Capital and Labor Taxes with Costly State Contingency*

Alex Clymo§ Andrea Lanteri¶ Alessandro T. Villa‖

April 2023

Abstract

We analyze optimal capital and labor taxes in a model where (i) the government makes noncontingent announcements about future policies and (ii) state-contingent deviations from these announcements are costly. With Full Commitment, optimal announcements coincide with expected future taxes. Costly state contingency dampens the response of both current and future capital taxes to government spending shocks and labor taxes play a major role in accommodating fiscal shocks. These features allow our quantitative model to account for the volatility of taxes in US data. In the absence of Full Commitment, optimal announcements are instead strategically biased, because governments have an incentive to partially constrain their successors. The cost of deviating from past announcements generates an endogenous degree of fiscal commitment, determining the average level of capital taxes.

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

*We thank the editor, the associate editor, and two referees for insightful comments. We also thank Gadi Barlevy, Davide Debortoli, François Gourio, Anastasios Karantounias, Nobu Kiyotaki, Albert Marcet, Morten Ravn, and Nicolas Werquin, as well as seminar and conference participants at the Chicago Fed, Computation in Economics and Finance 2021, and IAE-CSIC for helpful feedback. Jui-Lin Chen provided excellent research assistance. Disclaimer: The views expressed in this paper do not represent the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Declaration of conflicts of interest: none.

§University of Essex. Email: a.clymo@essex.ac.uk.

¶Duke University and CEPR. Email: andrea.lanteri@duke.edu.

‖Federal Reserve Bank of Chicago. Email: alessandro.villa@chi.frb.org.

1 Introduction

Changes in fiscal policy, such as reforms of the tax code, are costly endeavors for governments, as they typically require parliamentary approval, sometimes involving lengthy negotiations. As a result, there are often lags between tax-policy *announcements*, which are based on an expected evolution of the future state of the economy, and *realized* policies, which in part reflect original plans and in part respond to new information about the economy.[1]

This institutional 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 the state contingency of government debt, it is standard to make rather stark assumptions on the degree of state contingency of taxes, often for convenience, abstracting from the difference between policy announcements and implementation. In most cases, the literature assumes that the government can freely change taxes in response to shocks; in some cases, instead, it assumes that labor taxes can freely adjust, whereas capital taxes cannot adjust contemporaneously. Importantly, these common assumptions typically lead to optimal capital taxes that are substantially more volatile than in the data, motivating us to develop a model of frictions in state contingency.[2]

More broadly, quantitative macroeconomic models typically feature substantial adjustment costs for the dynamic decisions of both households and firms, but, to our knowledge, little is known about the role of this type of friction in the context of government policy. What are the effects of costly state contingency of tax plans on the optimal dynamic mix of capital and labor taxes? To what extent can fiscal plans in the presence of costly state contingency substitute for a commitment technology, by partially constraining future government policy? To address these questions, in this paper we develop a new framework to analyze optimal capital and labor taxes when governments make optimal noncontingent announcements about future policies, and state-contingent deviations from these announcements are costly.

[1] In a quote that exemplifies how fiscal announcements are often not contingent on the future state of the economy and are subject to ex-post modifications, in 1988 then presidential candidate George H.W. Bush famously stated "Read my lips: no new taxes," although his administration later increased several taxes to reduce the budget deficit. In their empirical analysis of the effects of tax changes in US post-war data, Mertens and Ravn (2012) find that approximately half of the changes in the tax code have an implementation lag that exceeds 90 days and the median implementation lag is 6 quarters. Moreover, Alesina, Favero, and Giavazzi (2015) document frequent deviations of realized tax policy with respect to pre-announced tax policy in a large sample of multi-year fiscal plan in OECD countries.

[2] We provide a detailed discussion of the related literature in Section 2.

A key feature of our model is that pre-announced tax plans are a state variable that partially constrains government policy. To fix ideas, for each tax instrument, at time t we allow the government to make an announcement $\overline{\tau}_t$ about the policy to be implemented at time $t+1$. This announcement cannot depend on the realization 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 instrument, τ_{t+1}, in response to the shock that materializes. However, the government incurs a cost that increases in the distance between the realized policy and the previous announcement.

We embed these assumptions in a standard model of optimal capital and labor taxation, which we calibrate to closely match salient features of US post-war data on fiscal variables. We use this model to both demonstrate how a government facing costly state contingency would choose fiscal announcements and policy in response to shocks, and to 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 first consider a government with Full Commitment to tax plans into the infinite future, but subject to a quadratic cost of state contingency. In this case, we show that optimal fiscal announcements are simply unbiased *forecasts* of future policies. Moreover, because it is costly to adjust current taxes in response to a government spending shock, if we assume a balanced-budget rule then the tax base must instead adjust to satisfy the government budget constraint. As a result, when the government needs additional tax revenue, it largely relies on its announcements about future policies to induce a higher current level of output.

This force prevents future capital taxes from rising as much as in other models, and generates a major role for labor taxes in accommodating government spending shocks. Overall, this mechanism allows our model to match the standard deviation of taxes on capital and labor income in US data (around 2%), which we show is a challenge both for models without costs of state contingency and for models that treat capital taxes as predetermined. In our quantitative analysis, we first illustrate this mechanism under the assumption of a government balanced-budget constraint, which allows us to explain the mechanism transparently, and then show that it remains relevant when the government can issue noncontingent debt to finance its expenditures.

