Capital and Labor Taxes with Costly State Contingency*

Alex Clymo[§] Andrea Lanteri[¶] Alessandro T. Villa[‖]

November 2020

Preliminary and Incomplete

Abstract

We analyze optimal capital and labor taxes in a model where (i) the government makes noncontingent announcements about future policies, and (ii) ex-post state-contingent deviations from these announcements are costly. We find that costly state contingency has important implications for the response of taxes and allocations to government spending shocks. Different from previous results based on freely state-contingent taxes, the volatility of capital taxes is low and labor taxes play a fundamental role in accommodating fiscal shocks, increasing persistently when government spending increases. Moreover, private consumption becomes highly responsive to government spending. We also characterize optimal fiscal announcements. Under Full Commitment, announcements are unbiased, i.e., they coincide with expected policies. When governments lack commitment, instead, fiscal announcements play a strategic role and governments use them to constrain future policies; as a result, optimal fiscal announcements are biased, but may sustain similar outcomes to the ones associated with Full Commitment.

Keywords: Optimal Fiscal Policy; Fiscal Announcements; Costly State Contingency; Time Inconsistency.

*We thank Albert Marcet, Morten Ravn, as well as conference and seminar participants at Computation in Economics and Finance 2018 and IAE-CSIC for helpful comments.

[§]University of Essex. Email: a.clymo@essex.ac.uk
[¶]Duke University and CEPR. Email: andrea.lanteri@duke.edu
[‖]Duke University. Email: alessandro.villa@duke.edu

"Read my lips: no new taxes." (George H.W. Bush, 1988)

1 Introduction

Changes in fiscal policy, such as reforms of the tax code, are costly endeavors for governments, as they typically require parliamentary approval, which may involve lengthy negotiations and lead to revisions. As a result, there are often lags and substantive differences between tax-policy announcements, which are based on an expected evolution of the state of the economy, and realized policies. This policy framework limits the degree to which fiscal policy can respond contemporaneously to shocks hitting the economy.

Whereas the literature on optimal fiscal policy has devoted considerable attention to limitations in state contingency of government *debt*, it is standard to make rather stark assumptions on the degree of state contingency of *taxes*, often for convenience. In most cases, the literature assumes that the government can freely change taxes in response to contemporaneous shocks; in some cases, instead, it assumes that labor taxes can freely adjust, whereas capital taxes cannot adjust at all.[1]

As a result, the implications of the realistic restriction that it is costly to adjust policies in response to shocks are not, to our knowledge, well understood. What are the positive and normative implications if the government finds it costly to set policies in a state contingent way, and easier to communicate simple promises for what policies will be in the future *regardless of future shocks*? What causes governments to deviate from these promises ex-post? What are the implications of *costly state contingency* of fiscal plans on optimal policy, equilibrium allocations, and strategic interactions across successive governments with limited commitment?

To fill this gap, in this paper we develop a new framework to analyze optimal capital and labor taxes when governments make non-contingent announcements about future policies, and ex-post state-contingent deviations from these announcements are costly. We find that costly state contingency has important positive and normative implications. We first develop our theory of *optimal fiscal announcements* in this framework, and characterize how governments make announcements depending on standard commitment assumptions. We then apply our framework to the workhorse model of capital and labor taxation, and show that costly state contingency brings the behavior of optimal taxes in response to fiscal

[1] We provide a detailed discussion of the related literature below.

shocks closer to the data. Crucially, we find that the difficulty in setting policy in a fully state contingent way does not simply induce governments to implement otherwise-optimal policies with a lag, but fundamentally changes the policy trade-off, inducing qualitatively and quantitatively different policy responses to shocks.

To fix ideas, in our costly-state-contingency framework, for a vector of policy instruments τ_t, we allow the government at time t to make an announcement $\overline{\tau}_t$ about the policy they will implement at time $t+1$. This announcement is not allowed to differentiate the policy depending on the realizations of shocks that may hit the economy at time $t+1$ and is therefore noncontingent. At time $t+1$, the government can choose any value for the instruments, τ_{t+1}, in response to the shock that materializes. However, to do so, the government must incur a cost that depends on the difference between realized policy and previous announcement. The higher this cost is, the more costly it is for the government to make policies state contingent, and the closer realized policies will be to the noncontingent announcement.

Our analysis proceeds in two parts. First, we develop a two-period model that allows us to provide transparent insights on the role of costly state contingency for optimal fiscal announcements, realized taxes, and allocations. We also characterize the strategic interactions that arise when governments lack commitment and analyze the welfare effects of costly state contingency. Under Full Commitment, announcements are *unbiased*, meaning that the government simply announces the "average" policy it expects to make next period. This announcement minimizes the ex-post cost of state contingency.

When governments lack commitment, instead, fiscal announcements play a strategic role and governments use them to influence the policies of future policymakers. As a result, optimal fiscal announcements are *biased*, meaning that the government makes an announcement which is different from the average policy that it expects the subsequent government to implement. For example, if the time-$t+1$ government is expected to set capital taxes inefficiently high from a time-t perspective, the time-t government will announce an even lower capital tax in order to try and drag down the realized tax the subsequent government sets. Thus, when governments lack commitment, they may make promises that they do not expect to be kept, in order to at least nudge the policies of future governments in their desired direction. The restrictions implied by costly state contingency therefore act as a form of limited commitment.

These results have implications for the design of constitutions that constrain governments, trading off flexibility versus commitment. When governments lack commitment, a constitution that makes it harder to quickly deviate from past commitments sacrifices short term flexibility—i.e., state contingency of policies—in exchange for the welfare benefits of greater commitment. We formalize this intuition in a constitutional design exercise. When governments have no commitment, we prove that the optimal degree of costly state contingency will be positive if the degree of uncertainty is sufficiently low. This is in contrast to Full Commitment, where there are by assumption no commitment benefits of constraining the government, and the optimal degree of costly state contingency is always zero.

In the second part of the paper, we consider an infinite-horizon model of optimal capital and labor taxation subject to a balanced budget, which we calibrate to closely match salient features of US post-war data on fiscal variables. We use this model to both (i) demonstrate how a government facing costly state contingency would choose fiscal announcements and policy in response to shocks, and to (ii) show that a realistically calibrated degree of costly state contingency brings the predictions of optimal policy in this framework closer to the data on the conduct of actual policies.

We explore how costly state contingency changes optimal government policy in response to shocks to government spending. We begin by describing our results for a government with Full Commitment. With free state contingency of taxes, when government spending increases, the optimal response would be to absorb the shock with a capital-tax hike. In contrast, in our model, adjusting current taxes from their promised value is costly, and the government chooses to move them less. Since governments adjust current taxes less in response to a spending shock and the government budget constraint must be satisfied, the tax base must adjust. Thus, the government relies on fiscal announcements about future policies that induce a higher current level of output. Specifically, the government announces that it will cut capital taxes and raise labor taxes next period.

This mechanism leads to a path for realized taxes that is significantly different from the standard prediction of optimal policy models where taxes can be costlessly adjusted, or models where capital taxes are predetermined. In particular, different from previous results, and consistent with empirical evidence, the volatility of capital taxes is low, and labor taxes play a fundamental role in accommodating fiscal shocks, increasing persistently when government spending increases. As the government finds it difficult to insure the allocation from fiscal shocks, private consumption becomes highly responsive to government spending.

Finally, we analyze the Time-Consistent policy, by deriving a Generalized-Euler-Equations representation of the optimality conditions when governments lack commitment, and solve our model under this alternative commitment assumption. In contrast to the model without costly state contingency, the dynamics of taxes under Full-Commitment and Time-Consistent policy are actually quite similar in our framework, although capital taxes are lower on average when there is Full Commitment. This is because the cost of deviating from past promises inherent in the costly state contingency framework partially substitutes for a commitment technology.

Related Literature. Our paper contributes to the large literature on optimal capital and labor income taxes, by focusing on a novel friction in the government problem, namely costly state contingency of tax plans. We highlight the importance of this friction both for models of fiscal policy with Full Commitment and for models of Time-Consistent policy. Furthermore, our paper contributes to the literature on the macroeconomic effects of fiscal rules, such as balanced-budget constraints, and fiscal announcements.

Optimal Capital and Labor Taxes with Full Commitment. Chari and Kehoe (1999) analyze optimal capital and labor income taxation under Full Commitment and complete markets and show that there is indeterminacy between the realization of state-contingent capital tax rates and the portfolio of state-contingent securities. Furthermore, average capital taxes are close to zero, consistent with the early findings of Judd (1985) and Chamley (1986). Since their work, several papers study optimal capital and labor taxes in the presence of incomplete financial markets. Closely related to our paper, Stockman (2001) studies optimal capital and labor taxes under a balanced-budget rule, assuming tax rates are fully state contingent; Fahri (2010) considers a more general incomplete-markets model with non-contingent debt as in Aiyagari, Marcet, Sargent, and Seppala (2002), and assumes that the capital tax is predetermined—i.e., not state contingent—whereas the labor tax is fully state contingent.[2]

We build on the insights of this body of work and generalize the benchmark model by introducing costs of state contingency for tax rates, thereby allowing for an intermediate, arguably more realistic, degree of state contingency, and nesting the previous assumptions on the timing of taxes. In our calibrated model, when government spending increases, the government engineers a short-lived capital-tax hike, which is then reversed, and increases the labor tax persistently. These dynamics are broadly consistent with the empirical evidence (Burnside, Eichenbaum, and Fisher, 2004).

[2]A related literature explores the degree to which imperfectly state-contingent debt instruments, such as government bonds with different maturities, can be used to absorb fiscal shocks. See, for instance, Faraglia, Marcet, Oikonomou, and Scott (2019).

Time-Consistent Fiscal Policy. Building on the insight that optimal capital taxation is generally time inconsistent, Klein, Krusell, and Ríos-Rull (2008) study optimal capital taxation and public-good provision when the government lacks commitment. They characterize the Time-Consistent optimal policy using a Generalized Euler Equation. We build on their approach and introduce a trade-off between partial commitment and state contingency in a stochastic environment. Klein and Ríos-Rull (2003) and Martin (2010) analyze time-consistent capital and labor taxes.[3]

Our work is closely related to the literature on intermediate notions of fiscal commitment. For instance, Debortoli and Nunes (2010, 2013) analyze models of fiscal policy with stochastic government re-optimizations. In our model, the degree to which governments renege on previous announcements is an endogenous choice that depends on the state of the economy. Clymo and Lanteri (2020) introduce a framework in which the government has Limited-Time Commitment—i.e., successive governments commit to tax plans over a finite future horizon—and find that a short commitment horizon may be sufficient to sustain Full-Commitment outcomes. In this paper, we generalize their framework by allowing governments to partially renege on previous noncontingent announcements, subject to a cost. In so doing, our theory introduces a meaningful distinction between optimal fiscal announcements and realized policies. Furthermore, in terms of application, this paper focuses on the trade-off between capital and labor taxes in a stochastic production economy.

Fiscal Announcements. Our analysis is related to the empirical literature that studies the macroeconomic effects of announcements about future fiscal plans, distinguishing them from actually implemented fiscal policies. See, for instance Mertens and Ravn (2012) and Alesina, Favero, and Giavazzi (2015). Our theoretical framework distinguishes between strategic announcements about future tax rates, and actually implemented tax rates, and allows us to develop of theory of optimal fiscal announcements under uncertainty, both with and without commitment. Thus, this paper builds a bridge between the theoretical literature on optimal taxation and the empirical literature on the effects of expectations about fiscal policy.

Fiscal Rules. Our work is also related to the theoretical and quantitative body of work that studies the macroeconomic effects of fiscal rules and their optimal design. In an early contribution, King, Plosser, and Rebelo (1988) find that balanced-budget rules amplify aggregate fluctuations. Schmitt-Grohe and Uribe (1997) find that a balanced-budget rule

[3]Relatedly, Karantounias (2019) uses a Generalized-Euler-Equation approach to analyze optimal taxation in a model with default on non-contingent debt.

may induce indeterminacy in a standard neoclassical production economy. We find that even in a model without indeterminacy, balanced-budget rules, combined with costly state contingency of taxes, induce large fluctuations in consumption.[4]

A theoretical literature studies the optimal design of policy rules, when the policymaker has private information about the economy and limited commitment. See, for instance, Athey, Atkeson, and Kehoe (2005) for a monetary model, and Halac and Yared (2014) for a model of fiscal policy with persistent shocks. The optimal institutional arrangements in these papers involve limits on the degree of state contingency in policy. Different from this literature, we do not explicit microfound the origins of limited state contingency, and focus instead on the effect of costly state contingency on the dynamic trade-offs between capital and labor taxes. Our approach is consistent with the notion that partial state contingency in fiscal policy may arise because of multiple reasons, such as institutional constraints on policy implementation, which may in turn result from the frictions highlighted in this literature, or partial information about the state of the economy, as in Hauk, Lanteri, and Marcet (2020).

The rest of the paper is organized as follows. Section 2 presents the two-period model and derives analytical insights on the role of costly state contingency. Section 3 analyzes the infinite-horizon model. Section 4 presents our numerical results. Section 5 concludes.

2 Two-Period Model

In this section we analyze a two-period model of optimal capital and labor taxes with costly state contingency. We use this simple framework to build intuition on the main trade-offs and we also establish some formal results on the role of costs of state contingency for optimal policy. Importantly, we distinguish between the case of Full Commitment and the case in which the government optimizes sequentially (Time-Consistent policy). In the interest of concreteness, we present a model with specific preferences and technology, but we prove the formal results in a general framework in Appendix A.

[4] Our findings under balanced budget may also be relevant for models with noncontingent debt, as long as the economy is sufficiently close to a debt limit, so that debt issuance is an imperfect substitute for a tax adjustment. A natural direction for future work is to combine our analysis of costly state contingency of taxes with a richer model of incomplete financial markets.

2.1 Competitive Equilibrium and Implementability Constraints

There are two dates, $t = 0, 1$. At $t = 0$, households make an investment decision and the government makes fiscal announcements. At $t = 1$, the stochastic level of government spending is realized, production takes place and the government raises capital and labor income taxes to finance government spending. We refer to variables as variables at $t = 1$ with no subscripts, and we index $t = 0$ variables with subscript 0.

A representative household has utility function

$$c_0 + \beta \mathbb{E} \left(\log(c) - \chi \frac{l^{1+\eta}}{1+\eta} \right), \tag{1}$$

where c_0 and c denote consumption at the two dates, and l is labor effort. We assume $\beta \in (0, 1)$, $\chi > 0$, and $\eta > 0$.

