

Optimal fiscal and monetary policy with distorting taxes

By CHRISTOPHER A. SIMS*

When government debt pays a lower return than private assets, the reasoning in Friedman’s essay on the optimal quantity of money suggests that it would be optimal to expand the debt until its return matched that on private assets. When the only other source of revenue is a distorting tax, however, this is not generally true. In a perfect foresight model, a benevolent government that can make credible commitments chooses a large gap in returns initially and high distorting taxation in the distant future. The optimal path of taxation is time-inconsistent, with ever-increasing temptation to abandon the path.

Blanchard (2019) and Mehrotra and Sergeyev (2021) have reminded us that real rates of return on government debt have in most years been so low that the debt-to-GDP ratio would decline or remain stable even if the debt were simply rolled over. That is, without any taxation to “back” the debt, with interest and principal payments being financed entirely by the issue of new debt, the debt-to-GDP ratio would not increase. They characterize this situation as one in which there is zero or negative “fiscal cost” to public debt.

The possibility that when debt pays a lower return than other assets permanent deficits are possible has been explored by Reis (2021) and by Mian, Straub and Sufi (2022), who do not include distorting taxes.

Conventional thinking about public debt sees debt as requiring fiscal backing, so that increased debt requires increased future taxes or reduced future expenditures. If the taxes are distorting, this is a burden. But if we can increase expenditures on beneficial programs, or reduce distorting taxes, without any need to offset the resulting deficits with future increased taxes or reduced expenditures, it appears it would be irresponsible not to exploit this possibility.

There is another approach to thinking about optimal fiscal policy that reaches an apparently opposite conclusion, however. Friedman (1969) argued that optimal monetary policy should set the nominal interest rate to zero, thereby making the real return on money match that on other assets. If “money” pays interest, Friedman’s prescription is that the interest rate on money should match that on other assets. In many models, this requires taxation, either to pay explicit interest or to contract the supply of outstanding government liabilities. So Friedman’s prescription becomes “increase government surpluses to raise the rate of return

* Department of Economics, Princeton University, sims@princeton.edu The paper has benefited from comments by participants at several seminar presentations. ©2024 by Christopher A. Sims. This document is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.

on government liabilities, until the rate on government liabilities matches that on other assets". Blanchard's reasoning, on the other hand, might be read as "so long as the rate of return on government liabilities is low enough, increased deficits are desirable."

While there are many papers on optimal fiscal and monetary policy, three of them illustrate the range of conclusions and contradictions in this area. Chari and Kehoe (1999)¹ show that in a model like that in this paper, an optimizing government that makes policy commitments at an initial date, will set the liquidity premium on government debt to zero at all times.² Woodford (1990) reaches the apparently contradictory conclusion that in the same sort of model, the optimal steady state may involve a positive liquidity premium. Chari and Kehoe suggest that Woodford's result arises because he considers only steady state equilibria (i.e. with constant tax rates), while their results apply when instead the government chooses a fully optimal, and non-stationary, path for policy. However, in a paper cited by neither the Chari/Kehoe paper nor the Woodford survey paper, Calvo (1978) showed in a similar model that optimal policy leads to a non-stationary equilibrium in which the liquidity premium does not converge to zero.³ This result arises despite the fact that the government could by appropriate policy induce a stationary equilibrium with a zero liquidity premium.

Chari and Kehoe argue that their result depends crucially on their assumption that balances and consumption enter the utility function separably, as a homothetic aggregate (so utility is $U(\phi(C, m), L)$ with ϕ homothetic). Calvo considers only models in which utility is additively separable in real government debt and consumption, which does not in general satisfy Chari and Kehoe's homotheticity assumption. It might appear, then, that this is why Calvo's result is different.

As we will see, though, this paper shows that the most important distinction between the Chari/Kehoe paper and the other two is that Chari and Kehoe assume the government has the option at "time zero" of backing the entire money stock with interest-bearing private-sector liabilities. This then allows financing the necessary interest payments on debt, or the contraction of the nominal stock, without any need to use distorting taxes. The papers by Woodford and Calvo, like this paper's model, assume instead that the government cannot hold large amounts of privately issued assets.

¹The main argument in the Chari/Kehoe paper follows that in the earlier paper by Chari, Christiano and Kehoe (1996). The earlier paper considered only cash-in-advance models, while the later one covered also money-in-utility-function models, which are equivalent to those in this paper.

²Their model, like the main version of the model in this paper, makes setting the liquidity premium to zero impossible, but it can be set arbitrarily close to zero, and welfare improves as the premium shrinks.

³Calvo actually discusses two models. In the first model, private sector behavior is represented by directly assuming a form for a "demand for money", without deriving it from private sector optimization. The government is assumed to have an objective function depending on aggregate consumption and real balances. It is in this model that he shows optimal policy has high initial inflation, low initial taxes, and convergence to a steady state not satisfying the Friedman rule. His second model instead postulates the same objective function for the government and private agents and derives private sector behavior from dynamic optimization. In this model he shows time inconsistency of optimal policy, but does not derive detailed results about the nature of the optimal policy path.

This paper works with a model that matches the assumptions on homotheticity and separability in Chari/Kehoe, matches their assumption that saturation of demand for liquidity is possible only asymptotically, and considers a possibly non-stationary Ramsey solution. The difference that alters the result is entirely in this paper's assumption that the government cannot hold private sector bonds that pay a premium return over government debt.

This paper's model reproduces Calvo's result that a benevolent government that can make firm commitments to future policy optimally initially runs deficits and generates a large gap between the return on government debt and the discount rate. The optimal policy promises an ever increasing tax rate. This type of result is probably generic. Postponing distorting taxes makes the required amount of taxation rise, but only at the real rate on government debt, which is less than the discount rate. So unless the distortion costs are rising rapidly with the tax rate, discounted utility rises when tax increases are postponed this way. In models like this one (and Calvo's) that assume perfectly flexible prices, optimal policy is likely to generate immediate inflation, while in a sticky-price model the inflation would be delayed. The same mechanism supporting delayed tax increases would nonetheless probably apply.

Angeletos, Collard and Dellas (2023) consider a system close to Calvo's, but provide deeper discussion of potential microfoundations for it. The Angeletos *et al* paper also assumes real debt is taken by the government as given at time 0. It therefore does not consider the possibility of a jump in the price level at the initial date and does not get the result in this paper and Calvo's that optimal initial real debt is uniquely determined.

Krishnamurthy and Vissing-Jorgensen (2012) present evidence that US Treasury securities have historically paid a lower yield than comparable corporate bonds, that this cannot be explained by default risk, and that the yield spread shrinks as the supply of treasury securities increases. They argue that Treasuries behave much like non-interest-bearing high-powered money in these respects.⁴ Our model in this section builds on their observations, treating all debt as providing liquidity services. Of course we are simplifying by having only one kind of debt, with a single interest rate and maturity. The argument of the previous section, supports the idea that an optimizing government would issue only liquidity-providing debt.

This paper works with a simple perfect-foresight, representative agent model in which there is just one type of duration-zero, nominal government debt. Real government debt provides a liquidity or transactions benefit that appears in the private sector budget constraint. There is a single, proportional, tax on labor and a linear production function. Prices are perfectly flexible. The government cannot acquire private sector assets that pay a higher rate than that on government debt.

⁴Angeletos, Collard and Dellas (2023) argue that if corporate bonds provide some liquidity services, the gap that Krishnamurthy and Vissing-Jorgenson measure, between corporate and government yields, does not capture the full value of liquidity services of government debt.

We pair this model of the private sector with three assumptions about government tax policy. First we consider a government that can credibly commit to an arbitrary, infinitely long, time path of the tax rate τ , which it announces at the initial date t_0 . Then we consider a form of limited commitment, in which the government can announce an arbitrary path for τ from an initial date t_0 until a later date $t_0 + T$, at which point a new government takes over for a new T -period term of office, etc. Finally, we consider results when the government can commit only to constant time paths for τ .