We then analyze optimal policy in an environment with commitment frictions, where successive governments make strategic one-period ahead noncontingent announcements. Each government inherits its predecessor's announced plan, but may reoptimize in a state-contingent fashion subject to a cost. We refer to this setup as Limited-Time Commitment, because it builds on and generalizes the framework of Clymo and Lanteri (2020). In this case, governments use fiscal announcements *strategically* not only to affect private sector allocations, but also, critically, to constrain future policy decisions, partly overcoming time inconsistency. Optimal fiscal announcements are no longer unbiased forecasts of future policies, and we derive a Generalized-Euler-Equation representation of the optimality conditions that highlights the presence of a strategic bias in fiscal announcements. We also leverage a two-period version of our model to relate this bias to model primitives in a transparent way. We then solve a quantitative version of the model and find that a calibrated degree of costly state contingency, consistent with empirical tax volatility, generates an endogenous level of fiscal commitment, sustaining allocations that are quite similar to the ones that we obtain with Full Commitment.

Noticeably, this regime leads to a positive, but small tax on capital income, of around 8%, while in the absence of costly state contingency governments would have a temptation to tax capital at confiscatory rates. Thus, the frictions that prevent the timely response of government policies to shocks may have the benefit of helping build commitment.

2 Related Literature

Our paper contributes to the large literature on optimal capital and labor income taxes, by introducing a new friction in the government problem, namely costly state contingency of tax plans. We highlight the relevance of this friction both for models of fiscal policy with Full Commitment and for models with commitment frictions. Furthermore, our paper contributes to the literature on the macroeconomic effects of fiscal announcements.

Optimal Capital and Labor Taxes with Full Commitment. Chari and Kehoe (1999) analyze optimal capital and labor income taxation in a stochastic economy under Full Commitment building on the early contributions 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; Farhi (2010) considers a more general incomplete-markets model with noncontingent 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.[3]

In these models, capital taxes are typically highly volatile in response to government spending shocks—significantly more so than in the data. Our contribution is to generalize the framework by introducing costs of state contingency for tax rates. By nesting previous assumptions on the measurability of taxes as special cases, we provide insights on the role of these assumptions for the optimal dynamic response of both current and future taxes to government spending shocks. In our calibrated model, costly state contingency allows us to account for the empirical volatility of capital and labor taxes. Moreover, our model produces empirically plausible conditional dynamics. In particular, when government spending increases, the government raises both capital and labor taxes persistently, consistent with the empirical evidence (Burnside, Eichenbaum, and Fisher, 2004). We also confirm and extend this evidence using the local-projection approach following Ramey and Zubairy (2018).

Partial Commitment in Fiscal Policy. Our framework with costly deviations from previous policy announcements is most closely related to the literature on intermediate notions of fiscal commitment. Specifically, Debortoli and Nunes (2010, 2013) analyze models of fiscal policy with stochastic government re-optimizations. A key contribution of our paper is that, in our framework, the degree to which governments renege on previous announcements is fully endogenous and depends on the state of the economy. In turn, this endogenous degree of commitment feeds back on strategic fiscal announcements. Related to our focus on the volatility of tax rates, Debortoli and Nunes (2010) obtain smooth capital taxes by assuming that capital utilization is endogenous.

Clymo and Lanteri (2020) introduce a framework in which the government has Limited-Time Commitment—i.e., successive governments fully 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 significantly generalize their framework by allowing governments to partially renege on previous noncontingent announcements, subject to a cost. Hence, our model 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.

Klein, Krusell, and Ríos-Rull (2008) characterize optimal capital taxation and public-good provision when the government lacks commitment using a Generalized Euler Equation.

[3] 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).

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. Relatedly, Karantounias (2019) uses a Generalized-Euler-Equation approach to analyze optimal taxation in a model with default and Ortigueira and Pereira (2021) use it to analyze the role of retroactive taxation for equilibrium multiplicity.

Fiscal Announcements. Our paper develops a theory of optimal fiscal announcements under uncertainty, when the government takes into account the effects of this announcements on the future costs of deviating from them in a state-contingent fashion. In so doing, the paper builds a bridge between the theoretical literature on optimal fiscal policy and the empirical and quantitative literature that studies the macroeconomic effects of announcements and expectations about future fiscal plans, distinguishing them from actually implemented fiscal policies, but often treating announcements as exogenous. See, for instance Mertens and Ravn (2012) and Alesina, Favero, and Giavazzi (2015) for empirical analyses, and Mertens and Ravn (2011) and Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015) for analyses based on quantitative macro models with several frictions.[4]

The rest of the paper is organized as follows. Section 3 describes our model. Section 4 characterizes optimal policy. Section 5 presents our quantitative analysis under Full Commitment. Section 6 analyzes the role of commitment frictions. Section 7 concludes.

3 Model

In this section, we describe an infinite-horizon model with capital and labor taxes and costly state contingency of fiscal plans.

[4]Our work is also related to the theoretical and quantitative body of work that studies the macroeconomic effects of fiscal rules, such as balanced-budget constraints. For instance, King, Plosser, and Rebelo (1988) find that balanced-budget rules amplify aggregate fluctuations; Schmitt-Grohe and Uribe (1997) find that a balanced-budget rule may induce indeterminacy. We find that even in a model without indeterminacy, balanced-budget rules, combined with costly state contingency of taxes, induce significant fluctuations in consumption. A theoretical literature studies the optimal design of policy rules. 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. We do not explicit microfound the origins of limited state contingency, and focus instead on the effect of costly state contingency on capital and labor taxes. Our approach is consistent with the notion that partial state contingency in fiscal policy may arise because of several reasons, including partial information about the state of the economy, as in Hauk, Lanteri, and Marcet (2021).

3.1 Environment

We consider a stochastic production economy populated by a continuum of identical households and a government.[5] Time is discrete and infinite, indexed by $t = 0, 1, 2, ...$ Households rank streams of consumption c_t and labor l_t according to the following utility function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - v(l_t) \right], \tag{1}$$

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

The resource constraint of the economy is given by

$$c_t + k_t + g_t = F(k_{t-1}, l_t) + (1 - \delta)k_{t-1}, \tag{2}$$

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

$$c_t + k_t + q_t b_t = w_t l_t (1 - \tau_t^l) + k_{t-1} \left[1 + r_t (1 - \tau_t^k) \right] + b_{t-1}, \tag{3}$$

where b_t are one-period risk-free bonds with price q_t, 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.