The resource constraints are

$$\begin{aligned} c_0 + k &= y_0, && (2) \\ c + g &= z k^{\alpha} l^{1-\alpha}, && (3) \end{aligned}$$

where y_0 is an exogenous endowment, which we assume to be sufficiently large to ensure positive consumption, k is capital, which fully depreciates in our period, and $\alpha \in (0, 1)$. Government spending g is a random variable with exogenous distribution $G(g)$.

Competitive firms hire labor and rent capital, resulting in each factor being compensated with its marginal product. The government budget constraint thus reads

$$\left(\alpha \tau^k + (1 - \alpha) \tau^l \right) z k^{\alpha} l^{1-\alpha} \geq g, \tag{4}$$

where τ^k and τ^l are proportional tax rates on capital income (for simplicity, without deduction for depreciation) and labor income respectively. We allow the left-hand side of equation (4) to be larger than the right-hand side, in which case the government transfers its positive surplus to households in a lump-sum fashion. In equilibrium, this transfer will equal zero. In principle, the government may set these taxes as state-contingent functions of the shock, g, and we suppress the dependence on g where notationally convenient.

The household optimality conditions with respect to labor supply at $t = 1$ and investment at $t = 0$, combined with equilibrium factor prices, give

$$\chi l^{\eta} c = (1-\alpha) z k^{\alpha} l^{-\alpha} \left(1-\tau^{l}\right), \tag{5}$$

$$1 = \beta \mathbb{E}\left[c^{-1}\left(1-\tau^{k}\right) \alpha z k^{\alpha-1} l^{1-\alpha}\right]. \tag{6}$$

Equation (6) is the standard Euler equation for capital, which implies the usual time-inconsistency government for a government setting taxes. In particular, time 1 capital taxes, τ^k, appear in the Euler equation, which constrains the time-0 government. We can combine equations (5) and (6) with the government budget constraint (4) to derive the following two implementability constraints. Firstly, a labor supply optimality condition,

$$\chi l^{\eta+\alpha}\left(z k^{\alpha} l^{1-\alpha} - g\right) = (1-\alpha) z k^{\alpha}\left(1-\tau^{l}\right), \tag{7}$$

which defines implicitly a function $l = h(k, g, \tau^l)$ with $h_{\tau^l} < 0$, and holds state-by-state for each realised g. Secondly, an Euler equation for capital investment

$$k \leq \beta\left[1 - \chi \mathbb{E}\left(h(k, g, \tau^{l})\right)^{1+\eta}\right], \tag{8}$$

which we express as an inequality due to the governments option of giving a lump sum tax rebate. Given a choice of labor tax τ^l for each realization of g, private-sector allocations must satisfy constraints (7) and (8), and the associated capital tax can be then obtained using (4).

2.2 Optimal Policy with Full Commitment

We now characterize optimal policy under the assumption that a government at $t = 0$ formulates a plan under Full Commitment, but faces costly state contingency. The key feature of our framework is that the government must make an announcement at $t = 0$ about policies to be implemented at $t = 1$, and then incurs a cost if it chooses to deviate from this announcement ex-post.

Specifically, at $t = 0$ the government makes noncontingent fiscal announcements for capital and labor taxes $\overline{\tau}^k$ and $\overline{\tau}^l$ respectively. The government also chooses state-contingent taxes τ^k and τ^l, to be implemented at $t = 1$. The government chooses announcements and policies, as well as allocations, to maximize

$$c_0 + \beta\mathbb{E}\left[\log(c) - \chi\frac{l^{1+\eta}}{1+\eta} - \frac{\gamma^k}{2}(\tau^k - \overline{\tau}^k)^2 - \frac{\gamma^l}{2}(\tau^l - \overline{\tau}^l)^2\right], \tag{9}$$

where $\gamma^k \geq 0$ and $\gamma^l \geq 0$ are parameters that determine the *costs of state contingency* in taxes. We model these as pure utility costs to the government, which act as a reduced form for any institutional features or frictions that constrain the ability of the government to set policies in a state contingent way. In particular, the costs are quadratic functions of the distance between realized tax rates and noncontingent announcements. When $\gamma^k = 0$ and $\gamma^l = 0$ the problem reduces to the standard Full Commitment problem. We consider generic cost functions in our infinite horizon model, and restrict to quadratic here for simplicity.[5]

The government maximization problem is subject to the resource constraints and the implementability constraints derived above. We denote by μ the multiplier on (8), ν the multiplier on (4), and directly substitute in $l = h(k, g, \tau^l)$ and the resource constraints, (2) and (3). The government chooses capital and labor taxes to implement contingent on the realized state. For each value of g, the first-order conditions with respect to capital and labor taxes give

$$\nu\alpha z k^\alpha l^{1-\alpha} = \gamma^k(\tau^k - \overline{\tau}^k) \tag{10}$$

$$\begin{aligned}\left[(1-\alpha)zk^\alpha l^{-\alpha}\left(c^{-1} + \nu(\alpha\tau^k + (1-\alpha)\tau^l)\right) - \chi l^\eta(1+\mu(1+\eta))\right] h_{\tau^l}(k, g, \tau^l)&\\ + \nu(1-\alpha)zk^\alpha l^{1-\alpha} = \gamma^l(\tau^l - \overline{\tau}^l),&\end{aligned} \tag{11}$$

For both taxes, the government trades off the effect of the tax on the private-sector allocation (on the left-hand side) with the marginal cost of state contingency (on the right-hand side). In particular, for a given realization of g, the government understand that deviating the tax τ^j from the promise $\overline{\tau}^j$ involves paying a cost, which is balanced against the benefit of achieving the best ex-post allocation. For capital taxes, the tax simply trades off state contingency costs versus the effect on the budget, through the multiplier ν. For the labor tax, there are additional effects on the direct allocation. Finally, the multiplier μ captures the government's forward looking understanding that time 1 policies affect investment at

[5]Equation (9) highlights that in assuming that the costs of state contingency appear in the government objective function, we make a slight departure from the standard assumption of purely benevolent government, and allow for a difference in the objective function of the government and that of households. However, because households take tax rates as given, nothing would change if we also added these costs in the household utility function (1).

time 0.

The announced taxes are not allowed to vary by state, and are thus effectively chosen in advance. The first order conditions with respect to the optimal tax announcements give

$$\overline{\tau}^k = \mathbb{E}\tau^k, \tag{12}$$

$$\overline{\tau}^l = \mathbb{E}\tau^l. \tag{13}$$

Thus, we find that optimal tax announcements under Full Commitment are unbiased forecasts of future tax rates. By formulating these announcements, the government minimizes the expected costs of state contingency. Intuitively, the government announces the average policy it intends to set next period, which minimizes the cost of its planned deviations from this announcement.

We can gain further intuition on the role played by costly state contingency by taking expectations of the first order conditions (10) and (11) and using (12) and (13) to give

$$\mathbb{E}\nu\alpha z k^{\alpha} l^{1-\alpha} = 0 \tag{14}$$

$$\mathbb{E}\Big\{ \left[(1-\alpha)zk^{\alpha}l^{-\alpha}\left(c^{-1}+\nu(\alpha\tau^k+(1-\alpha)\tau^l)\right)-\chi l^{\eta}(1+\mu(1+\eta))\right] h_{\tau^l}(k,g,\tau^l) + \nu(1-\alpha)zk^{\alpha}l^{1-\alpha}\Big\} = 0. \tag{15}$$

In the standard model with full state contingency, we have $\gamma^k = \gamma^l = 0$ and the right-hand sides of (10) and (11) are both zero. Comparing this to the above two equations, we see that with costly state contingency the first order conditions are now only equal to zero *on average*. Thus, there is a sense in which, under Full Commitment, costly state contingency preserves how policy is set on average, while introducing a wedge for each specific realization of the shock.[6]

2.3 Optimal Time-Consistent Policy

We now characterize the optimal policy when the government at $t = 0$ can make noncontingent announcements $\overline{\tau}^k, \overline{\tau}^l$, but cannot commit to future state-contingent taxes τ^k, τ^l; instead, the government at $t = 1$ acts under discretion. This assumption implies that the

[6]Notice that the averaging only applies to the first-order condition; since the model is nonlinear, it is possible for costly state contingency to have effects on the average policy choices themselves.

model can be described as a game with a strategic interaction between the government choosing ex-ante announcements and another government choosing ex-post taxes. We proceed by backward induction and start by discussing the government problem at $t = 1$, after government spending is realized.

Time-1 problem: The government at $t = 1$ takes as given the state variables $k, g, \overline{\tau}^k, \overline{\tau}^l$ and chooses taxes and allocations to maximize

$$\log(c) - \chi \frac{l^{1+\eta}}{1+\eta} - \frac{\gamma^k}{2}(\tau^k - \overline{\tau}^k)^2 - \frac{\gamma^l}{2}(\tau^l - \overline{\tau}^l)^2, \tag{16}$$

subject to the budget constraint (4) with multiplier ν, the implementability constraint (7), and the resource constraint for $t = 1$.

The first-order conditions with respect to the tax rates on capital (for $\gamma^k > 0$) and labor income are

$$\nu \alpha z k^\alpha l^{1-\alpha} = \gamma^k(\tau^k - \overline{\tau}^k) \tag{17}$$

$$\begin{aligned}\left[(1-\alpha) z k^\alpha l^{-\alpha}\left(c^{-1} + \nu(\alpha\tau^k + (1-\alpha)\tau^l)\right) - \chi l^\eta\right] h_{\tau^l}(k, g, \tau^l) \\ + \nu(1-\alpha) z k^\alpha l^{1-\alpha} = \gamma^l(\tau^l - \overline{\tau}^l)\end{aligned} \tag{18}$$

for all g. In choosing taxes, the government trades off the marginal costs of state contingency with the additional tax revenue, and, in the case of the labor tax, its effect on the allocation. Implicitly, these optimality conditions define a policy function $\tau^l = \tilde{\tau}^l(k, g, \overline{\tau}^k, \overline{\tau}^l)$. In the case $\gamma^k = 0$, the capital tax is effectively lump-sum, and thus the solution is $\tau^l = 0$ and τ^k satisfies the budget constraint, with $\nu = 0$.

Let $\tilde{W}(k, g, \overline{\tau}^k, \overline{\tau}^l)$ be the value function attained in the government maximization problem at $t = 1$. The envelope conditions with respect to fiscal announcements are given by

$$\tilde{W}_{\overline{\tau}^k} = \gamma^k(\tau^k - \overline{\tau}^k), \tag{19}$$

$$\tilde{W}_{\overline{\tau}^l} = \gamma^l(\tau^l - \overline{\tau}^l). \tag{20}$$

Time-0 problem: We now discuss the problem of the government at $t = 0$. The government chooses announcements $\overline{\tau}^k$ and $\overline{\tau}^l$, as well as, indirectly, private-sector investment k to maximize

$$c_0 + \beta \mathbb{E} \tilde{W}(k, g, \overline{\tau}^k, \overline{\tau}^l), \tag{21}$$

subject to the resource constraint at $t = 0$ and the implementability constraint

$$k \leq \beta \left[1 - \chi \mathbb{E} \left(h(k, g, \tilde{\tau}^l) \right)^{1+\eta} \right], \tag{22}$$

where we leave implicit the dependence of $\tilde{\tau}^l$ on the state variables at $t = 1$ to simplify notation. We denote by μ the multiplier on this constraint.

Taking the first-order condition with respect to the announcements and using the envelope conditions, we obtain the following optimality conditions:[7]

$$\overline{\tau}^k = \mathbb{E}\tau^k - \frac{\chi(1+\eta)\mu}{\gamma^k} \mathbb{E}\left[l^\eta h_{\tau^l}(k, g, \tau^l) \tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) \right], \tag{23}$$

$$\overline{\tau}^l = \mathbb{E}\tau^l - \frac{\chi(1+\eta)\mu}{\gamma^l} \mathbb{E}\left[l^\eta h_{\tau^l}(k, g, \tau^l) \tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) \right]. \tag{24}$$

Comparing these to the optimal promises made under Full Commitment, (12) and (13), we see a crucial difference. Under Full Commitment, the government announces taxes equal to the average of ex-post realized taxes. In the Time-Consistent case, this is no longer true, and the promises are "biased". These biases are introduced in order to manipulate the actions of the future government, who sets taxes "incorrectly" (from the perspective of $t = 0$) due to the lack of commitment.

Specifically, the government at $t = 0$ understands that by marginally decreasing future labor it can relax its current implementability constraint (22). Thus, intuitively, the biases in (23) and (24) depend on the product of the effects of each announcement on realized labor taxes and the effect of labor taxes on labor. These biases have an intuitive sign. Starting with capital taxes, we would expect that the government without commitment will set capital taxes too high, not internalizing that this will lower investment. To reduce this effect, the time-0 government will announce a lower capital tax, in order to try and bias downwards the capital taxes the time-1 government eventually sets. This intuition is confirmed in (23): We have $\mu > 0$ and $h_{\tau^l}(k, g, \tau^l) < 0$, and so as long as raising the capital tax promise, $\overline{\tau}^k$ raises the implemented capital tax and lowers the implemented labor tax ($\tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) < 0$) we have $\overline{\tau}^k < \mathbb{E}\tau^k$. A similar logic implies that the government biases upwards labor taxes ($\overline{\tau}^l > \mathbb{E}\tau^l$) in (24) for a symmetric reason.

It is worth noting that the choice to bias the announced taxes is not costless, and reflects a tradeoff the time-0 government faces. For example, if the time-0 government biases the

[7] In the interest of space, we formulate the first-order condition and the envelope condition with respect to capital in Appendix A.

capital tax promise downwards, it can lower average capital taxes, which benefits ex-ante welfare. However, this comes at the cost of raising the ex-post cost of state contingency, which is instead minimized by under Full Commitment.

2.4 Costly State Contingency and Welfare

When the government acts with discretion, the fact that the initial government biases its promised capital tax in order to influence the choices made by future governments suggests a new trade-off between commitment and state contingency in our model, which has not been explored in the literature.

In this section we explore this trade-off formally, by interpreting our costs of state contingency as a constitutional constraint which could be chosen ex-ante by a social planner trying to ensure the best outcome. We show that such a constraint can only have a negative effect on welfare if the government has Full Commitment, but that it may improve welfare if the government does not, and acts with discretion. In order to establish these results more transparently, we focus on the case where the cost of state contingency applies to only one tax instrument, for instance $\gamma^k = 0$, $\gamma^l = \gamma > 0$, so we vary the cost of state contingency with a single parameter, γ.

