bite of certain trading frictions, a function that gets reflected in their price. This indicates a formal connection between this new literature and the earlier literature on the Friedman rule: public debt in the former is similar to “outside money” in the latter literature.

This connection has not gone unnoticed.³⁹ Yet, the literature has lacked a theoretical analysis of how optimal fiscal policy is determined when the net indebtedness of the government influences the severity of financial frictions. The present paper makes a contribution towards filling this gap.

Similarly to the growing literature on “quantitative easing”, our analysis touches on how the balance sheet of the government can influence the level of risk-sharing within the private sector. But unlike that literature, we focus on the role played by the net debt position of the government, as opposed to the composition of its assets and liabilities. More importantly, we allow the govern-

³⁹See, for instance, the recent contributions of Stein (2012), Krishnamurthy and Vissing-Jorgensen (2012), and Brunnermeier and Sannikov (2016).

ment to ease financial frictions only at the expense of raising the net debt burden and future tax distortions. This leads to the intertwining of fiscal policy with the government’s role in facilitating private liquidity—an intertwining that is the key to understanding the nature of our results, their novelty vis-a-vis the existing literature, and their likely robustness.

Optimal policy balances three objectives: smoothing taxes; easing financial frictions; and managing the government’s cost of borrowing. In our preferred scenario, the analogue of the Friedman rule does not apply either in the short or in the long run: the financial distortion is preserved in steady state in order to depress the interest rate on public debt. In contrast to the standard Ramsey framework, the optimal dynamics around such a steady state features mean reversion and a departure from tax smoothing. Finally, new insights emerge about the optimal policy response to shocks and about the sense in which a financial crisis makes it cheap for the government to borrow.

As noted in Section 6, our findings are robust to an extension of our reduced-form framework that allows for multiple assets, all of which serve as vehicles of private liquidity. A more valuable exercise, however, would be to embed the trade offs identified in the present paper to multi-asset settings that have transparent, and empirically plausible, micro-foundations of the relevant compositional effects, such as that of Brunnermeier and Sannikov (2014, 2016). We leave this open for future research.

Another interesting direction for future research is to extend our analysis to settings that contain nominal rigidity. As noted in the Introduction, financial frictions have been suggested as an explanation for why the zero-lower bound became binding during the recent recession; but if public debt management can ease these frictions, it can also relax the constraint on monetary policy. This kind of interaction between fiscal and monetary policy remains to be explored.

Appendix A: Variant Model and Reinterpretation

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 $c_{it} - \nu(h_{it})$, 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 baseline model, he can do so by pledging ϕ and/or by posting his financial assets, a_{it} , as collateral. Moreover, he can use a fraction of the invested capital and/or the firm’s output as additional collateral. That is, the relevant constraint is

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

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

Relative to the baseline model, the model described above facilitates a different interpretation of the financial friction. It also features distinct positive implications. First, the financial friction can manifest itself as cross-sectional heterogeneity in the marginal product of capital. Second, by easing the reallocation of capital, higher public debt can lead to higher aggregate productivity, as in Holmström and Tirole (1998). Last but not least, by satisfying part of the economy’s demand for collateral, public debt can crowd out capital accumulation, an effect similar to that 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 12. *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:*

$$\max \sum_{t=0}^{\infty} \beta^t W(\tau_t, b_t) \tag{36}$$

$$\text{subject to } Q(\tau_{t+1}, b_{t+1})b_{t+1} = b_t + g - S(\tau_t, b_t) \tag{37}$$

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.⁴⁰ As we move from the baseline model to the new model,

⁴⁰In the baseline model, because the tax revenue, s , was merely a function of the tax rate only, we found it convenient to express the problem directly in terms of s . In the current model, this mapping is not as direct and, accordingly, we keep the tax rate as a control variable.