[5] Household heterogeneity in the context of limited state contingency of tax policy is an interesting avenue for research, that we leave for future work. In particular, when changing the tax code in response to shocks is costly, the government can use the degree of tax progressivity to obtain an endogenous degree of state contingency in average tax rates.

3.2 Household and Firm Optimality

Households maximize utility (1) subject to their budget constraint (3). The intratemporal labor-consumption margin and the Euler equations for savings in capital and bonds are

$$v_{l,t} = u_{c,t} w_t (1 - \tau_t^l), \tag{4}$$

$$u_{c,t} = \beta \mathbb{E}_t u_{c,t+1} \left[1 + r_{t+1}(1 - \tau_{t+1}^k) \right], \tag{5}$$

$$q_t u_{c,t} = \beta \mathbb{E}_t u_{c,t+1}. \tag{6}$$

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

$$w_t = F_{l,t}, \tag{7}$$

$$r_t = F_{k,t} - \delta. \tag{8}$$

Notice that our notation already imposes market clearing for labor and capital. The definition of a competitive equilibrium for a given government policy is standard.

3.3 Government

The government needs to finance spending g_t using capital and labor income taxes, as well as risk-free debt b_t, subject to the budget constraint

$$\tau_t^k r_t k_{t-1} + \tau_t^l w_t l_t + q_t b_t = g_t + b_{t-1}, \tag{9}$$

as well as the debt limits $b_t \geq b^{min}$ and $b_t \leq b^{max}$, with $b^{min} \leq 0 \leq b^{max}$. Note that the special case $b^{min} = b^{max} = 0$ imposes a balanced budget period by period; we will explore this case later in our analysis because its higher tractability allows us to obtain useful insights.

At date t, the government chooses current tax rates τ_t^k and τ_t^l, as well debt issuance b_t, which are 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 future state of the economy, and so are also measurable with respect to g^t.

Given initial conditions $k_{-1}, b_{-1}, \overline{\tau}_{-1}^k, \overline{\tau}_{-1}^l$, the government chooses stochastic sequences of current tax rates τ_t^k, τ_t^l, debt b_t, and future announcements $\overline{\tau}_t^k, \overline{\tau}_t^l$ to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[u(c_t) - v(l_t) - \Gamma^k(\tau_t^k, \overline{\tau}_{t-1}^k) - \Gamma^l(\tau_t^l, \overline{\tau}_{t-1}^l) \right], \tag{10}$$

where $\Gamma^j(\tau_t^j, \overline{\tau}_{t-1}^j)$ is a *cost of state contingency* for tax rate τ_t^j and $j \in \{k, l\}$, because the tax rate τ_t^j is measurable with respect to g^t, whereas the announcement $\overline{\tau}_{t-1}^j$ is measurable with respect to g^{t-1}. We assume that: (i) $\Gamma^j(\tau_t^j, \overline{\tau}_{t-1}^j) \geq 0$; (ii) $\Gamma^j(\tau_t^j, \overline{\tau}_{t-1}^j) = 0$ if $\tau_t^j = \overline{\tau}_{t-1}^j$; (iii) Γ^j is weakly increasing and weakly convex in a measure of distance between τ_t^j and $\overline{\tau}_{t-1}^j$. Thus, the cost functions Γ^j penalize deviations of state-contingent tax rates relative to the previously announced noncontingent plan. In our numerical application, we parameterize Γ^j as a quadratic function: $\Gamma^j(\tau_t^j, \overline{\tau}_{t-1}^j) \equiv \frac{\gamma^j}{2}(\tau_t^j - \overline{\tau}_{t-1}^j)^2$.[6] However, we emphasize that our framework is general and could accommodate other functional forms, including, for instance, fixed costs.

3.4 Discussion of Main Assumptions

In this section, we discuss the role of our main assumptions and relate costly state contingency with adjustment costs, that are often assumed in macroeconomic models.

3.4.1 Roles of Announcements and Costly State Contingency

Equation (10) 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 between the objective function of the government and that of households (1). However, because households take tax rates as given, nothing would change if we also added these costs in the household utility function. Moreover, we explore the difference between government welfare and household welfare quantitatively in Section 5.2 and find a negligible difference.

We highlight that costly state contingency plays two separate roles, depending on the government commitment regime. Specifically, when the government has Full Commitment, our assumption makes it costly for the government to let taxes differ depending on the realization of the government spending shock one period ahead. In this setup, the government can commit to future state-contingent plans and thus announcements are simply a modeling device to constrain the ex-post variation in realized tax rates.

[6] We also explore a case with asymmetric costs in Section 5.8.

When the government lacks commitment, costly state contingency plays the same role, as well as an additional role by creating partial commitment, because we assume that the announcements $\overline{\tau}^j_{t-1}$ are made by the government in power at $t-1$, whereas the realized taxes τ^j_t are chosen by the government in power at t. Hence, costly state contingency makes it costly to deviate from the announcements made by the previous government. In this regime, announcements play a more substantive role, as governments may potentially use them to relax current implementability constraints by manipulating the actions of the following government.

Reflecting this dual role of costly state contingency, in the remainder of the paper we first develop the Full-Commitment benchmark and show how costly state contingency allows the model to match the volatility of tax rates. Then we consider a framework with commitment frictions, which generalizes the Limited-Time Commitment model and combines costly state contingency with the additional channel of commitment building. We decompose these two roles in Section 6.3.

3.4.2 Costs of State Contingency vs. Adjustment Costs

We now clarify the distinction between our assumption of costs of state contingency and *adjustment costs*, which constitute a frequent assumption in dynamic models of household or firm behavior. Specifically, we argue that adjustment costs are a *special case* of our framework with costly state contingency and explain that our more general formulation is advantageous because it allows us to nest previous results in the literature on optimal taxation.