I. Simple model

A representative individual chooses the time paths of C (consumption), L (labor), A (nominal government debt) and B (government debt with possibly lower interest rate) while taking P (the price level) and τ (the tax rate on labor) paths as given (and known in advance). The individual's objective is to maximize

$$(1) \quad \int_0^{\infty} e^{-\beta t} (\log C_t - L_t) dt$$

subject to

$$(2) \quad C \cdot (1 + f(v)) + \frac{\dot{B} + \dot{A}}{P} = (1 - \tau)L - \phi + \frac{iB}{P} + \frac{rA}{P}$$

$$(3) \quad v = \frac{PC}{B}, \quad B \geq 0, \quad A \geq 0, \quad L \geq 0.$$

Here v is a version of “velocity”, and $f(v)/(1 + f(v))$ can be interpreted as the fraction of consumption spending that is absorbed by transactions or liquidity costs. (Henceforth we will treat “transactions costs”, “liquidity costs” as equivalent and “liquidity benefits” as minus the same thing.) B is government liabilities that provide a transactions or liquidity service, while A is government liabilities valued only for their return. Both may pay interest, but unless demand for liquidity or transaction services is saturated, $i < r$. τ is the proportional tax rate on labor, and ϕ is a lump-sum tax.

An equivalent specification would define

$$C^* = C \cdot (1 + f(v))$$

and then make utility depend on C^* and $b = B/P$, while omitting the $f(v)$ term in the budget constraint. Note that if utility is of the form $U(C, L)$, and we rewrite the problem with C^* but no v term in the budget constraint, and $V(C^*, b, L) = U(C, L)$ in the objective function, we have, by making B/P enter the definition of C^* only via $f(v)$, enforced homotheticity of V in $C^*, B/P$. This

is important for connecting this paper's results to Chari and Kehoe (1999), who emphasize that their results depend on homotheticity of utility in C^* and b .⁵ The money-in-budget-constraint formulation makes it easier to see what fraction of spending is absorbed by transactions costs, and in particular avoids allowing implied transactions costs becoming negative.

The linear disutility of labor, linear production function, and unbounded supply of labor are chosen to make the algebra of deriving our results simpler. As should become clear, the results don't depend on these details of the specification.

To derive the Euler first order conditions for the private agent's problem assuming none of the inequality constraints bind, we form a Lagrangian

$$(4) \quad \log C_t - L - \lambda \left(C \cdot (1 + f(v)) + \frac{\dot{B} + \dot{A}}{P} - (1 - \tau)L + \phi - \frac{iB}{P} - \frac{rA}{P} \right)$$

and apply calculus of variations to arrive at

$$(5) \quad \partial C : \quad \frac{1}{C} = \lambda(1 + f + f'v)$$

$$(6) \quad \partial L : \quad -1 = -\lambda(1 - \tau)$$

$$(7) \quad \partial B : \quad -\frac{\dot{\lambda}}{P} + \frac{\beta\lambda}{P} + \frac{\lambda \dot{P}}{P^2} = \frac{i\lambda}{P} + \frac{\lambda\gamma v^2}{P}$$

$$(8) \quad \partial A : \quad -\frac{\dot{\lambda}}{P} + \frac{\beta\lambda}{P} + \frac{\lambda \dot{P}}{P^2} = \frac{r\lambda}{P}.$$

These can then be simplified to

$$(9) \quad C \cdot (1 + f + f'v) = 1 - \tau$$

$$(10) \quad \frac{-\dot{\tau}}{1 - \tau} = i + f'v^2 - \beta - \frac{\dot{P}}{P}.$$

$$(11) \quad r - i = f'v^2.$$

We assume $f'v^2$ is monotone increasing in v . This is automatically true if $f''(v) \geq 0$ for all $v > 0$, but if we want to allow for the possibility of equilibria with $b = 0$, $f(v)$ must remain bounded as $v \rightarrow \infty$. This is possible with $f'v^2$ monotone increasing, but there are reasonable-looking choices of f that violate monotonicity of $f'v^2$. We rule them out because they raise the possibility of multiple equilibria, which would complicate analysis of the model.

⁵For example, with $f(v) = \gamma v$ and this model's $U(C, L) = \log(C) - L$, we get

$$V(C^*, b, L) = \log \left(\left(\sqrt{1 + \frac{4\gamma C^*}{b}} - 1 \right) \cdot \frac{b}{2\gamma} \right) - L.$$

The government has to finance a given path G of expenditures, which provide no utility to the private agents. Its budget constraint is

$$(12) \quad \frac{\dot{B} + \dot{A}}{P} + \tau L + \phi = G + \frac{iB + rA}{P}.$$

The social resource constraint

$$(13) \quad C \cdot (1 + f) + G = L$$

is derivable from the government and private budget constraints.

II. What are reasonable assumptions about initial conditions?

The government budget constraint (12) implies that total government liabilities, $A + B$ has a continuous time path. In other words, to change $A + B$ requires changing tax revenues $\tau L + \phi$, government expenditures G , or interest payments on debt $iB + rA$. It is unreasonable to think of any of these as being changeable at infinite rates. Treating the rate of growth $\dot{A} + \dot{B}$ as well defined at every date, which is implied by treating (12) as a non-forward-looking equation, makes economic sense.

The government budget constraint does not require that A and B individually have continuous time paths. With B_{-0} and B_{+0} (for example) as notation for the left and right limits of B at time zero, we have

$$(14) \quad B_{-0} + A_{-0} = B_{+0} + A_{+0},$$

while $B_{-0} \neq B_{+0}$ and $A_{-0} \neq A_{+0}$ are possible. Private agents see themselves as having to save out of current income to change their total assets $A + B$, but see themselves as able to trade instantly with the government to any desired ratio of A to B in their asset holdings. But then if $B_{-0} + A_{-0}$ is non-zero, to arrive at zero net worth at time zero requires a policy that makes P_{+0} infinite. If A is bounded below, this implies in turn real balances $b_{+0} = B_{+0}/P_{+0}$ are zero and that transactions costs are maximal.

While setting interest bearing debt to zero (or its minimum feasible value) has the beneficial effect of reducing required future taxes, in the model liquidity-providing balances are important. Unless barter equilibrium ($b = 0$) is optimal, an optimizing government, free to choose a time path of policy that fixes the price level at time zero, will choose a combination of inflation and tax policies that make $B_{+0} = B_{-0} + A_{-0}$ and thus $A_{+0} = 0$.

Another way to explain our initial date assumption is that we assume the government issuing nominal debt is recognized to have the option at any date of diluting the claims of holders of the debt and of money balances by deficit finance, but is assumed to be committed not to introduce a new currency, and

debt denominated in it, that has a higher seniority as a claim on future revenues.

Of course if the government at time zero can repudiate both existing debt and existing currency, then issue new currency and use it to buy private sector liabilities, and if it can do all this while preserving its perfect credibility, it can achieve a better outcome than is obtainable by just setting the P_{+0} so that agents choose $A_{+0} = 0$. However defaulting on outstanding debt and currency is an action that the government is assumed never to do in the future. On the other hand using unanticipated price level fluctuations to stabilize the real value $b + a$, and maintaining $a \doteq 0$, is not only optimal at time zero, in a stochastic version of the model the fully committed government would probably do this at every date.

In mapping the conclusions of this model into an actual economy, the optimality of $A \equiv 0$ should not be taken as implying that interest-bearing debt does not or should not exist. In reality government debt comes in many maturities, all or most of which pay real returns lower than what is apparently available on other investments. The $A = 0$ conclusion means that government debt that exists should pay an interest rate less than or equal to what is available on other investments, with equality only when demand for liquidity has been saturated.

If we removed the constraint that $A \geq 0$, we could implement Chari and Kehoe's conclusion on optimal government actions at the initial date: repudiate (or inflate away entirely) the existing $A_{-0} + B_{-0}$, then issue new liquidity-providing debt by purchasing private sector liabilities, so that $A_{+0} + B_{+0} = 0$. Then a policy that equalizes the return on debt with that on other assets can be financed without any use of taxation. We show in appendix section G that, in agreement with the Chari Kehoe result, under these assumptions it is indeed optimal for the government to equalize the return on private and government liabilities. Note that the initial step here, the repudiation of outstanding debt at time 0, is one that an optimizing government will be tempted to repeat, but must promise never to repeat.