Our first result is that if the government acts with Full Commitment, increasing the cost of state contingency, γ, must reduce welfare. Note that we have two notions of welfare: The welfare of the representative household does not include the cost of state contingency. The government objective includes household welfare and the cost of state contingency.

Proposition 1 *Under Full Commitment, household welfare (i.e., excluding the cost of state contingency) and government welfare (i.e., including the cost) are both decreasing in* γ.

The proof is in Appendix A, where we derive the result for a general model nesting this one. We provide an intuitive discussion here. Under Full Commitment, the government is already acting in the best interests of the household, given the institutional constraints. Increasing the cost of state contingency can only make the household worse off, because it forces the government to make policy less state contingent than it should be. In this case, the cost of state contingency is simply a utility cost placed on the government, which forces it to make sub-optimal decisions for the policies which affect actual household welfare. Accordingly, if one considers an institutional design problem, where a constitutional planner chooses the ex-ante optimal level of γ, it must be that the optimum is for no costs of state contingency ($\gamma = 0$) if the government acts under Full Commitment.

Our second result is that this is no longer the case under discretion, since now increasing the cost of state contingency might actually raise welfare by helping overcome time inconsistency issues.

Proposition 2 *Under the Time-Consistent policy, household welfare is increasing in γ at $\gamma = 0$, implying that the value of γ that maximizes welfare is positive, for a sufficiently low level of uncertainty.*

The proof is in Appendix A, where we derive the result for a more general model nesting this one. We provide an intuitive discussion here. Increasing the cost of state contingency has an additional effect when the government lacks commitment, which is to help twist the actions of the time-1 government, and prevent them from raising the capital tax ex-post. Since the allocation is not optimal from a time-0 perspective, it could be that the benefit of helping bind the future government to lower capital taxes is large enough to outweigh the cost of making the capital tax less state contingent.

Since there are now two opposing effects of raising γ on welfare, it is not obvious what the optimal level of γ would be, if set ex-ante by a constitutional planner. Could it be that a positive level of cost of state contingency ($\gamma > 0$) is now optimal, in contrast to the result under Full Commitment? It turns out that this is indeed the case. As the proposition argues, we show that, for a low enough level of uncertainty, the welfare gain to the representative household of raising γ is positive at $\gamma = 0$. This implies that for a low enough level of uncertainty, the optimal level of γ must be positive.

To understand this result, consider first a version of the model with no uncertainty, so that there is only a single known value of g at $t = 1$. In this case, marginally increasing γ when $\gamma = 0$ must be welfare enhancing. This is because there is no cost from loss of state contingency when there is no uncertainty, and raising γ makes the time-1 government marginally lower the capital tax down from its too-high level in the Time-Consistent game. If instead uncertainty was very large, it could conceivably be the case that any loss in state contingency would lead to welfare losses larger than the gains from lowering capital taxes. Thus, the level of uncertainty is crucial for determining whether there are gains to be had from raising γ.

We establish that, when there is no uncertainty, there must be a strict rise in welfare as γ is increased. By continuity, this implies that there must exist low enough level of uncertainty for which welfare is increased as γ is raised away from zero. These two propositions establish that costly state contingency is always welfare reducing under Full Commitment, but may raise ex-ante welfare under discretion if uncertainty is low enough.

3 Infinite-Horizon Model

In this section, we describe our infinite-horizon model and analyze the optimality conditions of the government in the presence of costs of state contingency for fiscal plans. An important difference relative to our two-period model is that in a dynamic framework governments optimally use future announcements to affect the current allocation in response to shocks hitting the economy.

3.1 Environment

We consider a stochastic production economy populated by a continuum of identical households and a government. Time is discrete and infinite, $t = 0, 1, 2, ...$ Households rank streams of consumption c_t and labor l_t according to the following utility function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - v(l_t)\right] \tag{25}$$

where $\beta \in (0, 1)$ is the discount factor, $u_c > 0$, $u_{cc} < 0$, $v_l > 0$, $v_{ll} > 0$.

The resource constraint of the economy is given by

$$c_t + k_t + g_t = F(k_{t-1}, l_t) + (1 - \delta)k_{t-1} \tag{26}$$

where k_t is capital, subject to a one-period time to build and depreciation rate $\delta \in (0, 1)$, f is a constant-returns-to-scale production function, and g_t is exogenous, stochastic government spending. We assume that g_t follows a discrete Markov process with transition probability matrix P_g. We denote by $g^t \equiv \{g_0, g_1, ..., g_t\}$ a history of realizations of government spending. To simplify notation, we avoid explicitly denoting allocations as functions of histories g^t, but it is understood that c_t, l_t, and k_t are measurable with respect to g^t.

Households demand consumption goods, supply labor and trade claims on the aggregate capital stock. The household budget constraint reads

$$c_t + k_t = w_t l_t(1 - \tau_t^l) + k_{t-1}\left[1 + r_t(1 - \tau_t^k)\right] \tag{27}$$

where w_t is the wage, r_t is the gross rate of return on capital, and τ_t^l and τ_t^k are proportional tax rates on labor and capital income respectively.[8]

3.2 Household and Firm Optimality

Households maximize utility (25) subject to their budget constraint (27). The intratemporal labor-consumption margin and the Euler equation for savings are

$$\begin{aligned} v_l(l_t) &= u_c(c_t)w_t(1-\tau_t^l) &(28)\\ u_c(c_t) &= \beta\mathbb{E}_t u_c(c_{t+1})\left[1+r_{t+1}(1-\tau_{t+1}^k)\right] &(29)\end{aligned}$$

Competitive firms rent capital and hire labor to maximize profits. Thus, factor prices are tied to marginal products as follows

$$\begin{aligned} w_t &= F_l(k_{t-1}, l_t) &(30)\\ r_t &= F_k(k_{t-1}, l_t) - \delta &(31)\end{aligned}$$

Our notation already implicitly imposes market clearing for labor and capital. The definition of a competitive equilibrium, for given sequences of tax rates, is standard.

3.3 Government

The government needs to finance spending g_t using capital and labor income taxes, subject to a balanced-budget constraint.[9]

$$\tau_t^k r_t k_{t-1} + \tau_t^l w_t l_t = g_t \tag{32}$$

At date t, the government chooses current tax rates τ_t^k and τ_t^l. It does so with knowledge of the current shock, and so these policies are measurable with respect to g^t. Furthermore, it formulates *announcements* about future (one-period ahead) tax rates, which we denote by $\overline{\tau}_t^k$ and $\overline{\tau}_t^l$. Importantly, these announcements are not allowed to be contingent on the

[8]We could also equivalently allow households to trade risk-free bonds among themselves. As we will impose a balanced-budget constraint on the government, these bonds would be in zero net supply. Because households are identical, in equilibrium no trade in these bonds would occur, making the presence of these bonds immaterial for equilibrium allocations.

[9]In Appendix D, we discuss the implications of costly state contingency in taxes in a model with state-contingent government debt.

future exogenous state of the economy, and so are also measurable with respect to g^t.

Given initial conditions $k_{-1}, \overline{\tau}^k_{-1}, \overline{\tau}^l_{-1}$, the government chooses sequences of current tax rates τ^k_t, τ^l_t and future announcements $\overline{\tau}^k_t, \overline{\tau}^l_t$ to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - v(l_t) - \Gamma^k(\tau^k_t, \overline{\tau}^k_{t-1}) - \Gamma^l(\tau^l_t, \overline{\tau}^l_{t-1})\right] \tag{33}$$

where $\Gamma^j(\tau^j_t, \overline{\tau}^j_{t-1})$ is a *cost of state contingency* for tax rate τ^j_t and $j \in \{k, l\}$. Recall that τ^j_t is measurable with respect to g^t, whereas the announcement $\overline{\tau}^j_{t-1}$ is measurable with respect to g^{t-1}. We assume that: (i) $\Gamma^j(\tau^j_t, \overline{\tau}^j_{t-1}) = 0 \geq 0$; (ii) $\Gamma^j(\tau^j_t, \overline{\tau}^j_{t-1}) = 0$ if $\tau^j_t = \overline{\tau}^j_{t-1}$; (iii) Γ^j is weakly increasing and weakly convex in a measure of distance between τ^j_t and $\overline{\tau}^j_{t-1}$. Thus, the cost functions Γ^j penalize deviations of state-contingent tax rates relative to the previously announced non-contingent plan. In our numerical application, we will parameterize Γ^j as a quadratic function of $(\tau^j_t - \overline{\tau}^j_{t-1})$. However, we emphasize that our framework is general and could accommodate other functional forms, including, for instance, fixed costs.

It is also important to note that our costly state contingency framework is starkly different from a standard adjustment cost framework. In an adjustment cost framework, the government would face an, e.g., quadratic cost of adjusting taxes relative to their value in the last period. Here, the cost is paid relative to the announced plan for taxes. Crucially, the government is free to announce any plan without cost: there is no penalty for setting next period's announcement, $\overline{\tau}^j_t$, differently from today's actual policy, τ^j_t. This means that policies are in principle able to be adjusted costlessly one period ahead, and it is only within the period that costly state contingency penalises adjustment. Thus, our framework truly restricts the governments ability to make policy state contingent one period in advance, while costlessly allowing longer term adjustments in policy in response to known information.[10]

3.4 Implementability Constraints

We now derive the implementability constraints of the government problem. This poses a novel challenge in our costly state contingency framework, since the government's objective

[10] To give a clear example of the difference, consider a world without uncertainty, and a change in required government spending which is perfectly anticipated one period in advance. Our costly state contingency framework allows the government to costlessly adjust its announced policy for the next period, in anticipation of this change in the economy. A normal adjustment cost would force policies to slowly adjust towards the new optimal policy, despite the new optimal policy being perfectly known in advance.

function now depends directly on taxes due to the cost functions. This rules out uses a pure "Dual approach", where the government's problem can be formulated directly in terms of choosing allocations rather than tax rates. We show how to adapt the Dual approach to our framework, allowing us to continue specifying the government's problem in terms of choosing allocations directly, with auxiliary functions linking these allocations back to implied policies.

In the absence of costs of state contingency, our general model specializes to the model analyzed by Stockman (2001), who shows that the balanced-budget constraint (32) can be combined with the private sector's first order conditions (28), (29), (30), and (31), to obtain a single implementability constraint expressed in terms of allocations only

$$u_c(c_t)k_t = \beta \mathbb{E}_t \left[u_c(c_{t+1})(c_{t+1} + k_{t+1}) - v_l(l_{t+1})l_{t+1} \right]. \tag{34}$$

for $t = 0, 1, \ldots$ This constraint, the resource constraint (26), and the balanced-budget constraint at $t = 0$ are sufficient to characterise the constraints placed on the government by competitive equilibrium.

Because of the presence of costs of state contingency, which explicitly depend on the realizations of the tax rates, the sequence of constraints (34) is not sufficient to characterize the government choice set in our general problem. Instead, we need to add the constraints (28) and (32), combined with equilibrium factor prices (30) and (31). As Clymo and Lanteri (2020) point out, these equations uniquely pin down the level of labor and consumption independently of the Euler equation (29), given (k_{t-1}, g_t) and a choice of contemporaneous tax rates (τ_t^k, τ_t^l). In particular, for any predetermined level of capital and exogenous realization of government spending, a choice of current tax rates must induce this allocation, in order to respect the government budget constraint and satisfy private sector optimality. In particular, for given $(k_{t-1}, g_t, \tau_t^l, \tau_t^k)$, there is a unique level of labor supply, l_t, required to generate a level of income consistent with the government budget constraint, given implicitly by the solution to

$$\tau_t^k \left(F_k(k_{t-1}, l_t) - \delta \right) k_{t-1} + \tau_t^l F_l(k_{t-1}, l_t) l_t - g_t = 0 \tag{35}$$

Given this level of labor supply, there is a unique level of consumption that ensures the level of labor obtained above is consistent with the household intratemporal optimality condition, given by

$$c_t = u_c^{-1}\left(\frac{v_l(h^l(k_{t-1}, g_t, \tau_t^l, \tau_t^k))}{F_l(k_{t-1}, l_t)(1-\tau_t^l)}\right) \tag{36}$$

We define two functions h^l and h^c to summarize the solutions for l_t and c_t to the above two equations for a given level of states and taxes:

$$l_t = h^l(k_{t-1}, g_t, \tau_t^l, \tau_t^k), \tag{37}$$

$$c_t = h^c(k_{t-1}, g_t, \tau_t^l, \tau_t^k). \tag{38}$$

These two equations can be solved for independently of any policy functions. We include them as additional constraints when formulating the government's maximization problem, in order to relate tax rates to allocation, and hence account for the contingency costs associated with policy choices.

3.5 Optimal Policy with Full Commitment

We now consider optimal fiscal policy under the assumption that the government has Full Commitment. The government chooses sequences of current taxes, future announcements, and allocations to maximize (33) subject to the resource constraint (26), with associated Lagrange multiplier λ_t, and the implementability constraint (34), with multiplier μ_t. To link the choice of allocations to the tax rates they require, we additionally must include the two constraints (37), and (38) with multipliers ν_t^l and ν_t^c respectively. These two constraints uniquely pin down the tax rates τ_t^k and τ_t^l required to implement a given allocation.[11]

The first-order conditions with respect to c_t and l_t are

$$\lambda_t = u_c(c_t) - \mu_t u_{cc}(c_t)k_t + \mu_{t-1}\left[u_{cc}(c_t)(c_t+k_t) + u_c(c_t)\right] - \nu_t^c, \tag{39}$$

$$v_l(l_t) = \lambda_t F_l(k_{t-1}, l_t) - \mu_{t-1}\left[v_{ll}(l_t)l_t + v_l(l_t)\right] + \nu_t^l. \tag{40}$$

The first-order condition with respect to k_t is

$$\begin{aligned}\lambda_t = \beta\mathbb{E}_t\lambda_{t+1}\left[F_k(k_t, l_{t+1}) + 1 - \delta\right] - \mu_t u_c(c_t) + \mu_{t-1}u_c(c_t) \\ - \beta\mathbb{E}_t\left[\nu_{t+1}^l h_k^l(k_t, g_{t+1}, \tau_{t+1}^k, \tau_{t+1}^l) - \nu_{t+1}^c h_k^c(k_t, g_{t+1}, \tau_{t+1}^k, \tau_{t+1}^l)\right]\end{aligned} \tag{41}$$

[11] An alternative "primal" formulation of the government problem can be obtained by using (28) and (32) to solve for (τ_t^k, τ_t^l) and plugging these values into the costs of state contingency. We find our formulation easier to interpret, because it explicitly allows us to analyze first-order conditions with respect to tax rates, but stress that the two formulations are equivalent.