In a model with adjustment costs on taxes, the government would face a cost that depends on the difference between current realized taxes and past realized taxes. This is clearly a special case of our general framework, which can be obtained by imposing the following constraints: $\overline{\tau}^j_t = \tau^j_t$, for $j = k, l$. In words, the noncontingent announcement for future taxes must coincide with current realized taxes.

In our model, this restriction is not present. The government can choose current realized taxes and future announcements independently from each other, implying that policies can be adjusted costlessly with a one-period lag.

The higher degree of generality of our framework, and associated smaller departure from a standard model without costs is not the only reason to prefer our formulation. A second reason is that costs of state contingency that relate taxes to previous announcements allow our framework to nest several models in the literature as limiting case. In particular, the common assumption of predetermined noncontingent taxes on capital income (e.g., Farhi, 2010) arises as a special case of our model with $\gamma^k = \infty$ and $\gamma^l = 0$, implying that ex-post realized capital taxes must coincide with the previous announcement. A model of adjustment costs would not have this desirable property. In a similar fashion, the assumption of (noncontingent) Limited-Time Commitment that Clymo and Lanteri (2020) consider corresponds to the special case $\gamma^k = \gamma^l = \infty$.

4 Optimal Policy

In this section we characterize optimal policy both with Full Commitment and Limited-Time Commitment.

4.1 Optimal Policy with Full Commitment

We now consider optimal fiscal policy under the assumption that the government has Full Commitment. We begin the analysis by deriving the implementability constraints of the government problem. Because taxes enter the government objective (10), we do not apply the standard primal approach formulating objective and constraints only in terms of allocations. While we substitute out all prices using equations (6), (7), and (8), we include the competitive equilibrium conditions that link allocations to tax rates as additional constraints in the government problem.

Combining the government budget constraint with the representative firm's optimality conditions, we obtain

$$u_{c,t}b_{t-1} = u_{c,t}\left(\tau_t^k\left(F_{k,t}-\delta\right)k_{t-1} + \tau_t^l F_{l,t}l_t - g_t\right) + \beta b_t \mathbb{E}_t u_{c,t+1} \tag{11}$$

and we attach a multiplier ν_t to this constraint.

We then substitute the marginal product of labor into the household intratemporal optimality condition to obtain

$$v_{l,t} - F_{l,t}u_{c,t}(1-\tau_t^l) = 0, \tag{12}$$

with multiplier ξ_t.

Finally, we attach multipliers μ_t to the capital Euler equation (5), $\beta\underline{\phi}_t$, and $\beta\overline{\phi}_t$ to the lower bound and upper bound on debt respectively.

The government chooses stochastic sequences of taxes, announcements, and allocations to maximize (10) subject to the resource constraint (2) with associated Lagrange multiplier λ_t, the implementability constraints (5), (11), (12), and the debt limits. Notice that the government can commit to future state-contingent taxes, but faces a cost of making these taxes different from previous noncontingent announcements.

The first-order conditions with respect to c_t, l_t, k_t, and b_t are:

$$\lambda_t = u_{c,t} - \mu_t u_{cc,t} + \mu_{t-1} u_{cc,t} \left(1 + (F_{k,t} - \delta)\left(1 - \tau_t^k\right)\right)$$
$$+ \nu_t u_{cc,t} \left(\tau_t^k (F_{k,t} - \delta) k_{t-1} + \tau_t^l F_{l,t} l_t - g_t - b_{t-1}\right) + \nu_{t-1} b_{t-1} u_{cc,t} - \xi_t F_{l,t}\left(1 - \tau_t^l\right) u_{cc,t} \quad (13)$$

$$\lambda_t F_{l,t} = v_{l,t} - \mu_{t-1} u_{c,t} F_{kl,t}\left(1 - \tau_t^k\right)$$
$$- \nu_t u_{c,t}\left(\tau_t^k F_{kl,t} k_{t-1} + \tau_t^l (F_{ll,t} l_t + F_{l,t})\right) - \xi_t \left(v_{ll,t} - F_{ll,t} u_{c,t}\left(1 - \tau_t^l\right)\right) \quad (14)$$

$$\lambda_t = \beta \mathbb{E}_t \lambda_{t+1} (1 + F_{k,t+1} - \delta) + \mu_t \beta \mathbb{E}_t u_{c,t+1} F_{kk,t+1}\left(1 - \tau_{t+1}^k\right)$$
$$+\beta \mathbb{E}_t \nu_{t+1} u_{c,t+1}\left(\tau_{t+1}^k (F_{kk,t+1} k_t + F_{k,t+1} - \delta) + \tau_{t+1}^l F_{kl,t+1} l_{t+1}\right) - \beta \mathbb{E}_t \xi_{t+1} F_{kl,t+1} u_{c,t+1} (1 - \tau_{t+1}^l)$$
$$(15)$$

$$\nu_t \mathbb{E}_t u_{c,t+1} = \mathbb{E}_t \nu_{t+1} u_{c,t+1} - \underline{\phi}_t + \overline{\phi}_t. \quad (16)$$

The first-order conditions with respect to the tax rates τ_t^k and τ_t^l are:

$$\Gamma^k_{\tau^k,t} = -\mu_{t-1} u_{c,t} (F_{k,t} - \delta) + \nu_t u_{c,t} (F_{k,t} - \delta) k_{t-1} \quad (17)$$
$$\Gamma^l_{\tau^l,t} = \nu_t u_{c,t} F_{l,t} l_t + \xi_t F_{l,t} u_{c,t}. \quad (18)$$

Equations (17) and (18) highlight that the government trades off the effects of current taxes on the cost of state contingency with their effects on allocations.