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

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

$$\max \int_0^{+\infty} e^{-\rho t} W(\tau, b) dt \quad (38)$$

$$\text{subject to } \dot{b} = [\rho - \pi(\tau, b)]b + g - S(\tau, b) \quad \forall t \quad (39)$$

$$b(0) = b_0 \quad (40)$$

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 $\pi(\tau, b) = W_b(b, \tau) = 0$ 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 $b_{seig}(\tau) \equiv \arg \max \{\pi(\tau, b)b + S(\tau, b)\}$; this is the analogue to the level of debt that maximized seigniorage in our baseline model, except that now we accommodate the possibility that the quantity of aggregate collateral affects the government budget, not only via the interest rate on public debt, but also via aggregate output and tax revenue. Adjusting the notion of “liquidity plus seigniorage” accordingly gives

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

We can express the planner’s Euler condition as

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

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


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

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

We illustrate this in Figure 9. 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. The boundaries of the regions are now curved, rather than vertical, reflecting the fact that b_{seig} 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. In the example under consideration, there happens to be 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 analysis can accommodate different micro-foundations of the

Figure 9: Entrepreneurial Model



financial friction and the role of public debt. Moreover, our results are robust to changing the specification to allow for the possibility that the issuance of public debt can crowd out private collateral.

Appendix B: Additional Results

B.1 Private versus social value of liquidity

Consider the micro-founded model of Section 2. Let $a(\theta, b)$ and $P(b)$ denote the allocation of the bond and the price of the afternoon good that obtains from solving the planning sub-problem (8)-(16) of Section 4 and, to simplify, let $a(\theta, b) > -\phi$ for all θ . From the definition of $V(\cdot)$ together with the fact that the aggregate net trade of the morning good is zero in equilibrium, we have that

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

and therefore

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

From the household's Euler condition (15), we have that $Q(b) = \beta + \mathcal{U}_a(\cdot)$ for all θ . From the bond market clearing condition (Equation 10), we have $\int a_b(\theta, b) \varphi(\theta) d\theta = 1$. Using these last two relations, we infer that


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

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

B.2 Additional Examples and Figures for Section 5

In Figure 3 in the main text, γ and ψ intersect three times. As noted there, this case imposes a certain structure, not only for the local dynamics around the three steady states, but also about the global dynamics. We now consider two alternative possibilities for the global dynamics, while maintaining the same configuration for the functions γ and ψ , the same steady states, and the same local dynamics around them.

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




Consider Figure 10. Unlike Figure 3, the stable arm corresponding to the highest steady state no longer cycles back to $b_M^\#$; instead, it extends past $b_L^\#$. This has the following important implication. If we consider $b_0 = b_L^\#$, then there are two candidate optimal plans, namely the plan of staying put at $b_L^\#$ and the plan that leads to $b_H^\#$. The former plan is dominated because it features $\mathcal{H}_\lambda = \dot{b} = 0$, whereas the latter features $\mathcal{H}_\lambda = \dot{b} > 0$. By continuity, the saddle path that leads to $b_L^\#$ is dominated also for any b_0 in an open neighborhood of $b_L^\#$. But then the path leading to $b_L^\#$ can *never* be 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 3, the path that leads to $b_H^\#$ is now the optimal path for all $b_0 < b_{skiba}$.

Figure 11 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 3 is the following: even though the ODE system continues to admit multiple saddle-path stable steady states, the optimal policy now features a unique and globally stable steady state in the region to the left of the satiation point, that is, optimal debt converges monotonically to the *same* long run value b^* for all initial values $b_0 \leq b_{bliss}$.