In their classic article on monetary and fiscal policy Lucas and Stokey (1983) consider a model without capital similar to the model in this paper. They assume the government at time 0 has real debt that cannot be repudiated and provides no liquidity service, and that it injects non-interest-bearing money into the system by open market operations in this real debt. They observe that, as in this paper, there is no way to avoid a time-inconsistency problem in their monetary economy when the government cannot commit to a time path for the price level. They discuss optimal monetary policy assuming the government *can* commit to a time path for the price level. Their assumption of non-zero real debt that carries no liquidity premium distinguishes their setup from that in this paper in Calvo (1978), and in Angeletos, Collard and Dellas (2023).

III. Optimal, time-inconsistent policy with full commitment

Here we assume a benevolent government that maximizes the same objective function as the representative agent. The government can credibly commit to any time path for τ , and we assume lump sum taxes are unavailable, i.e. $\phi \equiv 0$.

We assume the liquidity services function $f(v) = \gamma v$. We also assume there is no non-liquidity-providing debt A , only B .

Equation (10) with our linear form for f becomes

$$(15) \quad i = \beta + \frac{\dot{P}}{P} - \gamma v^2 - \frac{\dot{\tau}}{1 - \tau}.$$

We can rewrite the system with the real return on government debt $\rho = i - \dot{P}/P$ replacing i , and, because i and \dot{P}/P enter the system only as their difference, we then have a system of 5 equations in 5 unknowns, which we assemble here for convenience:

$$(16) \quad C \cdot (1 + 2\gamma v) = 1 - \tau$$

$$(17) \quad \dot{b} = \rho b + G - \tau L.$$

$$(18) \quad C(1 + \gamma v) + G = L$$

$$(19) \quad \rho = \beta - \gamma v^2 - \frac{\dot{\tau}}{1 - \tau}$$

$$(20) \quad v = \frac{C}{b}.$$

Note that by using ρ , we have eliminated both i and \dot{P}/P from the system. This implies that given the time path of τ , i affects the equilibrium only by changing the inflation rate. Changing i has no effect on the other five variables: C , v , ρ , L , and b .

This equation system can be solved to allow expressing all variables in the system as functions of v , τ , and $\dot{\tau}$. Dividing the government budget constraint (17) by b and expressing the left-hand side (\dot{b}/b) and ρ as functions of τ , $\dot{\tau}$ and v produces

$$(21) \quad \frac{\dot{b}}{b} = \frac{-\dot{\tau}}{1 - \tau} - \frac{\dot{v}}{v} \cdot \frac{1 + 4\gamma v}{1 + 2\gamma v} = \beta - \gamma v^2 - \frac{\dot{\tau}}{1 - \tau} + \frac{G - \tau L}{b}.$$

The terms in $\dot{\tau}$ cancel, allowing us to derive an expression for \dot{v}/v in terms of τ and v alone:

$$(22) \quad \frac{\dot{v}}{v} = \frac{\overbrace{1 + 2\gamma v}^S}{1 + 4\gamma v} \overbrace{(\gamma v^2(1 + \tau - 2G) + (\tau - G)v - \beta)}^R}.$$

The two expressions in the formula for \dot{v}/v are labeled R and S to aid understanding of the appendix derivation. Note that the locus in v, τ space of constant- v points is defined by $R = 0$.

Equation (22) captures all the constraints on the government imposed by the social resource constraint and private agent choice behavior. It is derived taking the time path of τ as given and makes no use of conditions derived from optimizing government behavior. An optimizing government with the same objective function as the private agent will therefore solve this problem:

$$(23) \quad \max_{\tau, v} \int_0^{\infty} e^{-\beta t} (\log C_t - L_t) dt$$

$$(24) \quad \text{subject to (22), (16), and (18).}$$

Note that (16) and (18) are needed here only to define C and L in terms of v and τ . We can therefore form a Lagrangian for the government's optimization problem, using a single constraint ((22)) and two variables, v and τ . The first-order conditions for this problem lead, after using (16) and (18) to leave just v and τ in the system, to two equations determining the optimal path for τ and v . One involves no derivatives (because $\dot{\tau}$ does not appear in the constraints or the objective function):

$$(25) \quad \frac{1}{1 - \tau} = \frac{1 + \gamma v}{1 + 2\gamma v} + (\gamma v^2 + v)\eta \frac{1 + 2\gamma v}{1 + 4\gamma v},$$

where η is the multiplier, or costate, associated with (22).

The other is a messy expression determining $\dot{\tau}$ as a function of τ and v . Computer code that builds up the expression from components is available in the R function `odevtau()`⁶ and in this paper's replication files. Algebraic expressions that track the computer code's logic are presented in appendix section D. We end up with a system of two ordinary differential equations in two unknowns, v and τ . These are the local (in time) first order conditions for the problem, or Euler equations.

Figure 1 plots information about the trajectories (time paths in (v, τ) space) followed by solutions to the Euler equations. It was generated for a numerical solution of the model with parameters set at $\beta = .02$, $\gamma = .01$, and $G = .25$. However, we show in the appendix that that the qualitative nature of the solution and this plot holds for all feasible parameter values.

The figure has the labor tax rate τ , bounded above by 1, on the vertical axis and velocity $v = C/b$ on the horizontal axis. The green line on the figure is the locus of points satisfying $\dot{v} = 0$. The line can be calculated from (22). The government can implement a steady state with v and τ constant at any point on the green line simply by announcing that τ will be constant. The blue triangle is the point on the $\dot{v} = 0$ line corresponding to the highest period utility $\log C - L$ among steady state equilibria on the green line.

To bring in information about the government's optimal policy, we add to Figure

⁶Available at <http://sims.princeton.edu/yftp/CostlessDebt/odevtau.R>.

1 a 100 by 100 point grid of red arrows representing $(\dot{v}, \dot{\tau})$ vectors implied by the Euler equations of the government's optimization problem. The lengths of the arrows indicate speed of movement in the directions the arrows point. The locus of (approximate) $\dot{\tau} = 0$ points on the grid is then the set of points where the Euler equations make $\dot{\tau} = 0$.⁷ It is the black line on the figure. Any intersection of the black and green lines is a steady state of the planner's problem.

From the figure, it might seem that the black and green curves do not intersect at all, i.e. that there is no steady state satisfying the Euler equations. This is an artifact of the finite resolution of the grid. Figure 4 zooms in on a thin sliver at the left edge of Figure 1 plot, using a 100 by 100 grid on this narrow slice of (v, τ) space, and it is clear that the two curves do intersect — at $\tau = 1$. (This is shown analytically in appendix section E.)

The planner can choose the starting point of the economy's path. By announcing the time paths of τ and deficits, the planner determines the initial price level and, thereby, both v and τ at the initial date. With the initial state freely chosen, the optimum is determined by the initial transversality condition $\eta_0 = 0$, combined with restrictions on the path's behavior as $t \rightarrow \infty$. η_0 is the Lagrange multiplier on the planner's constraint and is zero at t_0 because τ and v are unconstrained at t_0 .

The red arrows on Figure 1 make it clear that when starting from points above the $\dot{v} = 0$ line, these Euler equation paths imply eventually rapid increase in v , with τ converging to the $\dot{\tau} = 0$ line. In fact it can be shown, and is shown in Appendix E, that along such paths $v \rightarrow \infty$ in finite time. This means $b \rightarrow 0$ and $C \rightarrow 0$. L converges to G along the path, but if the $v = \infty$ limit is actually reached, the incentive to work disappears, as all of consumption expenditure is absorbed by transactions costs. If we interpret the exogenously given G as a requirement for the government over the whole infinite future, these paths are infeasible — once government liabilities are entirely inflated away, the economy cannot produce G .⁸

Paths that start well below the $\dot{v} = 0$ line lead to steadily declining v , toward $v = 0$. These paths are ruled out by the private agents' transversality condition. Along these paths τ steadily increases toward 1, and, as $v \rightarrow 0$, C and b also shrink toward zero, though since $v = C/b$ shrinks toward zero, C is decreasing faster. Along these paths, then, b becomes an ever-larger multiple of C . An agent in an economy on this path can consider at time t^* ceasing to work (setting $L_{t^*} = 0$) and from that point onward simply consuming her finite stock b_{t^*} over time. The lifetime utility achievable this way, assuming the real rate of return on debt $\rho \geq 0$,⁹ is a constant plus $\log(b^*)$ plus the discounted present value at t^* of

⁷It was plotted using the `R contour()` function.