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

$$\mathbb{E}_t \Gamma^k_{\overline{\tau}^k,t+1} = 0 \quad (19)$$
$$\mathbb{E}_t \Gamma^l_{\overline{\tau}^l,t+1} = 0. \quad (20)$$

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 (19) and (20) 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: $\Gamma^j(\tau^j, \overline{\tau}^j) \equiv \frac{\gamma^j}{2}(\tau^j - \overline{\tau}^j)^2$, with $\gamma^j > 0$. 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 is unbiased, coinciding with the expected realization of the future tax rate.

4.1.1 Special Case: Balanced-Budget Constraint

We consider as an instructive special case the assumption of government balanced budget (Stockman, 2001), that is, $b^{min} = b^{max} = 0$. We will then use this case to analyze the effects of commitment frictions.

With a balanced-budget constraint, the competitive-equilibrium conditions uniquely pin down the level of labor and consumption, given the state variables (k_{t-1}, g_t) and a choice of contemporaneous tax rates (τ_t^k, τ_t^l) (Clymo and Lanteri, 2020). The choice of current tax rates must induce this allocation of consumption and labor, in order to respect the government budget constraint and satisfy private sector optimality. To see this, notice that for given $(k_{t-1}, g_t, \tau_t^l, \tau_t^k)$, there is a unique level of labor supply l_t that satisfies the government budget constraint, given implicitly by the solution to

$$\tau_t^k \left(F_{k,t} - \delta\right) k_{t-1} + \tau_t^l F_{l,t} l_t - g_t = 0. \tag{21}$$

In turn, given this level of labor, there is a unique level of consumption consistent with the household intratemporal optimality condition, given by

$$c_t = u_c^{-1} \left(\frac{v_{l,t}}{F_l(k_{t-1}, l_t)(1 - \tau_t^l)} \right). \tag{22}$$

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:

$$l_t = h^l(k_{t-1}, g_t, \tau_t^l, \tau_t^k), \tag{23}$$

$$c_t = h^c(k_{t-1}, g_t, \tau_t^l, \tau_t^k). \tag{24}$$

Furthermore, using the balanced-budget constraint and the capital Euler equation (5) we obtain the following implementability constraint:

$$u_{c,t}k_t = \beta \mathbb{E}_t \left[u_{c,t+1}(c_{t+1} + k_{t+1}) - v_{l,t+1}l_{t+1} \right]. \tag{25}$$

The government problem is thus to maximize (10) subject to the resource constraint (2), as well as the implementability constraints (23), (24), and (25). We provide the optimality conditions under Full Commitment and balanced budget in Appendix A.1.

4.2 Optimal Policy with Limited-Time Commitment

We now consider a different assumption on the commitment technology. Specifically, we interpret the government sector as a succession of decision makers—one at each date t—without commitment to future realized policies. The government in power at t chooses current tax rates and makes announcements about future (one-period ahead) tax rates, thus internalizing their distortive effects on capital accumulation.

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 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.[7] 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.

To focus on the role of commitment frictions with costly state contingency, in this part of the analysis we impose the balanced-budget assumption.[8] Thus, 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

[7] The case of noncontingent Limited-Time Commitment that Clymo and Lanteri (2020) consider in a simpler model would be, in our framework, the case $\gamma^k = \gamma^l = \infty$. However, for sufficiently large fluctuations in government spending, an equilibrium of the model considered in the current paper does not exist when $\gamma^k = \gamma^l = \infty$, because of the balanced-budget constraint on the government. We verified that this is indeed the case under our calibration.

[8] There are additional complications in models with noncontingent government debt and lack of commitment; see, for instance (Krusell, Martin, and Ríos-Rull, 2004).

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:

$$\tilde{c}(x) = h^c(k, g, \tilde{\tau}^k(x), \tilde{\tau}^l(x)), \tag{26}$$

$$\tilde{l}(x) = h^l(k, g, \tilde{\tau}^k(x), \tilde{\tau}^l(x)). \tag{27}$$

Furthermore, let $\tilde{W}(x)$ be the present discounted value of government utility (10) associated with the policy functions introduced above, given the state of the economy x. Using this notation, we can state the optimization problem of a government as to choose allocations, taxes $(c, l, k', \tau^k, \tau^l)$, as well as announcements $(\overline{\tau}^{k,\prime}, \overline{\tau}^{l,\prime})$ to maximize

$$u(c) - v(l) - \Gamma^k(\tau^k, \overline{\tau}^k) - \Gamma^l(\tau^l, \overline{\tau}^l) + \beta \mathbb{E} \tilde{W}(x'), \tag{28}$$

subject to the resource constraint

$$c + k' + g = F(k, l) + (1 - \delta)k, \tag{29}$$

with associated multiplier λ, and the implementability constraints

$$u_c(c)k' = \beta \mathbb{E} \left[u_c(\tilde{c}(x')) \left(\tilde{c}(x') + \tilde{k}'(x') \right) - v_l(\tilde{l}(x'))\tilde{l}(x') \right], \tag{30}$$

with multipier μ, and

$$l = h^l(k, g, \tau^k, \tau^l), \tag{31}$$

$$c = h^c(k, g, \tau^k, \tau^l), \tag{32}$$

with multipliers ν^l and ν^c respectively.[9]

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

$$\lambda = u_c(c) - \mu u_{cc}(c)k' - \nu^c, \tag{33}$$

$$v_l(l) = \lambda F_l(k,l) + \nu^l, \tag{34}$$

$$\lambda = \beta \mathbb{E}\tilde{W}_k(x') - \mu u_c(c) + \mu\beta\mathbb{E}S_k(x'), \tag{35}$$

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 (30), 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 (30) are absent 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

$$\nu^c h^c_{\tau^k}(k,g,\tau^k,\tau^l) - \nu^l h^l_{\tau^k}(k,g,\tau^k,\tau^l) = \Gamma^k_{\tau^k}(\tau^k,\overline{\tau}^k), \tag{36}$$

$$\nu^c h^c_{\tau^l}(k,g,\tau^k,\tau^l) - \nu^l h^l_{\tau^l}(k,g,\tau^k,\tau^l) = \Gamma^l_{\tau^l}(\tau^l,\overline{\tau}^l). \tag{37}$$

As in the Full-Commitment case, these optimality conditions trade off the effects of taxes on allocations with their effects on the costs of state contingency.