Equations (39), (40), and (41) coincide with their respective counterparts in a model without costs of state contingency (Stockman, 2001), except for the presence of the multipliers ν_t^c and ν_t^l, which reflects the impact of marginal changes in the allocation on the costs of state contingency. Notice that the presence of multipliers on *past* implementability constraints (μ_{t-1}) in the optimality conditions is a manifestation of the time inconsistency of the plan under Full Commitment. We will thus explore alternative commitment assumptions in the next subsection.

The first-order conditions with respect to tax rates τ_t^k, τ_t^l are

$$\nu_t^c h_{\tau^k}^c(k_{t-1}, g_t, \tau_t^k, \tau_t^l) - \nu_t^l h_{\tau^k}^l(k_{t-1}, g_t, \tau_t^k, \tau_t^l) = \Gamma_{\tau^k}^k(\tau_t^k, \overline{\tau}_{t-1}^k), \tag{42}$$

$$\nu_t^c h_{\tau^l}^c(k_{t-1}, g_t, \tau_t^k, \tau_t^l) - \nu_t^l h_{\tau^l}^l(k_{t-1}, g_t, \tau_t^k, \tau_t^l) = \Gamma_{\tau^l}^l(\tau_t^l, \overline{\tau}_{t-1}^l). \tag{43}$$

Equations (42) and (43) highlight that the government trades off the effect of current taxes on allocations—and thus household utility—with their effect on the cost of state contingency. Notice that when costs of state contingency are removed ($\Gamma^k = \Gamma^l = 0$ for all inputs) the solution to the above equations requires that $\nu_t^c = \nu_t^l = 0$, and the FOCs reduce to those in Stockman (2001). Therefore, the above equations can be be interpreted as driving wedges relative to the standard allocation: For a given inherited promise $\overline{\tau}_t^j$, the government trades off the direct cost of setting the policy, τ_t^j, to a different value than the promise, against the wedges that a given choice induces in the FOCs relative to the standard solution.

The first-order conditions with respect to tax announcements $\overline{\tau}_t^k$, $\overline{\tau}_t^l$ are

$$\mathbb{E}_t \Gamma_{\overline{\tau}^k}^k(\tau_{t+1}^k, \overline{\tau}_t^k) = 0, \tag{44}$$

$$\mathbb{E}_t \Gamma_{\overline{\tau}^l}^l(\tau_{t+1}^l, \overline{\tau}_t^l) = 0. \tag{45}$$

Notice that the only effect of tax announcements on the government objective is through their effect on future costs of state contingency. Thus, as equations (44) and (45) show, the government optimally sets the expected future marginal cost of state contingency to zero to minimize the expected cost. To further interpret these conditions, consider the case of a quadratic function Γ^j, as we assume in our computations: $\Gamma^j(\tau^j, \overline{\tau}^j) \equiv \frac{\gamma_0^j}{2}(\tau^j - \overline{\tau}^j)^2$, with $\gamma_0^j > 0$. In this case, the optimal fiscal announcement with Full Commitment satisfies $\overline{\tau}_t^j = \mathbb{E}_t \tau_{t+1}^j$, i.e., the announcement coincides with the expected realization of the future tax rate, as in the two-period model of Section 2.

3.6 Optimal Time-Consistent Policy

We now consider a different assumption on the commitment technology. Specifically, we interpret the government sector as a succession of decision makers without commitment, one at each date t. Importantly, the government in power at t chooses current tax rates and makes announcements about future (one-period ahead) tax rates. Consistent with our assumptions in the previous subsection, these announcements are noncontingent with respect to future shocks and enter the cost of state contingency faced by the government in power at $t+1$. Thus, announcements provide an anchor for future tax rates, but do not amount to actual commitments, because the future government may choose state-contingent taxes subject to paying the costs of state contingency.

This institutional framework generalizes the Limited-Time Commitment model of Clymo and Lanteri (2020), in which future announcements are instead commitments and must coincide with ex-post realized policy. Here, we allow governments to endogenously choose the degree to which they desire to stick to their predecessors' announcements. In so doing, we develop a natural model to analyze the trade-off between partial commitment and partial state-contingency in optimal fiscal policy.

The state of the economy at date t is given by the physical state variables k_{t-1}, g_t, as well as the announced plan $\overline{\tau}^k_{t-1}, \overline{\tau}^l_{t-1}$ inherited by the previous government, which affects the costs of state contingency at t. We denote the state by $x_t \equiv (k_{t-1}, g_t, \overline{\tau}^k_{t-1}, \overline{\tau}^l_{t-1})$. We focus on a symmetric equilibrium. Building on the literature on Markov-perfect fiscal policy (e.g., Klein, Krusell, and Ríos-Rull, 2008), we restrict policies and allocations to be differentiable functions of a vector of "natural" state variables, and exploit differentiability to derive and interpret Generalized Euler Equations that characterize optimal policy.

Let all future governments set their policy according to functions $\tau^k = \tilde{\tau}^k(x)$, $\tau^l = \tilde{\tau}^l(x)$ and denote the associated allocations by $c = \tilde{c}(x)$, $l = \tilde{l}(x)$, $k' = \tilde{k}(x)$, where k' refers to capital productive in the following period. We highlight an important distinction between the functions $\tilde{c}, \tilde{l}$ and the functions h^c, h^l introduced above. Critically, the argument of $\tilde{c}$ and $\tilde{l}$ includes previously *announced* tax rates for the current period, which are part of the natural sate of the economy. In contrast, the argument of h^c and h^l includes currently *realized* tax rates. These functions are related as follows

$$\tilde{c}(x) = h^c(k, g, \tilde{\tau}^k(x), \tilde{\tau}^l(x)), \tag{46}$$

$$\tilde{l}(x) = h^l(k, g, \tilde{\tau}^k(x), \tilde{\tau}^l(x)). \tag{47}$$

Furthermore, let $\tilde{W}(x)$ be the present discounted value of government utility (33) associated with policy functions introduced above, when the state of the economy is x. Using this notation, we can state the optimization problem of a government as to choose allocations and taxes $(c, l, k', \tau^k, \tau^l)$ as well as announcements $(\overline{\tau}^{k,'}, \overline{\tau}^{l,'})$ to maximize

$$u(c) - v(l) - \Gamma^k(\tau^k, \overline{\tau}^k) - \Gamma^l(\tau^l, \overline{\tau}^l) + \beta \mathbb{E} \tilde{W}(x') \tag{48}$$

subject to the resource constraint

$$c + k' + g = F(k, l) + (1 - \delta)k \tag{49}$$

with associated multiplier λ, and the implementability constraints

$$u_c(c)k' = \beta \mathbb{E} \left[u_c(\tilde{c}(x')) \left(\tilde{c}(x') + \tilde{k}'(x') \right) - v_l(\tilde{l}(x'))\tilde{l}(x') \right] \tag{50}$$

with multipier μ, and

$$l = h^l(k, g, \tau^k, \tau^l) \tag{51}$$

$$c = h^c(k, g, \tau^k, \tau^l) \tag{52}$$

with multipliers ν^l and ν^c respectively. We also impose an upper bound on the capital tax $\tau^k \leq \tau^k_{max}$, with associated multiplier ξ, to ensure that the problem is well defined, even for small (or zero) costs of state contingency.[12]

The first-order conditions with respect to consumption, labor, and capital are

$$\begin{aligned} \lambda &= u_c(c) - \mu u_{cc}(c)k' - \nu^c && (53) \\ v_l(l) &= \lambda F_l(k, l) + \nu^l, && (54) \\ \lambda &= \beta \mathbb{E} \tilde{W}_k(x') - \mu u_c(c) + \mu \beta \mathbb{E} S_k(x') && (55) \end{aligned}$$

where we used shorthand notation $S(x') \equiv \left[u_c(\tilde{c}(x')) \left(\tilde{c}(x') + \tilde{k}'(x') \right) - v_l(\tilde{l}(x'))\tilde{l}(x') \right]$ to refer to the term in the square bracket of constraint (50), which relates the government primary surplus to the private-sector allocation. An important difference between these optimality conditions and their counterparts in the Full Commitment problem of the previous subsection is that *past* multipliers on the implementability constraint (50) are absent

[12] We will set this bound to a large enough value not to bind in the Full-Commitment problem.

here, because the government disregards the effects of current policy on past decisions of the private sector, and in particular past investment. Moreover, the derivatives of the future policy functions appear inside the term $\mathbb{E}S_k(x')$, rendering these optimality conditions Generalized Euler Equations.

The first-order conditions with respect to realized taxes are

$$\nu^c h^c_{\tau^k}(k,g,\tau^k,\tau^l) - \nu^l h^l_{\tau^k}(k,g,\tau^k,\tau^l) = \Gamma^k_{\tau^k}(\tau^k,\overline{\tau}^k) + \xi, \tag{56}$$

$$\nu^c h^c_{\tau^l}(k,g,\tau^k,\tau^l) - \nu^l h^l_{\tau^l}(k,g,\tau^k,\tau^l) = \Gamma^l_{\tau^l}(\tau^l,\overline{\tau}^l). \tag{57}$$

which coincide with their counterparts in the Full-Commitment problem, except for the multiplier on the upper bound on the capital tax.

The first-order conditions with respect to future tax announcements are

$$\mathbb{E}\tilde{W}_{\overline{\tau}^k}(x') + \mu\mathbb{E}S_{\overline{\tau}^k}(x') = 0, \tag{58}$$

$$\mathbb{E}\tilde{W}_{\overline{\tau}^l}(x') + \mu\mathbb{E}S_{\overline{\tau}^l}(x') = 0. \tag{59}$$

Furthermore, we have the following envelope conditions

$$\tilde{W}_k(x) = \lambda\left[F_k(k,l) + (1-\delta)k\right] - \nu^l h^l_k + \nu^c h^c_k, \tag{60}$$

$$\tilde{W}_{\overline{\tau}^k}(x) = -\Gamma^k_{\overline{\tau}^k}(\tau^k,\overline{\tau}^k), \tag{61}$$

$$\tilde{W}_{\overline{\tau}^l}(x) = -\Gamma^l_{\overline{\tau}^l}(\tau^l,\overline{\tau}^l). \tag{62}$$

By combining the envelope conditions (61) and (62) with equations (58) and (59), we obtain a key distinction with respect to the Full-Commitment problem: Optimal announcements do not just minimize the expected costs of state contingency. Instead, they trade off the incentive to reduce expected cost of state contingency with the possibility for the current government to manipulate the following government's problem by setting its inherited fiscal announcements. This strategic incentive is reflected in the presence of the terms $\mathbb{E}S_{\overline{\tau}^k}(x')$ and $\mathbb{E}S_{\overline{\tau}^l}(x')$ in these Generalized Euler Equations for the optimal announcements.

4 Numerical Results

In this section, we calibrate our infinite-horizon model and discuss our numerical results on optimal policy with Full-Commitment and under the Time-Consistent policy.

4.1 Calibration and Solution Method

We parameterize the utility function as follows: $u(c) \equiv \frac{c^{1-\eta_c}}{1-\eta_c}$ and $v(l) \equiv \chi \frac{l^{1+\eta_l}}{1+\eta_l}$, with $\eta_c = 1$, i.e., log utility from consumption, $\eta_l = 2$ to match the Frisch elasticity of labor supply, and $\chi = 0.7804$ to normalize average labor to one in the Full-Commitment model. The production function is Cobb-Douglas, with capital share α: $F(k,l) \equiv Zk^{\alpha}l^{1-\alpha}$, with $\alpha = 1/3$, and $Z = 0.338$ to normalize average capital to one.

We calibrate the Markov process for g_t as an AR(1) in logs, and match the average ratio of government spending to GDP (approximately 20%), as well as the standard deviation and autocorrelation of government spending. We then discretize this process with a two-valued Markov chain.

We parameterize the costs of state contingency as follows: $\Gamma^j(\tau^j, \overline{\tau}^j) \equiv \frac{\gamma_0^j}{2}(\tau^j - \overline{\tau}^j)^2$ with $\gamma_0^j > 0$ for $j = k, l$. We then restrict $\gamma_0^k = \gamma_0^l = \gamma_0$ to highlight the importance of treating the two instruments symmetrically. We calibrate this parameter to closely match the standard deviation of the capital tax rate (relative to a linear trend), which is approximately equal to 2%, using US data for the period 1970-2014 from Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015). This gives $\gamma_0 = 100$. Finally, we set the upper bound on the capital tax $\tau^k_{max} = 0.3$, ensuring it is not binding under Full Commitment.

For the Full Commitment problem, we solve the model using a generalization of the Parameterized Expectations Algorithm (den Haan and Marcet, 1990) proposed by Valaitis and Villa (2019). This method relies on a neural network (instead of standard regression) to approximate the key forward looking terms in the optimality conditions as functions of the state vector. For the Time-Consistent problem, we further adapt these methods to solve the Generalized Euler Equation directly. We provide more details on the solution method in Appendix B.

4.2 Full Commitment

We now present our numerical results and compare policies and allocations in our baseline model with costs of state contingency ($\gamma_0^k = \gamma_0^l = 100$) with two alternative comparison models. The first comparison model fully removes costly state contingency, and sets $\gamma_0^k = \gamma_0^l = 0$. We highlight that this model coincides with the model analyzed by Stockman (2001). The second comparison model sets $\gamma_0^k = \infty, \gamma_0^l = 0$. In this model, the capital tax must be chosen one period in advance, whereas the labor tax is freely adjustable within the period. This assumption is also common in the literature, and is adopted by, for example,

Table 1: Parameter Values

		Parameter	Value
Preferences	Discount factor	β	0.96
	Risk aversion	η_c	1
	Labor disutility	χ	0.7804
	Labor elasticity	η_l	2
Technology	Capital share	α	0.36
	Depreciation	δ	0.08
Government	Average g	μ_g	0.0682
	Volatility of $\log(g)$	σ_g	0.0161
	Autocorr. of $\log(g)$	ρ_g	0.9774
	Upper bound on τ^k	τ^k_{max}	0.3
	Cost of state cont. τ^k	γ^k_0	100
	Cost of state cont. τ^l	γ^l_0	100

Fahri (2010).