Let us now allow for the possibility that $\pi(b)b$ is not single-peaked. Suppose that $\pi(b)b$ has N local extrema, denoted by $\{b_1, b_2, b_3, \dots, b_N\}$, with $\underline{b} < b_1 < b_2 < \dots < b_N < b_{bliss}$, where N is an arbitrary finite number. First, note that $\sigma(b)$ crosses 1 whenever b crosses any of these points. Next, note that the last point, namely b_N , is necessarily a local maximum, because after that point $\pi(b)b$ falls to zero as b approaches b_{bliss} . It follows that $\sigma(b)$ is higher than 1 when $b \in (b_N, b_{bliss})$, lower than 1 when $b \in (b_{N-1}, b_N)$, higher than 1 when $b \in (b_{N-2}, b_{N-1})$, and so on. By the

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



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. The phase diagram now looks like the product of patching together multiple pairs of L and M regions from our earlier examples: within any of the regions in which $\sigma(b) > 1$, we can have an arbitrary number of intersections between γ and ψ , and a plethora of possibilities for dynamics.

We finally revisit the role of the condition $\gamma_{bliss} < \psi_{bliss}$. As shown in the main text, this condition is sufficient for satiation not to obtain in the long run, not only in our benchmark, but also in general. We now show, by example, that this condition is not necessary away from our benchmark.

Consider Figure 13. As in our benchmark (see Figure 1 in particular), letting $\gamma_{bliss} > \psi_{bliss}$ guarantees the *local* existence of a candidate path that leads to satiation: for some $\epsilon > 0$ and all $b_0 \in (b_0 - \epsilon, b_{bliss})$, there exists a path that satisfies the ODEs at all dates and that asymptotically converges to b_{bliss} . But unlike what was true in our benchmark, this type of path does not exist for sufficiently low b_0 . What is more, for all $b_0 < b_{bliss}$, there happens to exist another candidate optimal path, namely the one that leads to a steady state below b_{bliss} . Finally, note that the path leading to b_{bliss} features an initial value for \dot{b} that is arbitrarily close to 0 when b_0 is close enough to b_{bliss} , whereas the path leading to $b_L^\#$ features a \dot{b} bounded away from zero. Using once again Lemma 4, the convexity of \mathcal{H} in λ , and the fact that $\mathcal{H}_\lambda = \dot{b}$, we infer that the latter path dominates the former for b_0 in a neighborhood of b_{bliss} . But this also means that the path leading to satiation can not be optimal for *any* initial b_0 . Instead, there again exists a $b_{skiba} > b_{bliss}$ such that for all $b_0 < b_{skiba}$ it is optimal to converge either to $b_L^\#$ or to some point further below.

Figure 12: Multiple Regions





Figure 13: No Satiation Despite $\psi_{bliss} < \gamma_{bliss}$



Appendix C: Proofs

Proof of Proposition 1. Let us start from the initial problem

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

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

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

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

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

where $\phi \equiv \xi A (\nu')^{-1}((1 - \xi)A)$ and $z_{it} = p_t (x_{it} - \bar{e})$. Let us define $\tilde{c}_{it} = c_{it} + z_{it}$. We then have

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

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

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

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

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

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

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

Q.E.D. ■

Proof of Proposition 2. The proposition is a direct consequence of the discussion that precedes it and from the definitions of the functions $V(b)$, Q , U and S .

Q.E.D. ■

Proof of Lemma 1. Proved in the main text.

Proof of Lemma 2. Recall that $b_{seig} = \arg \max_b \pi(b)b$, so that b_{seig} solves $\pi(b)(1 - \sigma(b)) = 0$. Note that, for b_{seig} to be a maximum, the following has to hold: $\pi(b)(1 - \sigma(b)) \geq 0$ for $b \leq b_{seig}$. From the definition of γ and the assumption $V'(b) \propto \pi(b)$, we have

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

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

Q.E.D. ■

Proof of Lemma 3. From Lemma 2, we know that $\psi(b)$ and $\gamma(b)$ can only intersect in (b_{seig}, b_{bliss}) . Given (i) the monotonicity of $\sigma(b)$ and hence $\gamma(b)$, (ii) the fact that $\psi(b)$ is increasing and (iii) $\lim_{b \downarrow b_{seig}} \gamma(b) = \infty$, $\gamma(b)$ 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 $(b_{seig}, b_{bliss}]$ and therefore they never intersect. In $\gamma_{bliss} < \psi_{bliss}$, (i)–(iii) imply that they intersect only once.