⁸There is a separate question as to whether simply announcing a time path for current and future τ is enough to guarantee a unique initial price level. We take this up in appendix section C. Ruling out explosive paths for the price level usually requires that policy commit to higher taxes if the initial price level is above that consistent with the saddle path.

⁹We show in appendix section E that on these paths the real return ρ on government debt converges

$\log(1 - \gamma v_t)$. But, since C is decreasing, the utility on the original path is less than $\log(C_{t^*})/\beta$, which, as $v \rightarrow 0$, eventually becomes smaller than the utility obtainable by simply consuming b^* over time, with $C_t = b^* \beta e^{-\beta t}$.

There is only one path (the “saddle path”) that satisfies the Euler equations while neither diverging toward $v = 0$ nor diverging to $v = \infty$. The start of such a path is shown as the orange line on Figure 1. It was calculated by imposing $\eta_0 = 0$ and varying v_0 ¹⁰ to find a line that did not diverge either to the left or the right.¹¹ If we could calculate the full path, it would converge to $\tau = 1$, but beyond the part of it shown in orange, it would on the Figure 1 graph seem to coincide with the green line. They are so close, because in this range τ on the optimal path is increasing *extremely* slowly. The part of the optimal saddle path shown as the orange line is traversed over approximately 400 years. Though the economy proceeds, if the time-0 policy commitments are honored, toward $\tau = 1$, it does not get close to $\tau = 1$ for many centuries.

The time paths of several of the model’s variables are shown in Figure 2, for the first 400 years, and in Figure 3 for the first 10 years. Inflation is high at the start, taxes are very low, and utility is higher than in the steady-state optimum over the entire first 10 years.

In Calvo’s model with real balances entering separably in the utility function, it is optimal to set taxes to zero at time 0. In this model, with $\beta = .02$, $\gamma = .01$, it turns out that optimal $\tau_0 < 0$, i.e. a labor subsidy. The optimal path delivers considerably higher utility in the initial, low-tax periods, and utility approaching $-\infty$ in the far distant future. τ rises above the value required to support this level of G in the optimal steady state (the blue triangle, where $\tau = .27$), but not for a long time, and the rise of τ toward 1 is very slow. The near-term benefits of low taxes more than offset the heavy, but long-delayed, costs of high taxes in the distant future.

To visualize the last part of the optimal path, Figure 4 zooms in on the sliver of Figure 1 at the left edge, for v in $(0, .1)$. There is no separate orange line, because it so nearly coincides with the green $\dot{v} = 0$ line. Progress toward $\tau = 1$ in this region is extremely slow, as evidenced by the arrows appearing almost perfectly horizontal. The very small $\dot{\tau}$ values in this region are, however, all positive below and to the left of the black $\dots \tau = 0$ locus. The benefit of violating the time-zero commitment in this region is very large, as measured by η , which approaches infinity as the saddle path approaches 1. It is likely that the temptation for a policy maker to re-initialize by dropping taxes would be overwhelming. Note that this is different from the simple Phillips curve examples of time-inconsistency, where the benefits of abandoning commitment to no surprise inflation are constant

to zero.

¹⁰This method is called “multiple shooting”.

¹¹Actually, *every* path calculated diverged either to the left or the right. What is shown is just the longest path the differential equation solver could calculate before it diverged. Changes in the initial conditions in the fourth significant digit made the solution change from the path shown (which actually crosses the green line and shoots off to the right) to a path that parallels the green line for a while then shoots off to the left.




FIGURE 1. BEGINNING OF OPTIMAL PATH WITH $\beta = .02$, $\gamma = .01$, AND $G = .25$

Note: While the orange saddle path appears to hit the green constant- v locus, in fact it stays extremely close to it, but always below it, and the orange, green and black lines all meet at $\tau = 1$, never crossing before that point. The red arrows show the derivatives of the v, τ vectors. The derivative vectors change rapidly in the neighborhood of the green $\dot{v} = 0$ line. In fact they parallel the green line at points just below it. The blue triangle is the optimal steady state, corresponding to the $\gamma = .01$ line of table 3.

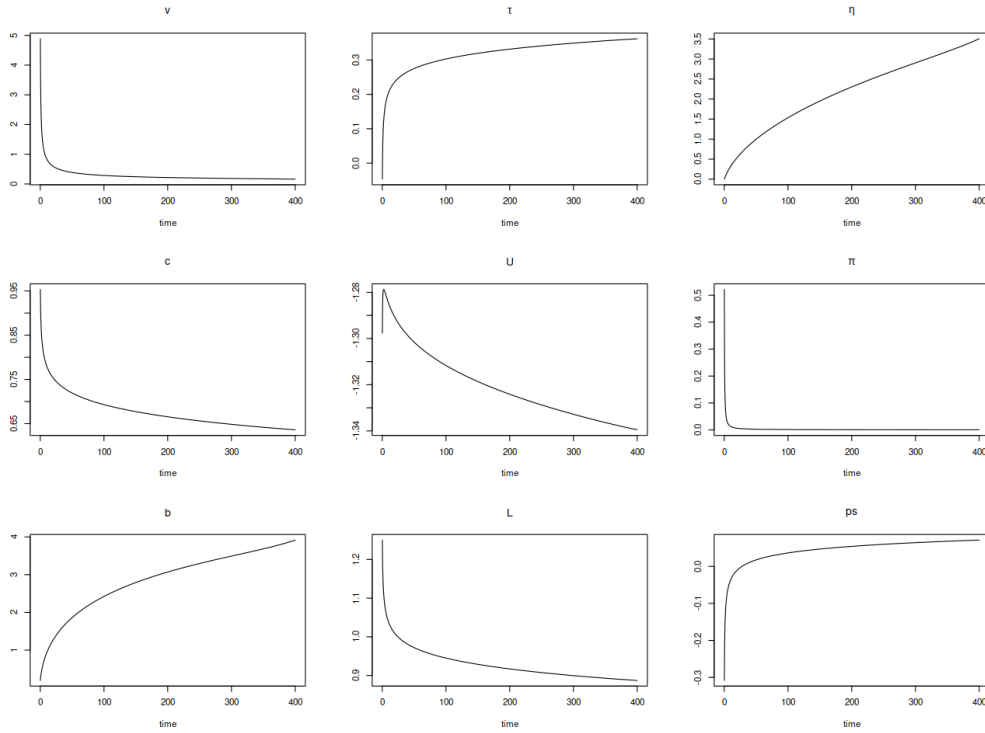


FIGURE 2. 400 YEARS ALONG THE SADDLE PATH

Note: v : velocity (C/b); τ : tax rate; η : costate; C : consumption; U : utility; π : inflation; b : real debt; L : labor input; ps : primary surplus. In the optimal steady state, $v=.85$, $C=.72$, $U=-1.30$, $\pi=.0072$, $b=.81$.