In Figure 1, we illustrate the response of taxes on capital and labor income to an exogenous increase in government spending in these three models. The figures plot the response of the economy to a switch from the low to high government spending states, after a long period in the low state.

In our baseline model with costly state contingency (solid line), the shock induces the government to increase capital taxes moderately and temporarily, and labor taxes in a more persistent way. In both comparison models, instead, the capital tax increases significantly, either contemporaneously (dashed line) or with a one period lag when they are predetermined (dashed-dotted line), and accounts for the bulk of the endogenous policy response to the shock.

We highlight that the government in our model could easily choose to increase capital taxes with a lag, because our costs of state contingency are relevant when the government is surprised by the shock, but do not prevent a lagged adjustment. Furthermore, our estimated government spending shock is very persistent, so it might appear optimal for the government to adjust policies towards the optimal level in the model without contingency costs, even with a lag. However, the government instead optimally chooses to promise a *lower* future capital tax. This is one of our main findings: when it is hard to set taxes in a

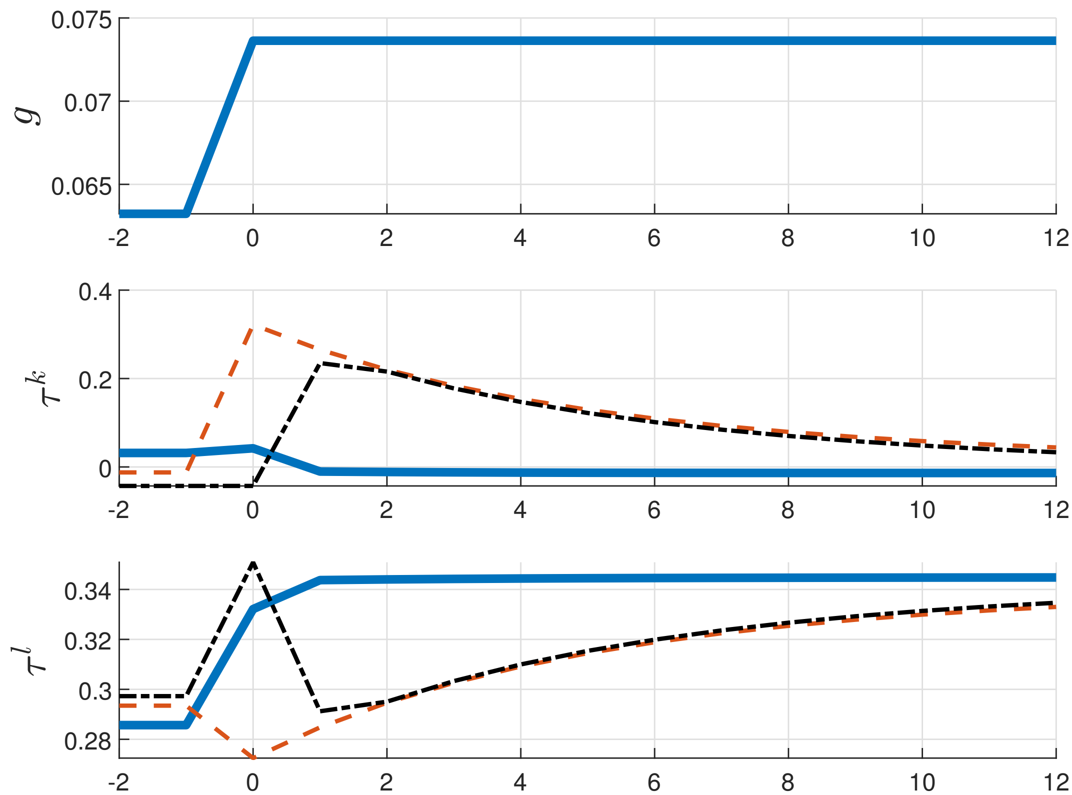

Figure 1: The figure displays the dynamics of fiscal variables around a shock that increases government spending. Top: government spending; middle: capital income tax rate; bottom: labor income tax rate. X-axes report time. Solid line: baseline model with costly state contingency ($\gamma_0^k = \gamma_0^l = 100$); dashed line: no costs of state contingency ($\gamma_0^k = \gamma_0^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma_0^k = \infty, \gamma_0^l = 0$).

state-contingent way, this does not just induce otherwise optimal policies to be implemented with a lag, but also fundamentally changes the policies that governments optimally choose to set.

To better understand this result, we now turn our attention to private sector allocations. In Figure 2, we show the dynamics of labor, consumption, capital, and output. We find that labor and capital are less responsive to the shock in our model than in either comparison model. Furthermore, our model produces a substantial consumption drop in response to the shock, whereas consumption is relatively smoother in the two comparison models.

Recall that in order to balance the budget for a given choice of current tax rates, the government must ensure a particular level of labor, given by equation (37). In turn, this level of labor dictates a level of consumption through equation (38), and a level of future

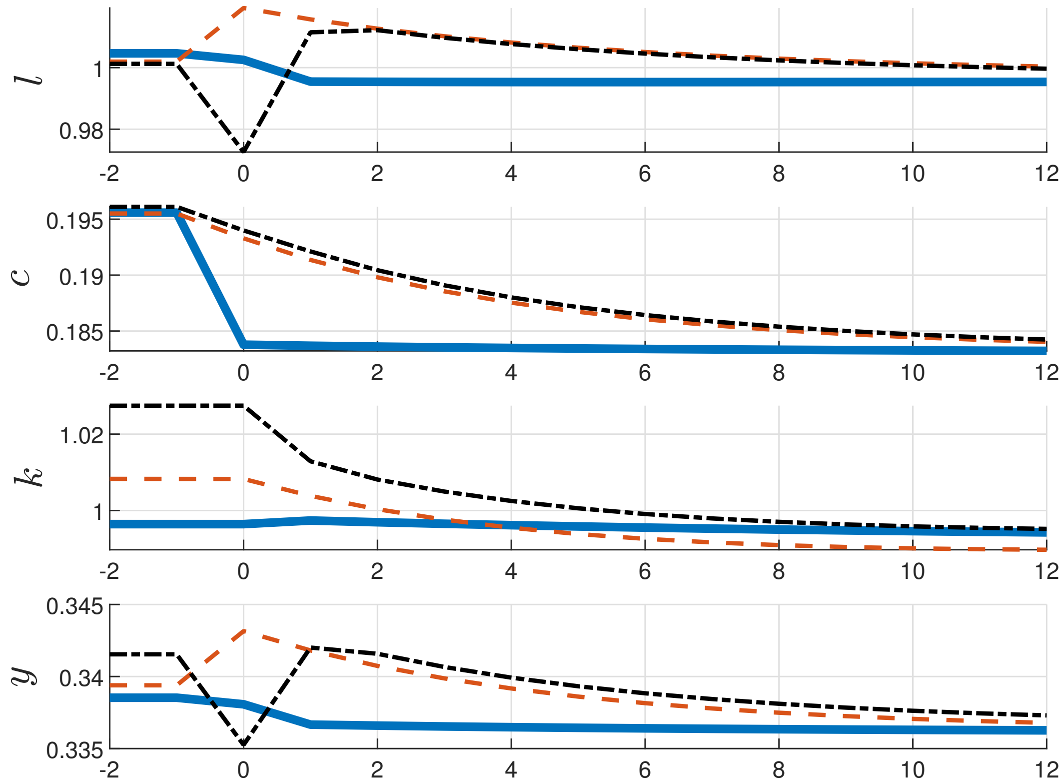

Figure 2: The figure displays the dynamics of allocations around a shock that increases government spending. From top to bottom: labor; consumption; capital; output. X-axes report time. Solid line: baseline model with costly state contingency ($\gamma_0^k = \gamma_0^l = 100$); dashed line: no costs of state contingency ($\gamma_0^k = \gamma_0^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma_0^k = \infty, \gamma_0^l = 0$).

capital through the resource constraint (26). In order to induce households to exert this required level of effort, when it is costly to adjust current taxes, the government engineers a drop in current consumption, thus leveraging the effect of the marginal utility from consumption on labor supply, and thus must induce an associated increase in investment to satisfy the resource constraint. Implementing this allocation requires low future capital taxes, driving a key difference with respect to the comparison models, in which capital taxes increase persistently. Instead of relying on high capital taxes to generate additional revenue, the government in our baseline model increases labor taxes persistently.

Overall, by contrasting the two comparison models, we find that the ability to adjust capital taxes contemporaneously or with a lag does not appear to make a large difference in terms of the government's ability to insure household consumption from government spending shocks. Indeed allocations are similar in the two comparison models, except for labor in the initial period of the shock.[13] In contrast, costs of state contingency on both capital and labor taxes change both the optimal policy response and the associated private-sector allocation. In particular, when changing current policy is costly, the government actively uses future policies to ensure that the current budget constraint is satisfied. The adjustment involves persistently higher labor taxes, and a and significant drop in private consumption.

In Figure 3 we focus on optimal announcements in our model (dashed), and contrast them to ex-post realized policies (solid). Because government spending is highly persistent, realized tax rates closely match promised tax rates, except in the periods in which the value of government spending increases, when the government deviates from the promise to generate additional revenue.

4.3 Time-Consistent Policy

We now discuss our numerical results for the case in which the government lacks commitment to state-contingent policies, but formulates noncontingent announcements. In Figures 4 and 5 we display fiscal variables and private-sector allocations respectively. As in the previous subsection, we focus on a shock that increases the level of government spending. We compare the Time-Consistent outcome (solid lines) with the Full-Commitment results (dashed lines).

The first difference we notice is that the capital tax is on average higher under the Time-Consistent policy. However, in absolute terms, the level of the capital tax is still small, approximately 5%. To put this result in context, we also solve the model under the Time-Consistent policy without costs of state contingency ($\gamma_0^k = \gamma_0^l = 0$), and we find that the capital tax is always equal to its upper bound, which we set equal to 30%. Thus, we find that costs of state contingency play a powerful role in building a chain of commitment across successive governments, thereby reducing each government's incentive to raise capital tax as if it was lump sum.

We also find the the dynamics of the labor tax are quite similar across the Time-Consistent and Full Commitment policy. In both cases, we find that an increase in government spending leads to a persistent increase in the labor tax, compared to a muted and short-lived response of the capital tax.

[13]In Appendix C we corroborate this finding by considering the case in which our calibrated degree of costs of state contingency applies only to capital taxes.

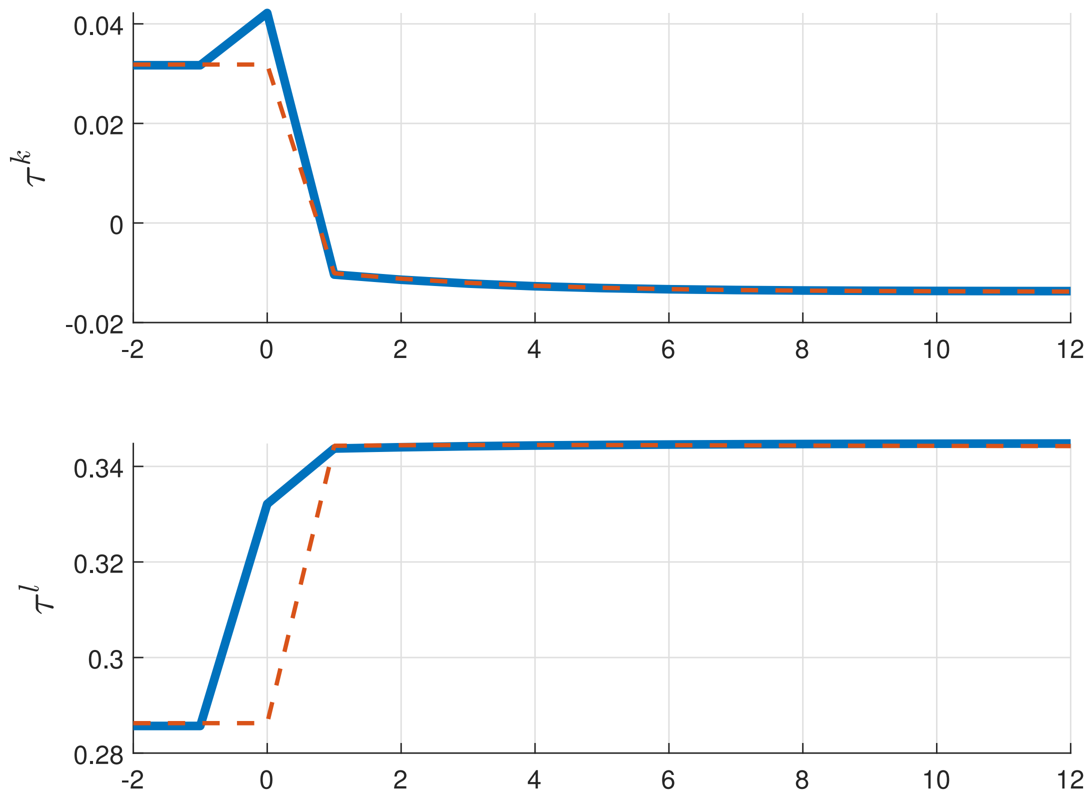

Figure 3: The figure compares the dynamics of realized tax rates (solid line) and announced tax rates (dashed line). Top: capital income tax rate; bottom: labor income tax rate. X-axes report time.

The allocations implied by the Time-Consistent policy are also similar to the ones that arise under Full Commitment, with the noticeable exception of investment, which drops more significantly under the Time-Consistent policy, because the capital tax does not decrease as substantially under Full Commitment in the periods following the shock.

In Figure 6, we compare realized taxes with noncontingent announcements in the Time-Consistent solution. As in the Full-Commitment case, the period in which the shock hits coincides with a large deviation between realized taxes and previous announcements. Different with the Full Commitment case, however, we also find noticeable deviations in periods in which the level of government spending stays constant. These differences arise because of the bias terms in equations (58) and (59): In formulating announcements, the government trades off a forecast of future taxes with an incentive to strategically manipulate future policies.