Q.E.D. ■

Proof of Proposition 3. Let us first consider $b_0 \geq b_{bliss}$. In this case, $V'(b) = \pi(b) = 0$ and the ODE system reduces to

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Q.E.D. ■

Proof of Lemma 4. See Brock and Dechert (1983).

Proof of Proposition 4. 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 $(b^\#, \lambda^\#)$.

The linear approximation of the system of the ODEs around a stationary point $(b^\#, \lambda^\#)$ 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 $(b^\#, \lambda^\#)$, 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 $(b^\#, \lambda^\#)$, 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 $b \in (b_{seig}, b_{bliss}]$. Therefore, given that $\psi'(b) > 0$, $\det(\mathbf{J}) < 0$ and hence the two eigenvalues are distributed around 0. Therefore, $(b^\#, \lambda^\#)$ 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 $(b^\#, \lambda^\#)$ increasing (resp. decreasing) both debt and taxes along the transition path.

Q.E.D. ■

Proof of Proposition 5. Proved in the main text.

Proof of Proposition 6. The linear approximation of the system of the ODEs around a stationary point $(b^\#, \lambda^\#)$ is given by

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

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

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

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

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

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

The two roots are complex if the discriminant is negative

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

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

Q.E.D. ■

Proof of Proposition 7. The proof is given in the main text.

Proof of Theorem 2. We prove this result with the help of Theorem 2 from Brock and Dechert (1983). Consider the optimal policy rule for the co-state variable, namely the correspondence from any given b_0 to the *optimal* value for λ_0 . Denote this correspondence by Λ^{opt} . Note that this is a selection from the correspondence Λ (which was defined in the context of Lemma 4). To illustrate, consider Figure 3. 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 $\lim_{b \uparrow \hat{b}} \Lambda^{opt}(b) > \lim_{b \downarrow \hat{b}} \Lambda^{opt}(b)$.⁴¹ In

⁴¹At 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.

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

Thus suppose there exists an initial point $b_0 = \tilde{b}_0$ such that it is optimal to converge to a point $b^* < b_{bliss}$. Clearly, b^* must be inside $B^\#$. Next, consider the set of points at which the policy rule of the co-state features a discontinuity and let \hat{b} be the highest such point below b^* ; if no such point exists, just let $\hat{b} = \underline{b}$. When $b_0 \in (\hat{b}, \tilde{b}_0)$, debt converges to b^* . When instead $b_0 < \hat{b}$ (which, of course, is relevant only insofar as $\hat{b} > \underline{b}$), debt converges to a point that is below \hat{b} , and hence also below b^* , but still inside $B^\#$. It follows that there exists a point $b_{skiba} \geq b^*$ such that, when $b_0 < b_{skiba}$, 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 7 that $\psi_{bliss} > \gamma_{bliss}$ is sufficient for $b_{skiba} > b_{bliss}$. It follows that $b_{skiba} < b_{bliss}$ is possible only insofar as $\psi_{bliss} < \gamma_{bliss}$, which in turn guarantees the existence of a candidate path that converges to b_{bliss} for any $b_0 \in [\hat{b}, b_{bliss})$ and some $\hat{b} < b_{bliss}$. Clearly, $\hat{b} \leq b_{skiba}$. By definition of b_{skiba} , the optimal path is one of the candidate paths that converge to a point inside $B^\#$ if and only if $b_0 < b_{skiba}$. Therefore, for any $b_0 \in [b_{skiba}, b_{bliss})$, either the path that leads to b_{bliss} is the unique candidate path, or it dominates any of the candidate paths that lead to a point inside $B^\#$.

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

Q.E.D. ■

Proof of Theorem 3. That any economy must belong to one of these three classes follows from Theorem 2. That the first and the third classes are not empty follows from the examples we have already considered; an example of the second class was provided in the Appendix. Finally, the claimed sufficiency of the condition $\psi_{bliss} > \gamma_{bliss}$ follows from Proposition 7.