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

$$\mathbb{E}\tilde{W}_{\overline{\tau}^k}(x') + \mu\mathbb{E}S_{\overline{\tau}^k}(x') = 0, \tag{38}$$

$$\mathbb{E}\tilde{W}_{\overline{\tau}^l}(x') + \mu\mathbb{E}S_{\overline{\tau}^l}(x') = 0. \tag{39}$$

[9]Technically, we also need to 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. However, we abstract from this constraint in the text because a sufficiently large bound (50%) never binds in equilibrium for the degrees of costly state contingency that we consider in our quantitative analysis.

Furthermore, we have the following envelope conditions:

$$\tilde{W}_k(x) = \lambda\left[F_k(k,l) + (1-\delta)\right] - \nu^l h_k^l + \nu^c h_k^c, \tag{40}$$

$$\tilde{W}_{\overline{\tau}^k}(x) = -\Gamma^k_{\overline{\tau}^k}(\tau^k, \overline{\tau}^k), \tag{41}$$

$$\tilde{W}_{\overline{\tau}^l}(x) = -\Gamma^l_{\overline{\tau}^l}(\tau^l, \overline{\tau}^l). \tag{42}$$

These optimality conditions reveal a key distinction with respect to the Full-Commitment problem: Optimal announcements are not set just to minimize the expected costs of state contingency. Instead, they are strategically biased, because they optimally trade off the incentive to reduce expected cost of state contingency with the possibility to manipulate the following government's problem by setting its inherited fiscal announcements, thus relaxing current implementability constraints.

In particular, by combining the envelope conditions (41) and (42) with equations (38) and (39) we derive the Generalized Euler Equations for the optimal announcements:

$$\mathbb{E}\Gamma^k_{\overline{\tau}^k}(\tau^k, \overline{\tau}^k) = \mu \mathbb{E} S_{\overline{\tau}^k}(x'), \tag{43}$$

$$\mathbb{E}\Gamma^l_{\overline{\tau}^l}(\tau^l, \overline{\tau}^l) = \mu \mathbb{E} S_{\overline{\tau}^l}(x'). \tag{44}$$

This strategic incentive is reflected in the presence of the terms $\mu \mathbb{E} S_{\overline{\tau}^k}(x')$ and $\mu \mathbb{E} S_{\overline{\tau}^l}(x')$ in these Generalized Euler Equations for the optimal announcements, which are absent in (19) and (20). As indicated by the presence of μ, which is the multiplier on the forward looking constraint (30), the strategic bias arises because the government has an incentive to use fiscal announcements to manipulate the choices of future governments. However, because these terms are of difficult interpretation in this general dynamic model, we now consider a simpler, two-period version of our model to gain further intuition on the strategic bias in fiscal announcements.

4.2.1 Inspecting the Mechanism: Analytical Results in a Two-Period Model

To further investigate the role of costly state contingency in building partial fiscal commitment through strategic government announcements, in Appendix B we develop a simplified, two-period version of our model and use it to derive analytical insights. This analysis complements the characterization of our infinite-horizon model, for which we must rely on numerical results.

In the interest of space, we focus here on a key result from the two-period model, which explicitly relates fiscal announcements to expectations and expresses the strategic bias in terms of primitives. For concreteness, we make the following assumptions. In the first period, households value consumption linearly and make an investment decision out of an exogenous endowment; in the second period, preferences are given by $u(c) \equiv \log(c)$ and $v(l) \equiv \chi\frac{l^{1+\eta}}{1+\eta}$ and the technology is $y = zk^{\alpha}l^{1-\alpha}$. The government makes noncontingent announcements for capital and labor tax rates before a government spending shock is realized in the second period. After observing the realization of the shock, the government chooses its policy subject to quadratic costs of state contingency, where γ^k and γ^l scale the size of these costs.

We denote by $h(k, g, \tau^l)$ the level of hours worked in the second period consistent with state variables k and g and implemented tax rate τ^l. Moreover, we denote by $\tilde{\tau}^l(k, g, \overline{\tau}^k, \overline{\tau}^l)$ the optimal labor tax policy as a function of states and fiscal promises.

We can then express optimal fiscal announcements—the counterparts of equations (43) and (44)—as follows:

$$\overline{\tau}^k = \mathbb{E}\tau^k - \frac{\chi(1+\eta)\mu}{\gamma^k}\mathbb{E}\left[l^{\eta}h_{\tau^l}(k, g, \tau^l)\tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l)\right], \tag{45}$$

$$\overline{\tau}^l = \mathbb{E}\tau^l - \frac{\chi(1+\eta)\mu}{\gamma^l}\mathbb{E}\left[l^{\eta}h_{\tau^l}(k, g, \tau^l)\tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l)\right], \tag{46}$$

where $\mu \geq 0$ is the Lagrange multiplier on the implementability constraint. The optimality conditions allow a clear interpretation of the bias between promises and expected tax rates. Firstly, this bias is inversely proportional to the cost of state contingency. For low costs (low γ^k, γ^l) the government must use larger biases to influence the behavior of the next government. Secondly, the sign of the bias depends on the product of two key terms: first, the marginal effect of realized taxes on the allocation, represented by the term h_{τ^l}; second, the marginal effect of announced taxes on realized taxes, represented by the partial derivatives of $\tilde{\tau}^l$ with respect to the announcements.