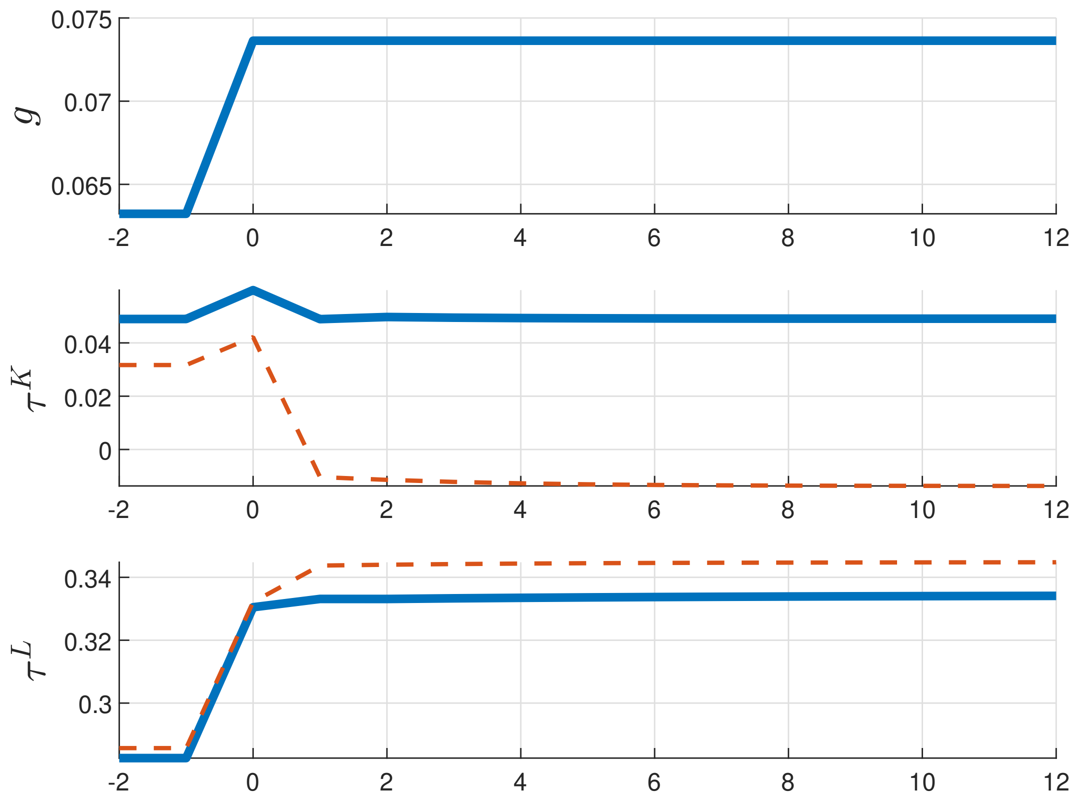

Figure 4: The figure displays the dynamics of fiscal variables around a shock that increases government spending. Top: government spending; middle: capital income tax rate; bottom: labor income tax rate. X-axes report time. Solid line: Time-Consistent policy; dashed line: Full Commitment.

5 Conclusion

In this paper, we have explored the role of costs of state contingency in fiscal plans for optimal policy and for the response of the economy to government spending shocks. In our framework, the government makes noncontingent announcements about future taxes. After shocks are realized, the government may deviate from these announcements, subject to a cost.

Under Full Commitment, these costs of state contingency reduce the volatility of taxes in response to shocks, but increase the volatility of private consumption. Whereas previous models of optimal fiscal policy imply that volatility in capital taxes should play a prominent role in absorbing fiscal shocks, we find an important role for persistent changes in labor taxes, consistent with empirical evidence. Furthermore, our model implies that future

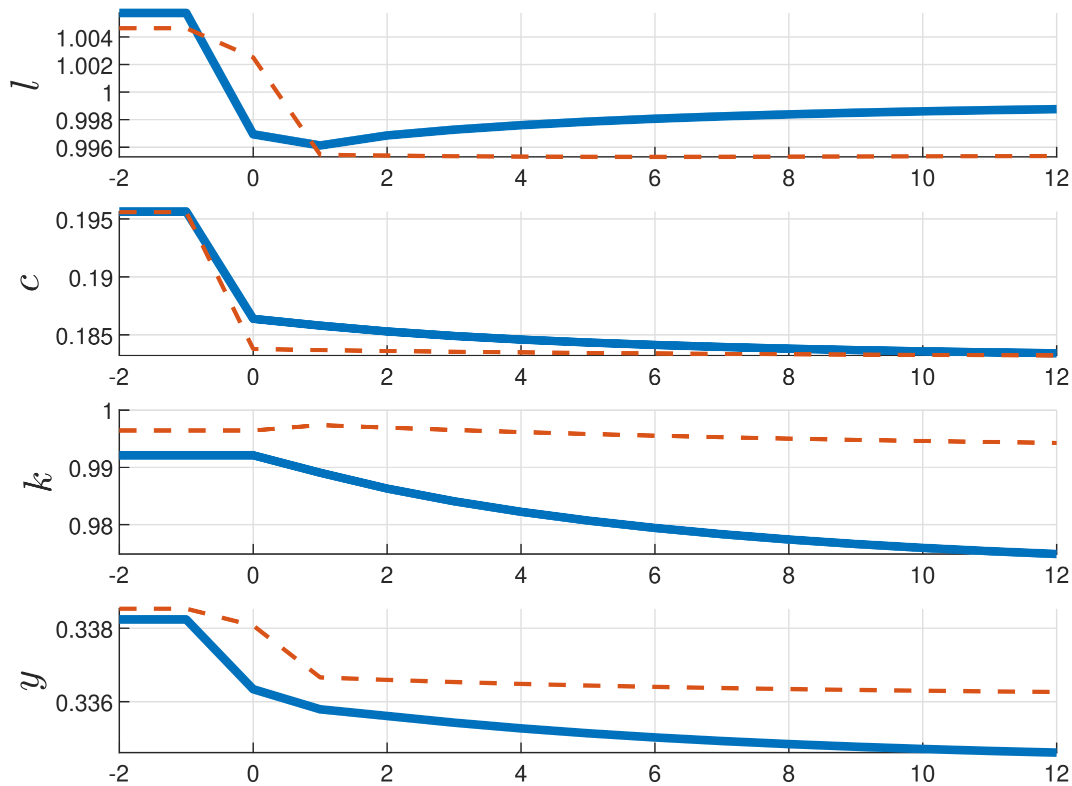

Figure 5: The figure displays the dynamics of allocations around a shock that increases government spending. From top to bottom: labor; consumption; capital; output. X-axes report time. Solid line: Time-Consistent policy; dashed line: Full Commitment.

announcements are used to give incentives to the private sector in order to generate a desired level of current tax base, when changing current taxes is costly. Crucially, we therefore find that the inability to set policy in a state contingent way does not simply induce governments to implement otherwise-optimal policies with a lag, but fundamentally changes the policies which are feasible, and hence optimal.

When the government lacks commitment, fiscal announcements play a strategic role and allow the government to affect future policies, by partially constraining its future decisions. As a consequence, we find that optimal announcements are biased forecasts of future policies, but may induce higher welfare than in the absence of costly state contingency.

We have explored the implications of our model under the assumption of a government balanced-budget constraint. In future work, we plan to combine these insights with a more general model in which fiscal shocks can be partially accommodated with fiscal deficits.

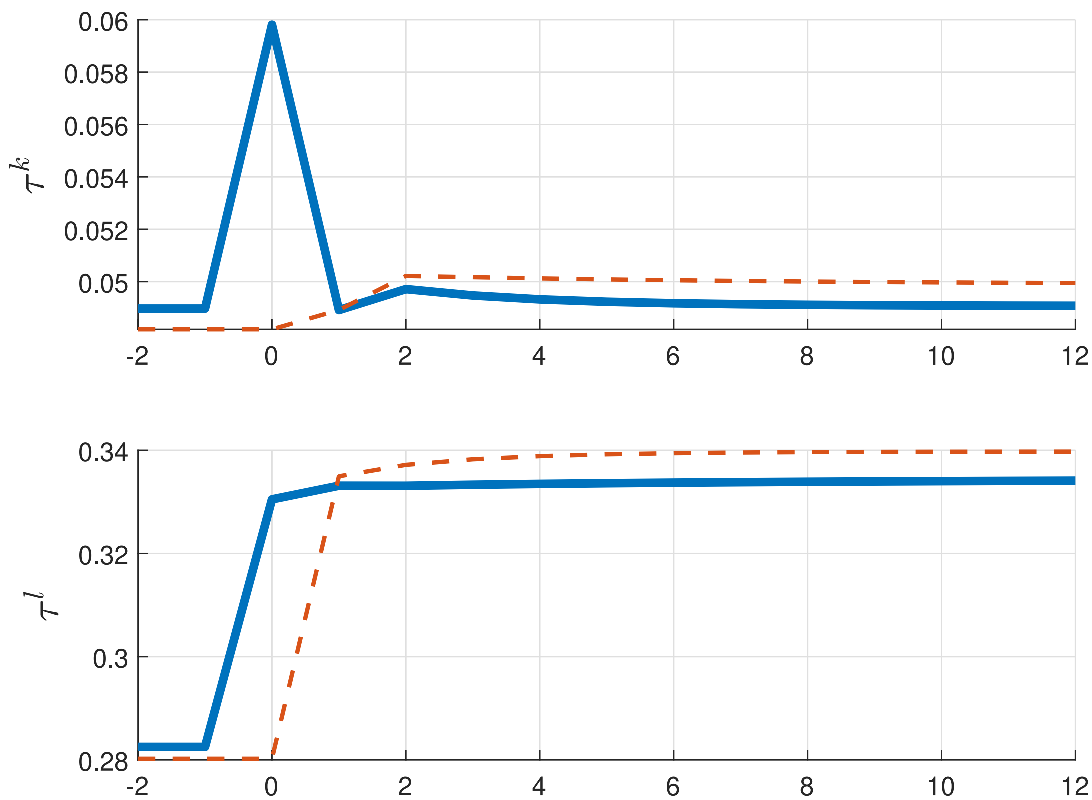

Figure 6: The figure compares the dynamics of realized tax rates (solid line) and announced tax rates (dashed line). Top: capital income tax rate; bottom: labor income tax rate. X-axes report time.

References

AIYAGARI, S., A. MARCET, T. SARGENT, AND J. SEPPALA (2002): "Optimal Taxation without State-Contingent Debt," *Journal of Political Economy*, 110(6), 1220–1254.

ALESINA, A., C. FAVERO, AND F. GIAVAZZI (2015): "The Output Effect of Fiscal Consolidations," *Journal of International Economics*, 96(S1), 19–42.

ATHEY, S., A. ATKESON, AND P. J. KEHOE (2005): "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 73(5), 1431–1475.

BURNSIDE, C., M. EICHENBAUM, AND J. D. FISHER (2004): "Fiscal shocks and their consequences," *Journal of Economic Theory*, 115, 89–117.

CHAMLEY, C. (1986): "Optimal taxation of capital income in general equilibrium with infinite lives," *Econometrica*, 54(3), 607–622.

CHARI, V. V., AND P. J. KEHOE (1999): "Optimal fiscal and monetary policy," in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford, vol. 1(C), pp. 1671–1745. Elsevier.

CLYMO, A., AND A. LANTERI (2020): "Fiscal Policy with Limited-Time Commitment," *Economic Journal*, 130(627), 623–652.

DEBORTOLI, D., AND R. NUNES (2010): "Fiscal policy under loose commitment," *Journal of Economic Theory*, 145(3), 1005–1032.

——— (2013): "Lack of Commitment and the Level of Debt," *Journal of the European Economic Association*, 11(5), 1053–1078.

DEN HAAN, W., AND A. MARCET (1990): "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics*, 8, 34–38.

FAHRI, E. (2010): "Capital Taxation and Ownership When Markets are Incomplete," *Journal of Political Economy*, 118(5), 908–948.

FARAGLIA, E., A. MARCET, R. OIKONOMOU, AND A. SCOTT (2019): "Government Debt Management: The Long and the Short of It," *Review of Economic Studies*, 86(6), 2554–2604.

FERNÁNDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, K. KUESTER, AND J. RUBIO-RAMÍREZ (2015): "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, 105(11), 3352–3384.

HALAC, M., AND P. YARED (2014): "Fiscal Rules and Discretion Under Persistent

Shocks," *Econometrica*, 82(5), 1557–1614.

Hauk, E., A. Lanteri, and A. Marcet (2020): "Optimal Policy with General Signal Extraction," Discussion paper.

Judd, K. L. (1985): "Redistributive taxation in a simple perfect-foresight model," *Journal of Public Economics*, 28(1), 59–83.

Karantounias, A. G. (2019): "Greed versus Fear: Optimal Time-Consistent Taxation with Default," Discussion paper.

King, R. G., C. I. Plosser, and S. T. Rebelo (1988): "Production, Growth, and Business Cycles, II. New Directions," *Journal of Monetary Economics*, 21, 309–341.

Klein, P., P. Krusell, and J. V. Ríos-Rull (2008): "Time-Consistent Public Policy," *Review of Economic Studies*, 75(3), 789–808.

Klein, P., and J. V. Ríos-Rull (2003): "Time-Consistent Optimal Fiscal Policy," *International Economic Review*, 44(4), 1217–1245.

Martin, F. (2010): "Markov-perfect capital and labor taxes," *Journal of Economics Dynamics and Control*, 34, 503–521.

Mertens, K., and M. Ravn (2012): "Empirical evidence on the aggregate effects of anticipated and unanticipated U.S. tax policy shocks," *American Economic Journal: Economic Policy*, 4(2), 145–181.

Schmitt-Grohe, S., and M. Uribe (1997): "Balanced-Budget Rules, Distortionary Taxes, and Aggregate Instability," *Journal of Political Economy*, 105(5), 976–1000.

Stockman, D. R. (2001): "Balanced-Budget Rules: Welfare Loss and Optimal Policies," *Review of Economic Dynamics*, 4, 438–459.

Valaitis, V., and A. T. Villa (2019): "Machine Learning Projection Methods for Macro-Finance Models," Discussion paper.

APPENDICES

A Appendix to Two-Period Model

A.1 General Framework

In this appendix we set up a general model for our formal results on welfare. Our approach is to formulate an indirect utility function, which expresses welfare in terms of policy instruments and endogenous state variables, subject to generic competitive-equilibrium constraints. The model in the main text fits into this framework.

There are two periods, $t = 0, 1$. There is a shock at time 1, denoted g. The government chooses policies τ at time 1, which may be allowed to depend on the shock depending on the game. We focus on a single continuous policy choice and a single state variable for simplicity. Because we have assumed in Section 2.4 that the cost of state contingency applies to only one tax instrument, the other tax instrument drops out from the objective function, and can be computed as a residual from the government budget constraint.