Q.E.D. ■

Proof of Lemma 5. Note that ψ_{bliss} is continuous and increasing in g as long as $g < g_{max}$ and diverges to $+\infty$ as $g \rightarrow g_{max}$. This is because a higher g requires higher taxes to balance the budget, and the marginal cost of these taxes explodes to infinity as we approach the peak of the Laffer curve. Furthermore, $\psi_{bliss} = 0$ if and only if $g = -\rho b_{bliss} < 0$. Finally, note that γ_{bliss} : is invariant to g ; is positive for the reasons offered above; and is finite by assumption. It then follows that there exists a threshold \hat{g} , necessarily less than g_{max} and possibly negative, such that $\psi_{bliss} > \gamma_{bliss}$ if and only if $g > \hat{g}$.

Q.E.D. ■

Proof of Proposition 8. By the definition of $b^* \in B^*$ and $b^* < b_{bliss}$, we know that the point $(b^*, \lambda^* \equiv \psi(b^*st))$ is locally stable. Similarly to Proposition 6, 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 6, that the eigenvalue associated with the stable arm is given by $\mu = \frac{\rho - \sqrt{\Delta}}{2}$ with $\Delta > 0$ (see proof of Proposition 6). It is then straightforward to obtain the eigenvector $\mathbf{v} = (v_1, v_2)$ satisfying

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

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

Q.E.D. ■

Proof of Proposition 9. Any $b^* \in B^*$ is such that $\gamma(b^*) = \phi(b^*)$. Therefore, it inherits the comparative statics of the γ and ϕ functions. The comparative statics are therefore identical to those already analyzed in proposition 5. Q.E.D. ■

Proof of Proposition 10. The proof is given in the main text.

Proof of Proposition 12. Let start with the entrepreneur. He chooses his production plan by solving the problem

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

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

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

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

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

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

The problem of the household is

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

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

$$\int a_{it} di = b_t$$

define aggregate capital as

$$\int \kappa_{it} di = \kappa_t$$

and define

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

The problem of the representative household is then

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

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

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

$$q_t = \beta + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1}; \theta') \tag{44}$$

$$1 = p_{t+1} (\beta + \Omega_x(x_{t+1}, p_{t+1}, w_{t+1}; \theta')) \tag{45}$$

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 $\Omega_x(\cdot) > 0$.

The labor market clearing condition ($h_t = n_t$), and the capital market clearing condition ($k_t = \kappa_t$) imply

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

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

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

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

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

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

$$q_t = Q(\tau_{t+1}, b_{t+1}; \theta')$$

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

and the government budget is

$$Q(\tau_{t+1}, b_{t+1}; \theta') b_{t+1} = b_t - S(\tau_t, b_t)$$

Hence, the problem of the central planner reduces to

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

subject to $Q(\tau_{t+1}, b_{t+1}; \theta') b_{t+1} = b_t - S(\tau_t, b_t)$

Q.E.D. ■

Proof of Proposition 11. The proposition is a direct consequence of the fact that $\dot{\lambda} = \bar{\pi}\lambda > 0$.