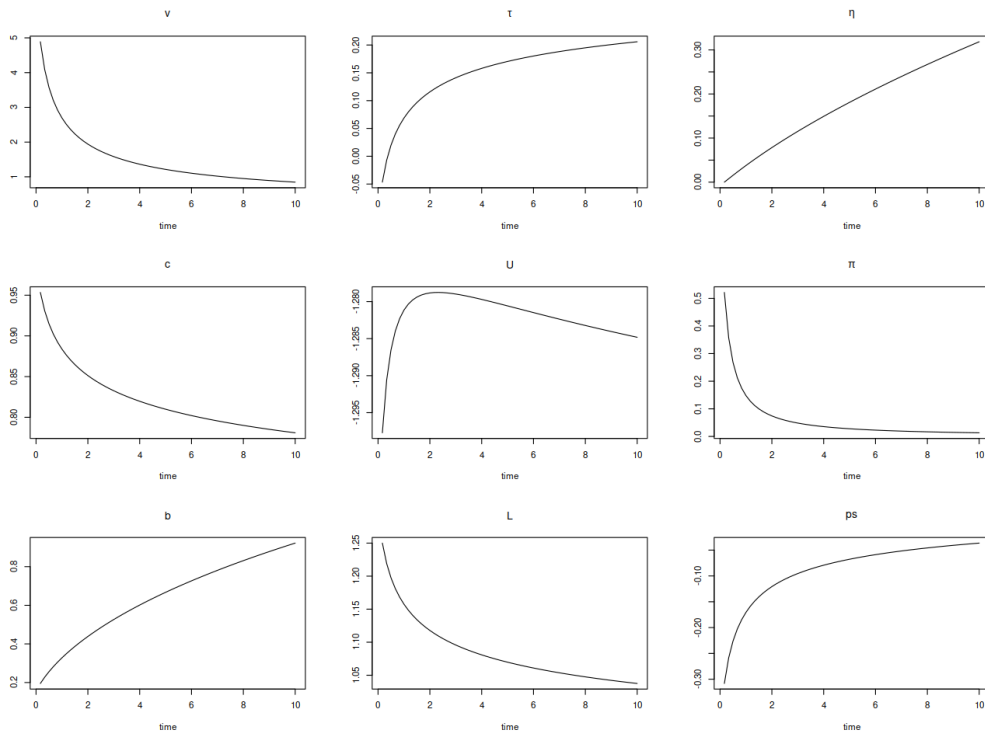


FIGURE 3. 10 YEARS ALONG THE SADDLE PATH

Note: See note to figure 2




FIGURE 4. END OF OPTIMAL PATH

Note: As in Figure 1, the green line is the $\dot{v} = 0$ locus and the black line the $\dot{\tau} = 0$ locus. The optimal path almost coincides with the green line, lying barely below it and visually indistinguishable from it in this plot.

over time after the initial period.

IV. Other values of G and γ

The qualitative nature of the phase diagram shown in Figures 1 and 4 does not change as G and γ are changed, but the magnitudes change a lot. If $G = 0$ while γ stays at .01, the initial positive rate of inflation is just 0.40% — that is, 40 basis points — and the initial negative tax rate is -0.44%. Thereafter, as in any solution, the tax rate steadily rises toward 1.0, but 250 years after the initial date has reached only 12%. The improvement in utility over the steady state value is very small — in the third significant digit — but persistent, lasting over 50 years. The initial inflation also dissipates slowly, taking around 75 years to drop from 0.40% to 0.10%,

With $G = .8$, numerical results are difficult to obtain, because the differential equation system, which is already “stiff”, becomes nearly unmanageably so. The initial inflation rate in this case is extremely high and drops extremely fast. I can’t obtain stable solutions for the time path of the system starting from the $\eta_0 = 0$ initial condition. I can, however, obtain a solution starting from a point farther along the saddle path, $v = 16.727$, $\tau = 6$. Here $\eta = .10$ rather than 0, but at this point on the saddle path the instantaneous inflation rate is 3.98 — in log units, not per cent, i.e. thousands of percent per year. After one year the rate is back to .53, and after another year it is back to .27, so it is declining rapidly, and it is below 1% after 54 years. The gains in welfare relative to the constant-tax solution with the same G and γ is large, .26 log units, and utility remains above the constant-tax solution level for the first 50 years. The tax rate rises from the initial .60 to .75 in the first 3 years, and goes above the constant-tax optimum level of .805 after 30 years.

From these results it is clear that the optimal initial inflation rate is quite sensitive to G , growing larger as the eventually necessary level of taxation increases. It is also somewhat sensitive to γ , increasing as γ increases.

V. Tax revenue vs. tax rate

Most of the existing literature in this area treats the fiscal policy variable as tax revenues, with tax distortions a function of revenue. In such a framework, it is puzzling to find optimal policy leading to tax rates that are on the far side of the “Laffer curve” — i.e. that deliver less revenue than would be available with a lower tax rate. But in this paper’s model the tax rate τ and its growth rate enter the determination of the real return on government debt, as we see in (10). Furthermore, the effect depends on $\dot{\tau}/(1 - \tau)$. so even when τ gets close to one and $\dot{\tau}$ itself is close to zero, $\dot{\tau}/(1 - \tau)$ can have large proportional effects on the $(\beta - \rho)b$ seigniorage term in the budget constraint. Of course as τ approaches one on the saddle path it would be much better to lower taxes, thereby increase revenues, and make the representative agent much better off. But doing so would

deviate from the original commitment generating the optimal path equilibrium. This is just one symptom of the strong time inconsistency in the optimal saddle path solution.

VI. What if b_0 , rather than B_0 , is taken as given?

In their closely related paper, Angeletos, Collard and Dellas (2023) treat real debt b_0 as a state variable. The model in that paper differs from this one in some other dimensions as well, and it is beyond the scope of this paper to trace out all the sources of differences between the two papers. Nonetheless it may be useful to consider what happens to this paper's results if we treat initial real debt b_0 as given.

This amounts to assuming that a price jump at time 0 is impossible. This constraint replaces the $\eta_0 = 0$ transversality condition, but leaves the Euler equation unchanged. From (D.10) we can see that $\eta_0 = 0$ implies a positively sloped relation between τ and v . This positively sloped line can have at most one intersection with the negatively sloped saddle path, which is the initial (v_0, τ_0) point on the optimal path. From the private sector FOC's and the definition of v , we have

$$(26) \quad b_0 = \frac{1 - \tau}{v \cdot (1 + 2\gamma v)}$$

This implies a negatively sloped concave relation for τ as a function of v . On our Figure 1, this would be a line starting in the upper left corner and decreasing with v . Since it is concave and the saddle path is convex, it will intersect the saddle path either not at all, at two points, or at a single tangency point. There are then usually two solutions, if any, to the Euler equations with this kind of initial condition — one at a high, and another at a low level of τ . From that point onward, the solution will be on the same saddle path as the $\eta_0 = 0$ solution. The government could choose between initial conditions, when there are two, based on which gave a higher value for the discounted utility objective function.

The differences in initial points are the only differences between the B_0 given and b_0 given solutions, though. The same pattern of ever-increasing τ converging toward 1, with velocity starting higher and heading downward, applies with either type of initial condition. The two solutions can even be exactly the same, if the given b_0 happens to match that chosen by the optimizing government with $\eta_0 = 0$.

The nature of time-inconsistency is different with b_0 given. So long as the chosen intersection of the saddle path with the concave (v, τ) locus defined by current b_t continues to correspond to higher discounted utility, there is no time-inconsistency. But as the economy progresses along the saddle path it may eventually arrive at a point where the objective function rankings of the two choices reverse. It would then be tempting to jump (without changing b , of course) to the lower-tax initial condition by deviating from the initially announced τ path.

This temptation would persist so long as the reversal of rankings persisted. This contrasts with the B_0 -given case, where at every moment it is tempting to jump back to the $\eta_0 = 0$ initial point.

VII. Limited commitment

The infinite horizon outcome with full commitment is unlikely to be even approximately implemented in practice because of the extremely strong and growing incentive to deviate from it as time goes on. On the other hand assuming no ability at all, even for an instant, to commit to a policy seems equally unrealistic. Realistic models with limited ability to commit would be much more complicated than this paper's setup, but we can gain some insight from this paper's model if we assume an ability to commit for a fixed, known span of time.

Suppose the policy maker remains in office for T years, after which a new policy maker, unbound by the previous policy regime's commitments, takes over, again for a new span of T years, with this rotation of policy makers going on forever. In most models of this type calculating optimal behavior depends on determining how the state variables at the end of a regime affect behavior of the policy maker in the following regime, but in our setup, because the new regime can start with a jump in the price level, there is no effect of the terminal state at time T on the choices of the new policy maker. Each policy maker has an initial transversality condition that makes $\eta = 0$.

However, there can be no jump in the price level with the regime change. The regime changes are deterministic and known by the public. A jump in the price level after the initial date would imply an infinite positive or negative rate of return on government debt at the regime-switch date. Private agents would attempt to buy or sell debt before the price-jump date and undermine the equilibrium. Note that this constancy of the price level across the regime switch date is not a constraint on the new policy-maker's behavior. The new policy maker is free to cause a jump in the price level. The constraint is on the previous regime: the price level at the end of the previous regime must be exactly the value that the new regime's policy will choose, so there is no jump. Since the level of nominal debt B is also assumed not to be able to jump, we can characterize this constraint equivalently as requiring that real debt b not jump. Then, since all policy makers have the same objective function and constraints, the equilibrium chosen by the policy maker in each regime is the same, and b in equilibrium takes the same value at all policy regime-switch dates.