Private agents make decisions correctly anticipating government policy. Substitute out all of these decisions into indirect utility, apart from an initial investment choice, k. Let k denote an endogenous state variable chosen by the household. This can be thought of as capital, chosen optimally given expected taxes. The variable is dynamic, and creating k units in period 0 costs $-k$ utility, but enters the utility function in period 1.

The game can be specified as follows. Let $u(k, \tau, g)$ denote the indirect utility function at time 1. Maintain that $u_\tau(\tau) > 0$ for some τ and $u_{\tau\tau} < 0$ so that the optimization has a unique interior solution in the Time-Consistent case. The welfare of the household is

$$-k + \beta \mathbb{E}_0 u(k, \tau(g), g) \tag{A1}$$

for some tax policy function $\tau(g)$. Household maximization leads to a choice of state given implicitly by the constraint

$$\beta \mathbb{E}_0 f(k, \tau(g), g) \geq 0 \tag{A2}$$

where $f(k, \tau, g)$ is a function evaluated at every realized g. Notice the classic problem of future policy choices in today's constraint.

Assumption 1 *Either $f_\tau(k, \tau, g) \geq 0$ for all values, or $f_\tau(k, \tau, g) \leq 0$ for all values.*

The intuitive meaning of this assumption is that increasing the policy τ always has consistently the same effect on either tightening or loosening the time-0 constraint.

Assume costly state contingency with parameter γ. Government welfare is

$$-k + \beta \mathbb{E}_0 \left[u(k, \tau(g), g) - \frac{\gamma}{2} (\tau(g) - \overline{\tau})^2 \right] \tag{A3}$$

A.2 Full Commitment with Costly State Contingency

The government value is given by

$$V^{FC} = \max_{k, \tau(g), \overline{\tau}} -k + \beta \mathbb{E}_0 \left[u(k, \tau(g), g) - \frac{\gamma}{2} (\tau(g) - \overline{\tau})^2 \right] \tag{A4}$$

s.t. $\beta \mathbb{E}_0 f(k, \tau(g), g) \geq 0$. Let $\mu \geq 0$ be the multiplier on the constraint. For $\gamma > 0$ the FOCs give

$$\frac{\partial}{\partial k} \implies \beta \mathbb{E}_0 \left[u_k(k, \tau(g), g) + \mu f_k(k, \tau(g), g) \right] = 1 \tag{A5}$$

$$\frac{\partial}{\partial \tau(g)} \implies \tau(g) = \overline{\tau} + \frac{1}{\gamma} \left(u_\tau(k, \tau(g), g) + \mu f_\tau(k, \tau(g), g) \right) \tag{A6}$$

$$\frac{\partial}{\partial \overline{\tau}} \implies \overline{\tau} = \mathbb{E}_0 \left[\tau(g) \right] \tag{A7}$$

When $\gamma = 0$, we have the standard Full-Commitment solution with $u_\tau(k, \tau(g), g) + \mu f_\tau(k, \tau(g), g) = 0$ and drop the equation for $\overline{\tau}$. When $\gamma > 0$, the policies are chosen to move towards the optimal choice but the higher is γ the closer they remain to $\overline{\tau}$. Finally $\overline{\tau}$ is chosen to be the average expected policy. The solution is continuous in the limit as $\gamma \to 0$, and converges to the standard solution. This is true for all variables except the promise $\overline{\tau}$, which is indeterminate at $\gamma = 0$. In the limit of $\gamma \to 0$, the promise is well defined and given by $\overline{\tau} = \mathbb{E}_0 \left[\tau(g) \right]$ where $\tau(g)$ is the policy function in the standard solution.

Taking expectations of the $\tau(g)$ FOC and plugging in $\overline{\tau} = \mathbb{E}_0 \left[\tau(g) \right]$ gives

$$\mathbb{E}_0 \left[u_\tau(k, \tau(g), g) + \mu f_\tau(k, \tau(g), g) \right] = 0 \tag{A8}$$

In the standard problem when $\gamma = 0$, the government sets the future policy optimally state by state to set $u_\tau(k, \tau(g), g) + \mu f_\tau(k, \tau(g), g) = 0$. This trades off the utility of the choice in period 1 with its effect on the constraint. When there is costly state contingency, this is no longer true state by state, but the above equation shows that the government still sets the policies to trade off future utility and the effect on the constraint on average.

A.3 Welfare under Full Commitment

We now prove Proposition 1 in this general model.

Proof. First take envelope for overall value:

$$\frac{\partial V^{FC}}{\partial \gamma} = -\frac{\beta}{2}\mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right] < 0 \tag{A9}$$

So overall value is decreasing in γ. Could welfare excluding the cost be increasing? Define value excluding cost as

$$\tilde{V}^{FC} = V^{FC} + \frac{\beta\gamma}{2}\mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right] \tag{A10}$$

We show a proof by contradiction that this must be decreasing too. Suppose γ rises to γ'. Consider the old optimal policy $\tau(g)$ and new one $\hat{\tau}(g)$. We know from the Envelope Theorem that overall value decreases: $V^{FC}(\hat{\tau}) < V^{FC}(\tau)$. Suppose value excluding cost increases: $\tilde{V}^{FC}(\hat{\tau}) > \tilde{V}^{FC}(\tau)$. Then we know that cost component must have decreased: $\frac{\beta\gamma'}{2}\mathbb{E}_0\left[\left(\hat{\tau}(g) - \hat{\overline{\tau}}\right)^2\right] < \frac{\beta\gamma}{2}\mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right]$.

Given that γ has increased, it must be that $\mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right]$ is strictly lower at the new policy: $\mathbb{E}_0\left[\left(\hat{\tau}(g) - \hat{\overline{\tau}}\right)^2\right] < \mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right]$. This implies a contradiction: the new policy must deliver higher value than the old policy at the original parameter values. The new policy at the old γ gives total value

$$\tilde{V}^{FC}(\hat{\tau}) - \frac{\beta\gamma}{2}\mathbb{E}_0\left[\left(\hat{\tau}(g) - \hat{\overline{\tau}}\right)^2\right] \tag{A11}$$

And the old policy at the original γ gives

$$\tilde{V}^{FC}(\tau) - \frac{\beta\gamma}{2}\mathbb{E}_0\left[(\tau(g) - \overline{\tau})^2\right] \tag{A12}$$

According to the inequalities above, the new policy gives higher value, contradicting that the original policy was optimal. Intuitively, we know that value including the cost is decreasing in γ. The value excluding the cost must therefore also fall, otherwise the quadratic cost itself must have fallen when the cost parameter got larger. ■

A.4 Time-Consistent Policy with Costly State Contingency

Time-1 problem: The second period problem is static, and solved ex-post for any realized shock:

$$V^1(k,\overline{\tau},g) = \max_{\tau} u(k,\tau,g) - \frac{\gamma}{2}(\tau-\overline{\tau})^2 \tag{A13}$$

The solution is a policy function $\tilde{\tau}(k,\overline{\tau},g)$ which satisfies the FOC

$$\frac{\partial}{\partial\tau} \implies \tilde{\tau}(k,\overline{\tau},g) = \overline{\tau} + \frac{1}{\gamma}u_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g) \tag{A14}$$

for $\gamma > 0$ and $u_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g) = 0$ for $\gamma = 0$. Notice the choice is different than under Full Commitment, since the government does not consider effect on the time-0 constraint. The policy is always increasing in the announcement:

$$\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g) = \frac{1}{1-\frac{1}{\gamma}u_{\tau\tau}(.)} \tag{A15}$$

Since $u_{\tau\tau}(.) < 0$ by assumption, the denominator is positive, so raising the promise raises the policy at every value of the shock. The partial derivative with respect to γ will also be useful (dependence on γ is suppressed in the notation unless needed)

$$\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma) = \frac{\overline{\tau}-\tilde{\tau}(k,\overline{\tau},g;\gamma)}{\gamma - u_{\tau\tau}(.)} \tag{A16}$$

Since $u_{\tau\tau}(.) < 0$ by assumption, the denominator is positive. This means that the effect of the announcement on realized policy is intuitive: If the current choice $\tilde{\tau}$ is below the promise $\overline{\tau}$ then making the deviation more costly by raising γ will raise $\tilde{\tau}$.

The following envelopes will also be useful:

$$V^1_k(k,\overline{\tau},g) = u_k(k,\tau,g) \tag{A17}$$

$$V^1_{\overline{\tau}}(k,\overline{\tau},g) = \gamma(\tilde{\tau}(k,\overline{\tau},g)-\overline{\tau}) \tag{A18}$$

$$V^1_\gamma(k,\overline{\tau},g;\gamma) = -\frac{1}{2}(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau})^2 \tag{A19}$$

Time-0 problem: The first period government problem is

$$V^{TC} = \max_{k,\overline{\tau}} -k + \beta \mathbb{E}_0 V^1(k, \overline{\tau}, g) \tag{A20}$$

s.t. $\beta \mathbb{E}_0 f(k, \tilde{\tau}(k, \overline{\tau}, g), g) \geq 0$. Let $\mu \geq 0$ be the multiplier on the constraint. For $\gamma > 0$ the FOCs give

$$\frac{\partial}{\partial k} \implies \beta \mathbb{E}_0 \left[u_k(k, \tilde{\tau}, g) + \mu \left(f_k(k, \tilde{\tau}, g) + \tilde{\tau}_k f_\tau(k, \tilde{\tau}, g) \right) \right] = 1 \tag{A21}$$

$$\frac{\partial}{\partial \overline{\tau}} \implies \overline{\tau} = \mathbb{E}_0 \left[\tilde{\tau}(k, \overline{\tau}, g) \right] + \frac{\mu}{\gamma} \mathbb{E}_0 \tilde{\tau}_{\overline{\tau}}(k, \overline{\tau}, g) f_\tau(k, \tilde{\tau}(k, \overline{\tau}, g), g) \tag{A22}$$

When $\gamma = 0$, we have the standard Time-Consistent solution and drop the equation for $\overline{\tau}$. The Time-Consistent policy with costly state contingency has two differences with respect to its Full-Commitment counterpart. Firstly, the choice of endogenous state is chosen differently since now it is used to influence future policy. Secondly, the first period government uses the promise to try and influence the second period government. Instead of setting the promise equal to the average realized policy, the first period government biases the promise to influence the policies. Since $\mu > 0$ and $\tilde{\tau}_{\overline{\tau}} > 0$, the sign of the bias depends only on f_τ, i.e. how the policies influence the constraint on choosing the endogenous state. If increasing the policies relaxes the constraint ($f_\tau > 0$) then the first period government would like to raise the policies chosen ex-post by the second period government and will set the promise above the average value. Vice versa if $f_\tau < 0$. Recall we assume that it is always only one or the other at all values.

The solution is continuous in the limit as $\gamma \to 0$, and converges to the standard Time-Consistent solution. This is true for all variables except the promise $\overline{\tau}$, which is indeterminate at $\gamma = 0$. Interestingly, the promise $\overline{\tau}$ does have a well defined finite limit at $\gamma = 0$, since $\overline{\tau}_{\tilde{\tau}}/\gamma$ has a well defined limit. The implication of this is that for very small values of γ, just an epsilon above zero, it is possible for the time-0 government to influence the time-1 government despite the small γ: It just has to set an extremely large (in absolute terms) promise. However, this is too costly, since it would imply a large quadratic penalty, despite the vanishing γ, because the required size of the promise grows faster than γ shrinks.

We can also get a characterization of the trade-off of raising γ. Plugging the τ FOC from the second period into the $\overline{\tau}$ FOC to remove $\mathbb{E}_0\tau$ gives

$$\mathbb{E}_0\left[u_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)+\mu\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)\right]=0 \tag{A23}$$

Notice that this is very similar to the condition under Full Commitment, (A8), except with the $\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)$ term added. We have a formula for $\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)$ which can be plugged in if desired. When $\gamma=0$, $\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)=0$ and this is just the Time-Consistent policy where τ is chosen ex post. When $\gamma\to\infty$, then $\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)\to 1$ and we recover the Full-Commitment FOC with infinite costs. Thus we have our trade-off: Increasing γ moves the policy to the correct level on average, but at the cost of reducing state contingency.

A.5 Welfare under the Time-Consistent Policy

We want to see if raising γ could raise overall welfare, in contrast to the Full-Commitment case. We know that from time-1, raising γ must lower welfare, holding k and $\overline{\tau}$ fixed. Could raising γ improve ex-ante welfare by helping overcome time inconsistency?

We now prove Proposition 2.

Proof. The envelope theorem gives the change in time-0 welfare as we change γ as

$$V^{TC\prime}(\gamma)=\beta\mathbb{E}_0\left[V^1_\gamma(k,\overline{\tau},g;\gamma)+\mu\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g;\gamma),g)\right] \tag{A24}$$

Plug in envelope for V^1_γ:

$$V^{TC\prime}(\gamma)=\beta\mathbb{E}_0\left[-\frac{1}{2}(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau})^2+\mu\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g;\gamma),g)\right] \tag{A25}$$

Value excluding cost is

$$\tilde{V}^{TC}=V^{TC}+\beta\mathbb{E}_0\left[\frac{\gamma}{2}(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau}(\gamma))^2\right] \tag{A26}$$

Derivative with respect to γ is

$$\tilde{V}^{TC\prime}(\gamma)=V^{TC\prime}(\gamma)+\frac{\beta}{2}\mathbb{E}_0\left[(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau}(\gamma))^2\right]+\beta\gamma\mathbb{E}_0\left[(\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)-\overline{\tau}'(\gamma))(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau}(\gamma))\right] \tag{A27}$$

Plugging in $V^{TC\prime}(\gamma)$ gives

$$\tilde{V}^{TC\prime}(\gamma) = \beta\mathbb{E}_0\left[\mu\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g;\gamma),g) + \gamma(\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)-\overline{\tau}'(\gamma))(\tilde{\tau}(k,\overline{\tau},g;\gamma)-\overline{\tau}(\gamma))\right] \tag{A28}$$

To see that value might be increasing in γ at $\gamma = 0$, set $\gamma = 0$ in the above to yield

$$\tilde{V}^{TC\prime}(0) = \beta\mathbb{E}_0\left[\mu\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g;\gamma),g)\right] \tag{A29}$$

where it is understood that $\gamma = 0$, but we leave some γ in the formula for notational clarity. Recall that we know the sign of f_τ: it is either always non-negative or always non-positive by assumption. We have $\mu \geq 0$ since it's a KT multiplier. Next look at $\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma)$ when $\gamma = 0$. This is

$$\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma) = \frac{\overline{\tau}-\tilde{\tau}(k,\overline{\tau},g;\gamma)}{-u_{\tau\tau}(.)} \tag{A30}$$

From the $\overline{\tau}$ FOC we have $\overline{\tau} = \mathbb{E}_0\left[\tilde{\tau}(k,\overline{\tau},g)\right] + \frac{\mu}{\gamma}\mathbb{E}_0\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)$.