In this model $h_{\tau^l}(k, g, \tau^l) < 0$ as raising labor taxes must reduce labor in equilibrium. Hence (45) implies that the sign of the bias in capital tax announcements is determined by the sign of $\tilde{\tau}^l_{\overline{\tau}^k}$. As long as announcing lower future capital taxes induces lower realized capital taxes (and hence higher realized labor taxes, meaning $\tilde{\tau}^l_{\overline{\tau}^k} < 0$), the optimal non-contingent announcement for the capital tax is therefore lower than the average realized capital tax. Therefore, because of costly state contingency, fiscal announcements sustain a high level of investment by announcing low future capital taxes, thus partially constraining the future government to set relatively low taxes on capital income. In Appendix B we prove analytically that this is necessarily the case under some regularity conditions.

5 Quantitative Analysis with Full Commitment

In this section, we calibrate our model and discuss our quantitative results on optimal policy with Full Commitment, comparing them with the empirical evidence.

5.1 Calibration and Solution Method

We parameterize the utility function as follows: $u(c) \equiv \log(c)$ and $v(l) \equiv \chi\frac{l^{1+\eta}}{1+\eta}$, with $\eta = 2$, which is in a standard range considered in this literature (e.g., Bhandari, Evans, Golosov, and T.Sargent, 2017). We then set the value of χ to normalize average labor to one in the steady state of the Full-Commitment model. The production function is Cobb-Douglas, with capital share α: $F(k,l) \equiv zk^{\alpha}l^{1-\alpha}$, with $\alpha = 0.36$ to match the labor share of output, and set the value of z to normalize average capital to one in the steady state of the Full-Commitment model.

We calibrate the Markov process for g_t as an AR(1) in logs, formally: $\log g_{t+1} = (1 - \rho_g)\log\mu_g + \rho_g \log g_t + \epsilon_t$, with ϵ_t normally distributed with mean zero and standard deviation σ_g. We set the value of μ_g to match the average ratio of government spending to output, which is around 20%. We then estimate ρ_g and σ_g using linearly detrended US annual data, and discretize this process with a two-valued Markov chain. In Appendix C.2, we also consider an alternative parameterization of this process, with less persistence and higher variance of innovations, as in Farhi (2010), and show that our main mechanism is robust to this modification.

Our main sample for fiscal variables is 1971-2013 Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015), which focuses on tax volatility in the US.[10] The data moments on tax rates are summarized in Table 2, along with moments of our model. We use this evidence to calibrate the costs of state contingency, which we parameterize as follows: $\Gamma^j(\tau^j, \overline{\tau}^j) \equiv \frac{\gamma^j}{2}(\tau^j - \overline{\tau}^j)^2$ with $\gamma^j > 0$ for $j = k, l$. As a baseline case, we consider a parsimonious formulation imposing $\gamma^k = \gamma^l = \gamma$, thus treating the two instruments symmetrically. We then decompose the role of each cost in Section 5.4. Using our model under

[10] We also verify that our main moments of interest are similar in a longer sample. Appendix D provides details on the data and our procedure to calculate average tax rates, which follows the empirical and quantitative literature on fiscal policy.

Full Commitment and balanced budget, we calibrate γ to closely match the standard deviation of the capital tax rate in the data, which is approximately equal to 2%.[11] This gives $\gamma = 39$. We thus adopt a similar calibration strategy to the one typically adopted in models of firm investment, in which convex costs are often identified by investment volatility. We then recalibrate the parameter γ in the version of the model with noncontingent debt to match the same target, obtaining $\gamma = 5$. In this version of the model, we set the limits on debt b^{min} and b^{max} so that debt fluctuates between (approximately) -10% and 100% of steady-state output.

To interpret the magnitude of our calibrated costs of state contingency, we can express the consumption equivalent of a 1% deviation—i.e., $\tau_t^j - \overline{\tau}_{t-1}^j = 0.01$—as $1 - \exp\left(-\frac{\gamma}{2} \times 0.01^2\right)$, which yields a consumption equivalent of approximately 0.2% with $\gamma = 39$ and 0.025% with $\gamma = 5$.

Table 1: Parameter Values

		Parameter	Value
Preferences	Discount factor	β	0.96
	Labor disutility	χ	0.78
	Labor elasticity	η	2
Technology	Capital share	α	0.36
	Depreciation	δ	0.08
Government shock	Average g	μ_g	0.068
	Volatility of $\log(g)$	σ_g	0.016
	Autocorr. of $\log(g)$	ρ_g	0.977
Cost of state cont.	τ^k (balanced budget)	γ^k	39
	τ^l (balanced budget)	γ^l	39
	τ^k (debt)	γ^k	5
	τ^l (debt)	γ^l	5
Debt limits	Lower bound (balanced budget)	b^{min}	0
	Upper bound (balanced budget)	b^{max}	0
	Lower bound (debt)	b^{min}	-0.035
	Upper bound (debt)	b^{max}	0.338

Notes: The table reports the calibrated parameter values. See text for details.

[11] In Appendix D.4 we discuss the role of different components of capital taxes (e.g., corporate taxes and personal income taxes) for this volatility.

Using our balanced-budget model, we show numerically in Figure C1 in Appendix C.2 that parameter γ is identified by the volatility of the capital tax rate. We also show that to match this moment it is critical to assume that costly state contingency applies to both capital and labor taxes. Moreover, our calibrated value for γ delivers a standard deviation of labor taxes close to the data (1.8% in the model, 1.5% in the data), which is an untargeted moment. We then verify that this value for γ induces volatilities for capital and labor taxes close to their empirical counterparts also under Limited-Time Commitment.

We solve the model using a generalization of the Parameterized Expectations Algorithm (den Haan and Marcet, 1990) proposed by Valaitis and Villa (2023). This method relies on a neural network to approximate the forward looking terms in the optimality conditions, as functions of the state vector. For the Limited-Time Commitment problem, we further adapt these methods to numerically approximate the derivatives of the policy functions in the Generalized Euler Equation. To assess the accuracy of our solution method, we verify that two limiting cases of our model converge to an independent solution obtained using a projection-based approach.[12] We provide more details on the solution method in Appendix C.1.