We can calculate the equilibrium in this limited-commitment model using the same set of differential equations we used in the infinite-horizon, full commitment model, replacing the non-explosiveness terminal condition with a constraint on the terminal value b_T . Among paths satisfying the two differential equations of our model over $(0, T)$ satisfying $\eta_0 = 0$ and $b_T = b^*$, we choose the one with $b_0 = b^*$,

reflecting our knowledge that all policy makers are identical.¹²

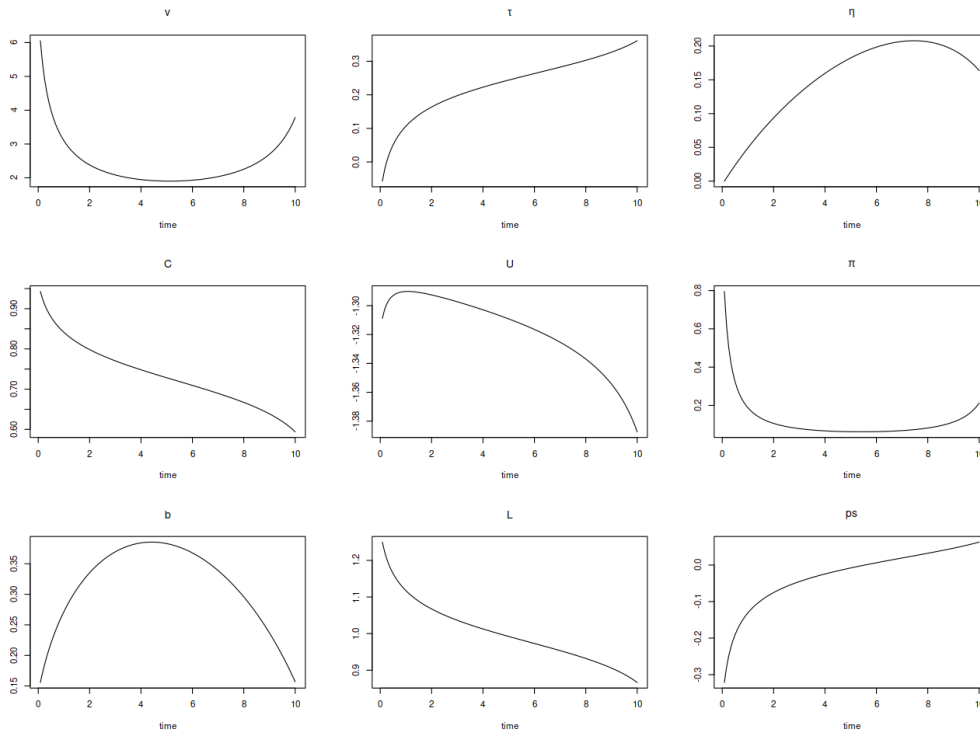


FIGURE 5. PATH WITH 10-YEAR COMMITMENT

See notes to Figure 2

Figure 5 shows time paths of model variables when $T = 10$. Compared with the first 10 years of the full commitment solution, the initial inflation is even higher. The tax rate rises much more rapidly. Though utility $U(C, L)$ briefly goes above the steady-state optimal level, the average over the 10 years is much lower than the optimal steady state value, whereas the full-commitment solution delivers higher utility than the steady-state solution over the first 10 years.

That the outcome is worse when commitment is possible only over a finite span is not surprising. Perhaps not so obvious is that recognizing the finite commitment span does not damp the incentive to run initial large primary deficits and high inflation.

If the commit period T is long, the solution for the first few years looks similar

¹²It is important that the policy optimization problem treats b_0 as freely chosen and $b_T = b^*$ as a constraint, with no connection to b_0 . That $b_T = b_0$ is an equilibrium condition, not perceived as a constraint by the policy maker.

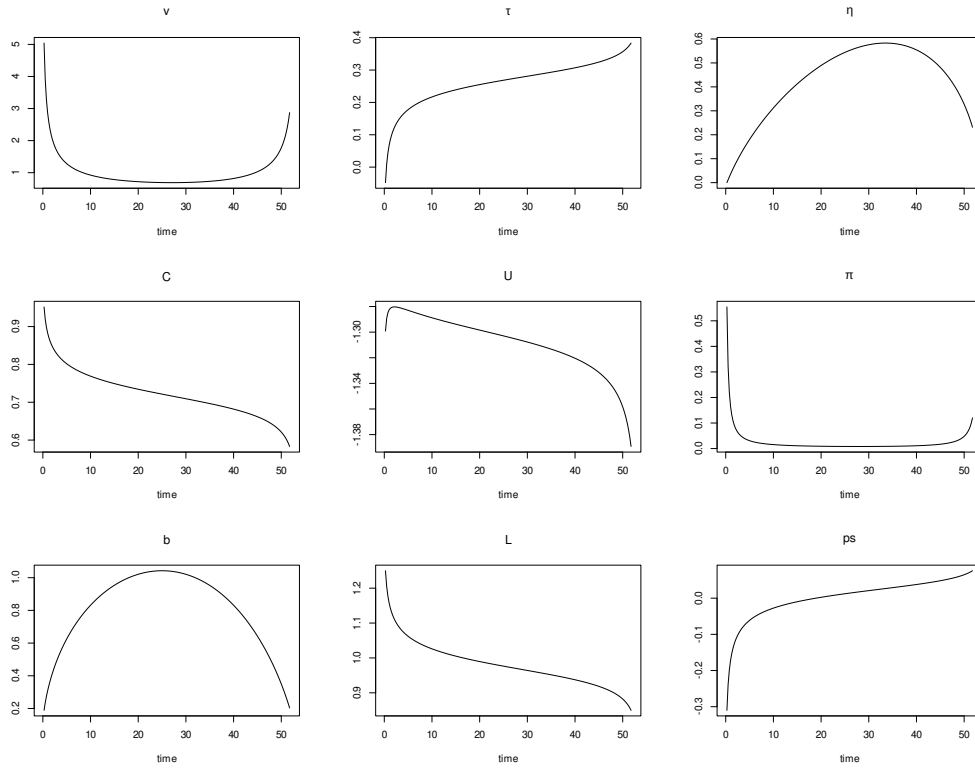


FIGURE 6. PATH WITH 52-YEAR COMMITMENT

See notes to Figure 2

to that for the infinite-horizon commitment case. Figure 6 shows solution paths when $T = 52$. The first third or so of the plots are close to the corresponding parts of the paths for the infinite horizon case shown in Figure 2. But in the middle of the 52 year commitment span debt starts shrinking, taxes start rising faster, and utility starts an accelerating decline.

uss	uinf	u10	u52
-1.3043	-1.2955	-1.3145	-1.2995

TABLE 1—UTILITY LEVELS BY POLICY SETTING

uss: commit to a constant τ ; uinf: full infinite horizon commitment; u10: 10 year finite time span commitment; u52: 52 year finite time span commitment. Calculated by numerical integration of the time paths of utility, weighted by $e^{-.02t}/.02$. Parameter values are $G = .25$, $\beta = .02$, $\gamma = .01$.

Weighted mean utility values (the objective function multiplied by $\beta = .02$) for policies that commit to a constant τ , commit to optimal path for τ over an infinite horizon, or commit over a fixed time span of 10 or 52 years are shown in Table 1. The full and flexible commitment produces the best result, as expected; the 10-year commitment is worse even than a constant- τ commitment; and the 52-year commitment, though worse than the infinite horizon case, is better than a fixed- τ commitment. The differences in average utility across these cases are modest — they can be thought of as in log consumption units. The wider swings in utility in the time paths we have displayed above shift the timing of C , L and thereby U substantially, but the high initial and lower later values tend to cancel out.

VIII. Optimal steady states

Why study the constant-tax case when we know that, for a government that can commit, it is not optimal? One reason is that the unconstrained solution involves trading high utility now for much lower utility in the distant future. If all agents in the economy are alive now and live forever, this is not a problem. But it is also possible to interpret the representative agent economy as consisting of finite-lived agents who are born at a constant rate and discount the future only because of a constant risk of death. A benevolent government might then not discount the future, in order to treat current and unborn future agents equally, which would lead to choosing the optimal constant tax steady state.