Define $x(g) \equiv \overline{\tau} - \tilde{\tau}(k,\overline{\tau},g;\gamma)$. Then we use the $\overline{\tau}$ FOC to give

$$\mathbb{E}_0 x(g) = \mathbb{E}_0\left[\overline{\tau}-\tilde{\tau}(k,\overline{\tau},g;\gamma)\right] = \frac{\mu}{\gamma}\mathbb{E}_0\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g) \tag{A31}$$

Then we can finally define the error of $x(g)$ realization from mean as $\epsilon(g) \equiv x(g) - \mathbb{E}_0 x(g)$, giving

$$x(g) = \frac{\mu}{\gamma}\mathbb{E}_0\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g) + \epsilon(g) \tag{A32}$$

Plug this is to the $\tilde{\tau}_\gamma$ formula to get

$$\tilde{\tau}_\gamma(k,\overline{\tau},g;\gamma) = \frac{\frac{\mu}{\gamma}\mathbb{E}_0\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)+\epsilon(g)}{-u_{\tau\tau}(.)} \tag{A33}$$

Finally plug this into our object of interest, the derivative of value wrt γ:

$$\tilde{V}^{TC\prime}(0) = \beta\mathbb{E}_0\left[\mu f_\tau(k,\tilde{\tau}(k,\overline{\tau},g;\gamma),g)\frac{\frac{\mu}{\gamma}\mathbb{E}_0\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)+\epsilon(g)}{-u_{\tau\tau}(.)}\right] \tag{A34}$$

Recall that $\mu > 0$, $\tilde{\tau}_{\overline{\tau}} > 0$, and $-u_{\tau\tau} > 0$. First check the limit of no uncertainty, which means that $\epsilon(g) = 0$ for all g. Then the formula reduces to

$$\tilde{V}^{TC\prime}(0) = \beta\frac{\frac{\mu^2}{\gamma}\tilde{\tau}_{\overline{\tau}}(k,\overline{\tau},g)f_\tau(k,\tilde{\tau}(k,\overline{\tau},g),g)^2}{-u_{\tau\tau}(.)} > 0 \tag{A35}$$

This is positive because the f_τ term is squared and $\tilde{\tau}_{\overline{\tau}} > 0$. Notice that this is well defined as $\gamma \to 0$ despite the γ in the denominator since $\tilde{\tau}_{\overline{\tau}}(k, \overline{\tau}, g) = \frac{1}{1-\frac{1}{\gamma}u_{\tau\tau}(.)}$ so the ratio $\tilde{\tau}_{\overline{\tau}}(k, \overline{\tau}, g)/\gamma = \frac{1}{\gamma - u_{\tau\tau}(.)}$ converges to $\frac{1}{-u_{\tau\tau}(.)}$

The derivative is strictly positive as long as the KT multiplier is binding. For the case of no uncertainty we are therefore done, and have proved that raising γ increases welfare in the Time-Consistent case when $\gamma = 0$. By continuity, this must also hold for some positive amount of uncertainty in the neighbourhood of zero uncertainty as well, since the rise in welfare at zero uncertainty is strictly positive. ■

B Solution Method

We solve the model using a generalization of the Parameterized Expectations Algorithm (den Haan and Marcet, 1990) proposed by Valaitis and Villa (2019).

We begin by describing the case of Full Commitment. The state variables of the government problem are $x^{FC} \equiv (k, g, \overline{\tau}^k, \overline{\tau}^l, \mu)$. We approximate the integrands in the expectation terms of the optimality conditions with a neural network with a single hidden layer with five neurons and hyperbolic tangent sigmoid transfer functions. We train the neural network to reproduce initial conditions that correspond to either the economy with no costs of state contingency, or the economy with predetermined capital taxes. We perform a long simulation of our economy ($T = 1000$), solving the optimality conditions for the current allocation and policies, given the approximated expectation terms. Next, we use our simulated sample for obtain a new iteration of our approximating neural network. We proceed up to convergence of our approximation.

In the case of the Time-Consistent policy, the state variables of the government problem are $x \equiv (k, g, \overline{\tau}^k, \overline{\tau}^l)$. The structure of the algorithm is similar to the one we use for the case of Full Commitment; the key distinction is that we now need to also approximate the derivatives S_x for state variables x. We perform this step by numerical approximation.

C Additional Numerical Results

We now provide a decomposition of the results in Figures 1 and 2, by considering the intermediate case in which costs of state contingency apply only to the capital tax. As the figures illustrate, this scenario leads to intermediate outcomes between our baseline model and the case with fully predetermined taxes.

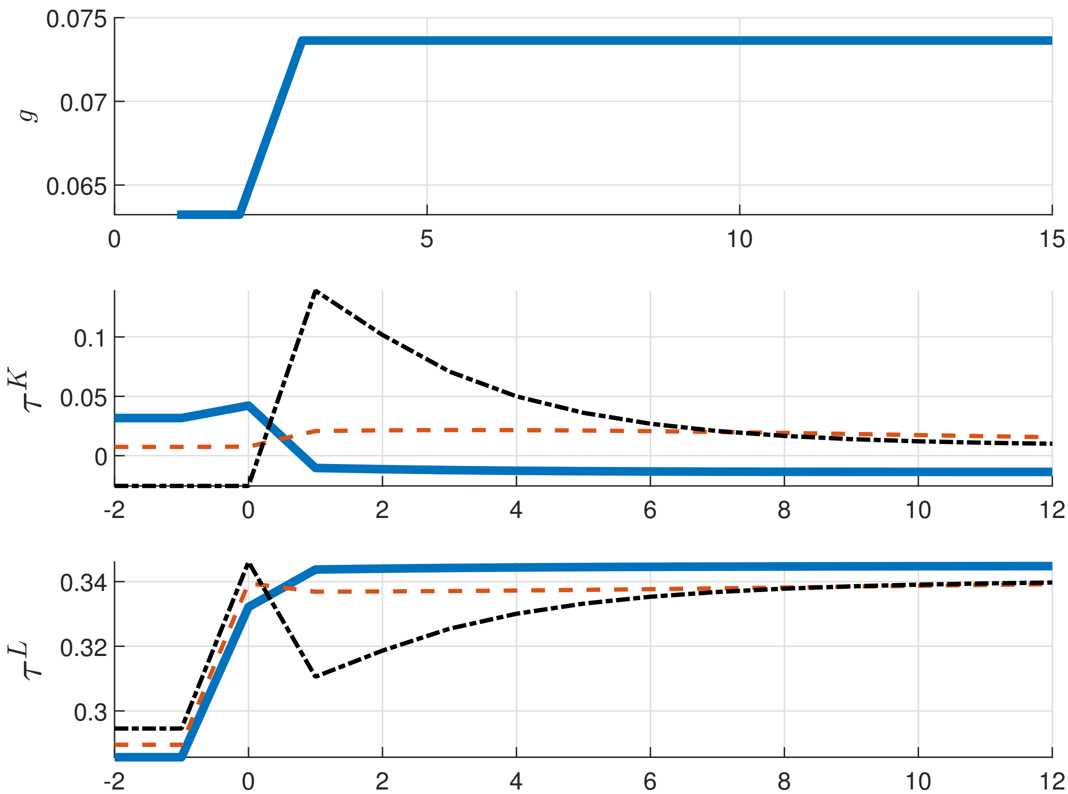

Figure C1: The figure displays the dynamics of fiscal variables around a shock that increases government spending. Top: government spending; middle: capital income tax rate; bottom: labor income tax rate. X-axes report time. Solid line: baseline model with costly state contingency ($\gamma_0^k = \gamma_0^l = 100$); dashed line: costs of state contingency on capital tax ($\gamma_0^k = 100, \gamma_0^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma_0^k = \infty, \gamma_0^l = 0$).

D State-Contingent Government Debt

We now consider the case in which the government can issue state contingent debt $b_t(g^t)$ and discuss the effects of costs of state contingency of tax instruments in this context. The government budget constraint is

$$b_t(g^t) = \tau_t^k r_t k_{t-1} + \tau_t^l w_t l_t - g_t + \sum_{g^{t+1}} q_t(g^{t+1}|g^t) b_{t+1}(g^{t+1}), \tag{D1}$$

where $q_t(g^{t+1}|g^t)$ is the price at time t of a debt instrument that pays one unit of consumption at $t+1$ contingent on the realization of history g^{t+1}. Household optimality implies that this price satisfies $q_t(g^{t+1}|g^t) = \beta p(g^{t+1}|g^t) \frac{u_c\left(c_{t+1}(g^{t+1})\right)}{u_c(c_t(g^t))}$, where $p(g^{t+1}|g^t)$ denotes the

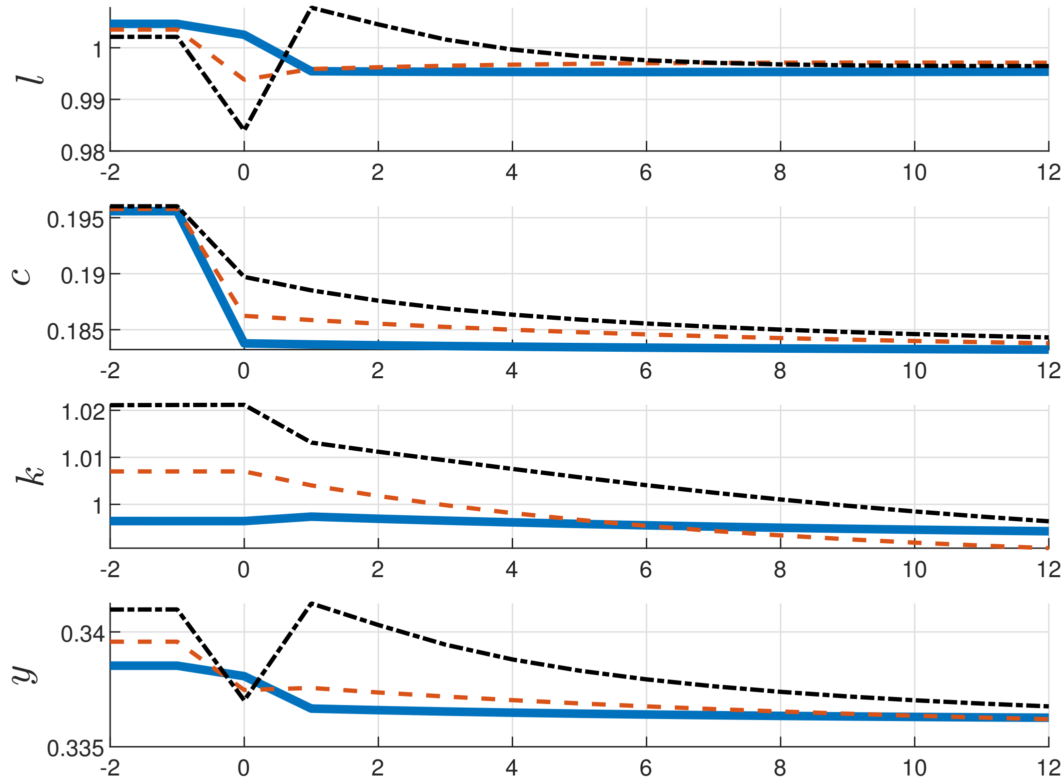

Figure C2: The figure displays the dynamics of allocations around a shock that increases government spending. From top to bottom: labor; consumption; capital; output. X-axes report time. Solid line: baseline model with costly state contingency ($\gamma_0^k = \gamma_0^l = 100$); dashed line: costs of state contingency on capital tax ($\gamma_0^k = 100, \gamma_0^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma_0^k = \infty, \gamma_0^l = 0$).

conditional probability of this history. In the interest of space, we avoid reformulating the rest of the household problem, which is unchanged.

By following standard steps (e.g. Chari and Kehoe, 1999), i.e., substituting in private sector optimality conditions and iterating forward on equation (D1) by recursively substituting out state-contingent debt, we obtain a single implementability constraint:

$$u_c(c_0)\left[b_{-1} + k_{-1} + (F_k(k_{-1}, l_0) - \delta)(1 - \tau_0^k)k_{-1}\right] = \mathbb{E}_0 \sum_{t=0}^{\infty} (u_c(c_t)c_t - v_l(l_t)l_t) \tag{D2}$$

We parameterize preferences and costs of state contingency consistent with our baseline calibration, that is: $u(c) \equiv \frac{c^{1-\eta_c}}{1-\eta_c}$, $v(l) \equiv \chi \frac{l^{1+\eta_l}}{1+\eta_l}$, and $\Gamma^j(\tau^j, \overline{\tau}^j) \equiv \frac{\gamma_0^j}{2}(\tau^j - \overline{\tau}^j)^2$ for $j = k, l$.

Consider first the case of no costs of state contingency, i.e., $\gamma_0^j = 0$ for $j = k, l$. In this case, given our utility function, the results of Chari and Kehoe (1999) imply that the labor tax rate is constant across states and over time. Furthermore, there is indeterminacy between state-by-state realizations of the capital tax and values of state-contingent debt. Multiple combinations of these variables are consistent with the same optimal allocation.

In particular, one implementation of the optimal allocation features a constant capital tax rate across states and over time, and the government using only state-contingent debt to absorb fluctuations in government spending. Next, notice that this policy with constant tax rates on both capital and labor is indeed optimal also when $\gamma_0^j > 0$ for $j = k, l$. Specifically, the government supports it by making non-contingent announcements about future tax rates that are equal to these constant realized tax rates, implying that the realized costs of state contingency are always equal to zero.

Other implementations of the allocations that are also optimal when $\gamma_0^j = 0$ imply variation in the capital tax rate across states. Thus, they are no longer optimal when $\gamma_0^j > 0$, because they involve positive costs of state contingency and are strictly dominated by the implementation with noncontingent taxes.

Hence, we find that costs of state contingency do not affect the optimal allocation when the government has access to state-contingent debt, but they do select the optimal implementation of this allocation, resolving the fundamental indeterminacy between the role of debt and capital taxes in absorbing fiscal shocks.