References

- Aiyagari, S. Rao, Albert Marcet, Thomas Sargent, and Jesus Seppälä.** (2002) “Optimal Taxation without State-Contingent Debt.” *Journal of Political Economy*, 110(6): 1220–1254.
- Aiyagari, S. Rao, and Ellen McGrattan.** (1998) “The Optimum Quantity of Debt.” *Journal of Monetary Economics*, 42(3): 447–469.
- Barro, Robert J.** (1979) “On the Determination of the Public Debt.” *Journal of Political Economy*, 87(5): 940–971.
- Brock, William A., and W.D. Dechert.** (1983) “The Generalized Maximum Principle.” University of Wisconsin mimeo.
- Brunnermeier, Markus K., and Yuliy Sannikov.** (2014) “A Macroeconomic Model with a Financial Sector.” *American Economic Review*, 104(2): 379–421.
- Brunnermeier, Markus K., and Yuliy Sannikov.** (2016) “The I-Theory of Money.” Princeton University mimeo.
- Buera, Francisco J.** (2008) “Persistency of Poverty, Financial Frictions, and Entrepreneurship.” Northwestern University mimeo.
- Chari, V.V., Lawrence Christiano, and Patrick Kehoe.** (1996) “Optimality of the Friedman Rule in Economies with Distorting Taxes.” *Journal of Monetary Economics*, 37(2): 203–223.
- Correia, Isabel, and Pedro Teles.** (1999) “The Optimal Inflation Tax.” *Review of Economic Dynamics*, 2(2): 325–346.
- Cúrdia, Vasco, and Michael Woodford.** (2011) “The central-bank balance sheet as an instrument of monetary policy.” *Journal of Monetary Economics*, 58(1): 54–79.
- Dávila, Eduardo.** (2015) “Dissecting Fire Sales Externalities.” New York University, Stern School of Business mimeo.
- Dechert, W.D., and K. Nishimura.** (1981) “A Note on Optimal Growth with a Convex-Concave Production Function.” State University of New York at Buffalo Department of Economics Discussion Papers.
- Eggertsson, Gauti, and Paul Krugman.** (2012) “Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo approach.” *Quarterly Journal of Economics*, 127(3): 1469–1513.

- Gertler, Mark, and Nobuhiro Kiyotaki.** (2010) “Financial Intermediation and Credit Policy in Business Cycle Analysis.” In *Handbook of Monetary Economics*. Vol. 3 of *Handbook of Monetary Economics*, , ed. Benjamin M. Friedman and Michael Woodford, Chapter 11, 547–599. Elsevier.
- Gertler, Mark, and Peter Karadi.** (2011) “A model of unconventional monetary policy.” *Journal of Monetary Economics*, 58(1): 17–34.
- Guerrieri, Veronica, and Guido Lorenzoni.** (2009) “Liquidity and Trading Dynamics.” *Econometrica*, 77(6): 1751–1790.
- Guerrieri, Veronica, and Guido Lorenzoni.** (2011) “Credit Crises, Precautionary Savings, and the Liquidity Trap.” NBER Working Paper 17583.
- Holmström, Bengt, and Jean Tirole.** (1998) “Private and Public Supply of Liquidity.” *Journal of Political Economy*, 106(1): 1–40.
- Kiyotaki, Nobuhiro, and John Moore.** (1997) “Credit Cycles.” *Journal of Political Economy*, 105(2): 211–248.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** (2012) “The Aggregate Demand for Treasury Debt.” *Journal of Political Economy*, 120(2): 233–267.
- Lagos, Ricardo, and Randall Wright.** (2005) “A Unified Framework for Monetary Theory and Policy Analysis.” *Journal of Political Economy*, 113(3): 463–484.
- Lorenzoni, Guido.** (2008) “Inefficient Credit Booms.” *Review of Economic Studies*, 75(3): 809–833.
- Lucas, Robert E., Jr., and Nancy L. Stokey.** (1983) “Optimal Fiscal and Monetary Policy in an Economy without Capital.” *Journal of Monetary Economics*, 12(1): 55–93.
- Shleifer, A., and R.W. Vishny.** (1992) “Liquidation Values and Debt Capacity: A Market Equilibrium Approach.” *Journal of Finance*, 47(4): 1343–1366.
- Skiba, A. K.** (1978) “Optimal Growth with a Convex-Concave Production Function.” *Econometrica*, 46(3): 527–539.
- Stein, J.C.** (2012) “Monetary Policy as Financial Stability Regulation.” *Quarterly Journal of Economics*, 127(1): 57–95.
- Woodford, Michael.** (1990) “Public Debt as Private Liquidity.” *American Economic Review*, 80(2): 382–388.