The fully optimal, non-constant-tax solution generates severe time-inconsistency — much more severe than that in the constant-tax-rate solutions. In fact, it generates, after a very long time, utility that tends to minus infinity. It seems this policy could actually never be recommended, because it requires maintaining commitment to the optimal path despite the benefits of abandoning commitment

growing without bound. So commitment to a constant tax policy might be possible, even though sustaining a commitment to the optimal non-constant tax path is not.

In this section we solve numerically for optimal stationary equilibrium with a given constant government spending G and lump-sum tax ϕ . If ϕ were freely variable, it would be optimal to set $\tau \doteq 0$ and nearly saturate the demand for liquidity. We include ϕ in the model because it may clarify thinking about conditions for existence and uniqueness of equilibrium, which we take up in the appendix. But for numerical solutions we set $\phi = 0$ for convenience.

The optimizing government sets constant values for the labor tax rate τ and i , the nominal interest rate on B . It maximizes the equilibrium value of the private agent's objective function (1), which in this case, restricted to steady states, is the same as maximizing the undiscounted $U(C, L) = \log(C) - L$. The government takes (16)-(20) as constraints. However, because we are now assuming $\dot{\tau} \equiv 0$, the term in τ in (19) drops out. An argument in appendix section B shows that transversality and feasibility guarantee that when τ is constant, b must also be constant. Equilibrium for given constant τ and G is then defined by 5 ordinary equations in 5 unknowns.

Because b can (via a jump in initial price level P_0) jump at the initial date, there is no transition path from equilibrium with one level of constant τ to another if a one-time, permanent, unanticipated change in τ occurs. In this flex-price model, the economy goes immediately to the new equilibrium. The same is true if there is a one-time, unanticipated shift in a parameter, like the transactions cost or liquidity service parameter γ , so long as τ remains constant or simultaneously moves to a new constant value.¹³

The equation system for the constant- τ case can be solved analytically, resulting in a univariate quadratic equation. Existence of constant- τ equilibrium requires $\tau > 2G - 1$. In that case, there is a unique steady state for every feasible τ , and there are no non-steady-state equilibrium paths¹⁴, so the initial price level is uniquely determined. Note that $\tau < 1$, because otherwise it is optimal to set $L = 0$, and that therefore constant- τ steady states require $G < 1$.

In the tables below τ is the optimal constant τ for the table's assumed value of G . It was found by a one-dimensional grid search on τ . In all the tables $\beta = i = .02$, which means that the gap between β and ρ is generated entirely by inflation.

Tables 2 3, and 4 show the model's optimal steady state at various levels of government spending G and transactions cost parameter γ . Even in Table 2, where there is no government spending to finance, it is optimal to have a positive rate of labor tax τ , in order to induce low inflation, a positive real return ρ on

¹³This contrasts with models (like Angeletos, Collard and Dellas (2023)) that assume the real value of government debt at the initial date b_0 is given. In that case, there would be transition dynamics and the optimal constant value of τ would depend on the initial b_0 value.

¹⁴As is shown in appendix section A, uniqueness may depend on trigger policies, or small positive values for ϕ that are negligible in equilibria but important on off-equilibrium paths.

γ	C	b	L	\dot{P}/P	U	τ	$psurp$	$\frac{\gamma v}{1+\gamma v}$	P_0
0.001	0.972	1.31	0.973	0.0006	-1.001	0.026	0.0254	0.0007	0.76
0.010	0.940	2.70	0.943	0.0012	-1.005	0.054	0.0507	0.0035	0.37
0.100	0.868	5.27	0.882	0.0027	-1.024	0.103	0.0912	0.0162	0.19
1.000	0.721	9.27	0.778	0.0061	-1.104	0.166	0.1293	0.0722	0.11

TABLE 2—OPTIMAL STEADY STATE WITH $G = 0$, $\beta = .02$, $i = .02$

γ : transactions cost parameter; C : consumption; b : debt; L : labor; \dot{P}/P : inflation rate; U : utility; τ : labor tax rate; $psurp$: primary surplus; $\gamma v/(1 + \gamma v)$: proportion of consumption expenditure absorbed by transaction costs; P_0 : initial price level, assuming $B_0 = 1$; G non-productive government expenditure; β : discount rate.

γ	C	b	L	\dot{P}/P	U	τ	$psurp$	$\frac{\gamma v}{1+\gamma v}$	P_0
0.001	0.741	0.28	0.992	0.0069	-1.293	0.256	0.0037	0.0026	3.56
0.010	0.721	0.83	0.977	0.0076	-1.304	0.266	0.0103	0.0086	1.21
0.100	0.668	2.19	0.938	0.0093	-1.342	0.291	0.0234	0.0296	0.46
1.000	0.547	4.68	0.861	0.0137	-1.464	0.325	0.0295	0.1047	0.21

TABLE 3—OPTIMAL STEADY STATE WITH $G = .25$, $\beta = .02$, $i = .02$

See notes to Table 2

government debt, and hence a reduced incentive to conserve on real balances. However even in this case the inflation rate does not go to zero, as would be required to make $\rho = .02$ and thereby implement the Friedman rule.

The column labeled $psurp$ is the steady state primary surplus, $\tau L - G$. Steady states with negative primary surplus exist in this model, even though in its simple form it has no growth. A permanent primary deficit arises when the real rate of return on government debt is negative, and by choosing a low enough constant τ , the return can be driven negative. However as can be seen from the three tables, the primary surplus turns negative here only when G is very large, in Table 4.

The most realistic of these cases, might be Table 3's $\gamma = .01$ row. This produces a "debt to GDP" ratio of about 0.8, transactions costs absorbing about 1% of consumption and a tax rate of 27%. In this case the optimal real rate on debt ρ (the 2% nominal rate i minus the inflation rate) is 1.24%.

In Table 4 taxes and inflation substantially depress consumption in order to accommodate the high level of G . Inflation is over 8%. Nonetheless the steady state primary deficit delivers only a small fraction (about half of one percent when $\gamma = .01$) of total revenue.

With $\tau = 0$, the primary deficit is increasing in v over the whole $(0, \infty)$ range, and approaches .5 from below as $v \rightarrow \infty$. Even in Table 4, where government

γ	C	b	L	\dot{P}/P	U	τ	$psurp$	$\frac{\gamma v}{1+\gamma v}$	P_0
0.001	0.196	0.02	0.998	0.0823	-2.627	0.800	-0.0013	0.0090	46.27
0.010	0.188	0.06	0.994	0.0855	-2.665	0.801	-0.0042	0.0284	15.55
0.100	0.167	0.17	0.983	0.0955	-2.775	0.801	-0.0129	0.0890	5.87
1.000	0.123	0.37	0.965	0.1137	-3.060	0.794	-0.0342	0.2522	2.74

TABLE 4—OPTIMAL STEADY STATE WITH $G = .8$, $\beta = .02$, $i = .02$

See notes to Table 2

expenditure is taking up around 80% of output and the tax rate is about 80%, the primary deficit remains well below its upper bound. Even though the Friedman rule is not optimal in this model, the inflation tax produces little revenue relative to the distortion it induces, so that it is optimal to use it to generate revenue only when the other tax available is at highly distortionary levels.

Note the contrast between these steady state results and those for the fully optimal equilibrium. Optimal steady states involve positive primary surpluses and small gaps between β and ρ , except when a very large proportion of output is taken up by G . Even with $G = .8$, though, the proportion of revenue raised through seigniorage is small. The steady state results suggest that, though the Friedman rule is not optimal, it is not a bad approximation to optimal policy.

In contrast, the non-stationary optimum involves negative primary surpluses and a large gap between β and ρ over a decade or more after the initial date. Over time, it does drive $\beta - \rho$ toward a small, but positive limit, and in this sense approximates the Friedman rule. But it does this by driving C and B/P both toward limiting values of zero, and thus also $U(C_t, L_t)$ toward a limiting value of $-\infty$.

IX. Allowing for growth

The arithmetic of “zero fiscal cost” debt might seem to rest on there being positive growth. The idea is that if the economy is growing at a rate exceeding the real rate of return on government debt, debt can be increased today, without any actual or expected future increases in taxation (or reductions in expenditure), with debt nonetheless shrinking relative to output over time.

This line of reasoning is misleading. It ignores wealth effects and inflation. People must be induced to hold the additional debt. They can be induced to do so, even without increased taxation, but only by lowering its real value through inflation. This is true whether or not there is economic growth.

In fact, in this paper’s model if we introduce constant labor augmenting technical progress at the rate θ and interpret C , G and $psurp$ as ratios to $e^{\theta t}$, the model is unchanged except that the discount rate now has to be interpreted as the discount rate β of private agents plus the rate of technical progress θ . None

of the entries in the steady-state table then change. The “ $\beta = .02$ ” line in the captions would change to $\beta + \theta = .02$.

Of course, tables calculated with private agents’ $\beta = .02$, $\theta = .02$, would differ from those shown. But the difference would not be lower labor taxes at a given G . Because the Friedman rule requires a higher return on B with an increase from β to $\beta + \theta$ in the discount rate, the optimal steady state is likely to have a positive primary surplus (that is, $\tau L > G$) even with positive θ . If we start with the economy of the second row of Table 3 and increase θ from 0 to .02 or .04, the optimal steady state still implies a positive primary surplus, and the primary surplus is larger for larger θ . The increased liquidity costs from a greater gap between ρ and $\beta + \theta$ more than offset the benefits from the reduction in τ that would be possible with a permanent primary deficit.

X. Allowing for $\rho = \beta$ at finite v

This paper, like many related papers, specifies the liquidity services of government debt as positive no matter how high the ratio of real debt to consumption, so an equilibrium with saturated demand for liquidity services ($\rho = \beta$) is not possible. The main results of this paper hold, though, even in versions of the model that do allow saturation. In the online appendix section F we first show that if the generic transactions cost term $f(v)$ in equation (2) is smooth enough, then the optimal steady state has $\rho < \beta$, even when a steady state with saturation is feasible. The reason is that an f that goes very smoothly (with many derivatives) to zero as v shrinks makes the marginal reduction in transactions cost per unit reduction in v very small for small f , while the fiscal cost of additional reductions of v shows no such shrinkage.

Then we consider the fully optimal policy in a version of the paper’s main model in which the $f(v) = \gamma v$ assumption is replaced by $f(v) = \max(0, \gamma(v - v^*))$. In this case there is a path for the economy that satisfies the Euler equations and, as in the main model, has τ increasing very slowly to one, while v decreases to a limit strictly greater than v^* . Though in principle it might seem that there could be a steady state with v just below v^* that, by eliminating transactions costs, might dominate the interior solution, this is not true for any $v^* < \bar{v}$, where \bar{v} is the lower bound and limiting value of v on the interior solution path. Though transactions costs are gone, maintaining a $v < \bar{v}$ in steady state requires such high taxation that for $v^* < \bar{v}$ there is no feasible steady state. A sketch of the details of these results is in appendix section F.

XI. Are there generally useful insights from this simple, stylized, model?

This model makes strong functional form assumptions that affect its conclusions, but it embodies principles that may apply more widely. Why is it optimal to inflate early and postpone taxation? One reason is that the gap between the

return on government debt and the discount rate β means that when revenue collection is postponed, the ratio of required future revenue to the revenue initially postponed grows slower than β . Its discounted present value therefore shrinks the more it is postponed. The solution is not indefinite postponement, because as the required revenue grows, the costs of the distortion induced by taxation grow more than linearly with the revenue. Another reason is that high future seigniorage revenue, since it is anticipated, raises the current price level and thereby reduces the base for the current inflation tax, while high future taxes have the opposite effect on the base for the current inflation tax. These are general principles that are likely to apply in other models.

Calvo's model produces a qualitatively similar result. Though his framework, unlike this paper's, made satiation of demand for money possible in a steady state, he showed that with full commitment the economy has high initial inflation and converges to a steady state that preserves a gap between the return on money and the discount rate.

Since abandoning the full-commitment path eventually becomes extremely beneficial, it is unlikely that any government would actually follow the path for a long time. But as we have seen, a solution with temporary commitment can provide lower discounted utility than a steady-state solution. The optimal steady state solution also requires commitment, but that commitment is arguably more sustainable. The incentive to deviate from the steady state commitment is constant over time, rather than ever-growing, and it might be easier to enforce the simpler behavioral norm of constancy in tax rates.

Nonetheless, the tradeoffs the full-commitment path brings out are likely to affect policy discussion, especially in times of fiscal stress. The argument that a gap between the return on government debt and market interest rates is a fiscal resource that can be used to avoid high current levels of taxation is correct. That this will require increased fiscal effort in the future, but only a modest increase, and one that can be repeatedly postponed at modest cost, is also correct.

The paper's model and our discussion of it leave plenty of open questions. Perhaps among the most important are: How would conclusions change if there were capital, or another asset? What if prices are sticky?

XII. Real world complications

Blanchard (2019), based on his Figure 15, argues that there is some evidence for a decline in the real return on capital, measured as the rate of profit relative to market capitalization. If there is such a decline¹⁵, we could interpret it in this paper's model as a decline in θ , the rate of labor-augmenting technological improvement. This would imply that the gap between the return on debt and private asset returns might have been declining despite the stable rate of infla-

¹⁵The rate was in the 6-8% range in the 1960's in the US, and has been in the same range since 2001. There is some apparent downward trend, but it is not strong or uniform.

tion, and therefore that the economy has been moving toward less reliance on seigniorage finance. This does not in itself, of course, imply that returning to previous levels of seigniorage would be optimal.

This paper uses a representative agent model, and therefore cannot consider intergenerational tax-shifting or “crowding out” issues. Blanchard and Mehrotra/Sergeyev instead assume away tax distortions in order to focus on intergenerational issues. Both aspects of debt finance are important, and should be considered jointly. Furthermore, liquidity premia on government debt vary somewhat across the term structure and across time, and currency, bearing no interest, does exist. A more serious quantitative evaluation of the effects of debt finance should consider all of these potentially important factors as operating jointly.

REFERENCES

- Angeletos, George-Marios, Fabrice Collard, and Harris Dellas.** 2023. “Public Debt as Private Liquidity: Optimal Policy.” *Journal of Political Economy*, 131(11): 3233–3264.
- Blanchard, Olivier.** 2019. “Public Debt and Low Interest Rates.” *American Economic Review*, 109(4): 1197–1229.
- Calvo, Guillermo A.** 1978. “On the Time Consistency of Optimal Policy in a Monetary Economy.” *Econometrica*, 46(6): 1411–1428.
- Chari, V.V., and Patrick J. Kehoe.** 1999. “Chapter 26 Optimal fiscal and monetary policy.” In . Vol. 1 of *Handbook of Macroeconomics*, 1671–1745. Elsevier.
- Chari, V.V., Lawrence Christiano, and Patrick J. Kehoe.** 1996. “Optimality of the Friedman rule in economies with distorting taxes.” *Journal of Monetary Economics*, 37: 203–223.
- Friedman, Milton.** 1969. “The Optimum Quantity of Money.” In *The Optimum Quantity of Money and Other Essays*. , ed. Milton Friedman, 1–50. Aldine Publishing Company.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** 2012. “The Aggregate Demand for Treasury Debt.” *Journal of Political Econom* vol. 120, no. 2, 120(2): 233–267.
- Lucas, Robert E. Jr., and Nancy Stokey.** 1983. “Optimal Fiscal and Monetary Policy in an Economy without Capital.” *Journal of Monetary Economics*, 12(1): 55–93.
- Mehrotra, Neil R., and Dmitriy Sergeyev.** 2021. “Debt Sustainability in a Low Interest Rate World.” *Journal of Monetary Economics*, 124(S): S1–S18.

Mian, Atif, Ludwig Straub, and Amir Sufi. 2022. “A Goldilocks Theory of Fiscal Deficits.” Princeton University.

Reis, Ricardo. 2021. “The constraint on public debt when $r < g$ but $g < m$.” Bank for International Settlements Working Paper 939.

Woodford, Michael. 1990. “The optimum quantity of money.” In *Handbook of Monetary Economics*. Vol. 2, , ed. Benjamin Friedman and Frank Hahn, Chapter 20, 1067–1152. Elsevier.