5.2 Volatility of Taxes with Balanced Budget

We now discuss the dynamics of taxes with balanced budget, focusing first on long-run simulation moments. We compare policies in our calibrated model with two alternative parameterizations of our model, which nest previous work in the literature. The first comparison model sets $\gamma^k = \infty, \gamma^l = 0$; in this case, the capital tax must be chosen one period in advance, whereas the labor tax is freely adjustable within the period.[13] This timing assumption is common in the literature (e.g., Farhi, 2010). The second comparison model removes costly state contingency entirely, setting $\gamma^k = \gamma^l = 0$, and thus coincides with the model that Stockman (2001) analyzes. For clarity we refer to our calibrated model as the baseline model, and the comparison models as "predetermined capital taxes" and "no costly state contingency" respectively.

[12] Specifically, we compute the solution to Stockman (2001) using a projection-based approach and we verify that the solution of our model solved with Valaitis and Villa (2023) converges to that solution for $\gamma^k \to 0$ and $\gamma^l \to 0$. Moreover, we compute the solution to a predetermined capital taxes model, in a similar fashion to Farhi (2010), using a projection-based approach and we verify that the solution of our model solved with Valaitis and Villa (2023) converges to that solution for $\gamma^k \to \infty$ and $\gamma^l \to 0$.

[13] We approximate the case $\gamma^j = 0$ with $\gamma^j = 0.01$. This approximation allows us to compute the alternative models using the exact same solution method as in our calibrated model, thus facilitating the comparison.

We consider a long simulation (10,000 periods) of our calibrated model under costly state contingency and the two comparison models. In Table 2, we report first and second moments of capital and labor income tax rates and compare them under different policy regimes. The last column reports the empirical counterparts in US data.

All three models considered generate a substantially lower average capital tax than the one in the data, due to the assumption of Full Commitment. This moment is approximately zero under all three Full Commitment models, whereas the empirical counterpart is 36%. Accordingly, the labor tax tends to be higher in all models than in the data. In Section 6, we explore the role of commitment frictions and costly state contingency for the average capital tax rate.

Table 2: Full Commitment: First and Second Moments of Tax Rates

Moment	BB CSC	BB no CSC	BB pred. τ^k	Debt CSC	Debt pred. τ^k	Data
$\mathrm{E}\tau^k$	0.001	-0.003	-0.009	0	-0.003	0.355
$\mathrm{E}\tau^l$	0.314	0.314	0.314	0.312	0.314	0.226
St. Dev. $\log(1+\tau^k)$	0.021	0.064	0.045	0.022	0.227	0.022
St. Dev $\log(1+\tau^l)$	0.018	0.017	0.017	0.021	0.016	0.015
Autocorr. $\log(1+\tau^k)$	0.835	0.819	0.832	0.913	0.15	0.868
Autocorr. $\log(1+\tau^l)$	0.97	0.991	0.937	0.998	0.986	0.876

Notes: The table reports first and second moments of the tax rates on capital and on labor income in the models with Full Commitment. The first two rows report the means; the third and fourth row reports the standard deviations; the fifth and sixth row report autocorrelations. The first column refers to the baseline calibration with balanced budgets and Full Commitment ($\gamma^k = \gamma^l = 39$); the second column refers to the balanced budget model without costs of state contingency ($\gamma^k = \gamma^l = 0$); the third column refers to the balanced budget model with predetermined capital tax ($\gamma^k = \infty$,$\gamma^l = 0$); the fourth column refers to the baseline calibration with debt ($\gamma^k = \gamma^l = 5$); the fifth column refers to the predetermined capital tax calibration with debt ($\gamma^k = \infty$,$\gamma^l = 0$); the sixth column refers to US data from 1971-2013 at annual frequency.

However, it is in the second moments that costly state contingency plays a crucial role. In particular, our baseline model matches both the standard deviation of the capital tax (targeted) and also closely matches that of the labor tax (untargeted). In contrast, the two comparison models (with no cost of state contingency, or with predetermined capital taxes) overstate the volatility of taxes by factors of three and two respectively. The autocorrelation coefficient of both tax rates are instead closely aligned in all models and data, displaying high persistence.

Furthermore, the standard deviations of labor taxes appear similar across all three models, and are all in line with the data. However, this hides significant differences in the dynamics of labor taxes in response to a government spending shock, which will become evident in the next subsection, which focuses on conditional dynamics.

We also use our simulations to assess the difference in utility between household and government due to the presence of costs of state contingency in the government objective. We find that this cost is small and equals approximately 0.05% of permanent consumption. This finding suggests that our main results do not depend on assuming a large wedge between household and government welfare, nor on an implausibly large calibrated cost γ.

5.3 Conditional Dynamics: Government Spending Shock

In Figure 1, we illustrate the response of taxes on capital and labor income to an exogenous increase in government spending in our calibrated model and in the two comparison models. Specifically, the figure plots the response of the economy to a switch from the low to high government spending state, after a long spell in the low state.

In our baseline model with costly state contingency (solid line), the shock (at $t = 0$) induces the government to increase future capital taxes moderately, and labor taxes in a persistent way. In both comparison models, instead, the capital tax increases more substantially, either contemporaneously in the model without costs of state-contingency (dashed line) or with a one period lag when it is 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 the costs of state contingency apply when the shock hits, but do not prevent a lagged large adjustment. Furthermore, our estimated government spending shock is highly 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. Nevertheless, the government instead optimally chooses to promise a lower future capital tax than in the two comparison models.

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 immediate consumption drop in response to the shock, whereas consumption is relatively smoother in the two comparison models.[14]

[14] In Table C1 in Appendix C.2 we complement this analysis by showing the first and second moments of allocations, comparing our baseline model with the two alternative models. Consistent with our analysis

Figure 1: Full Commitment with Balanced Budget: Taxes

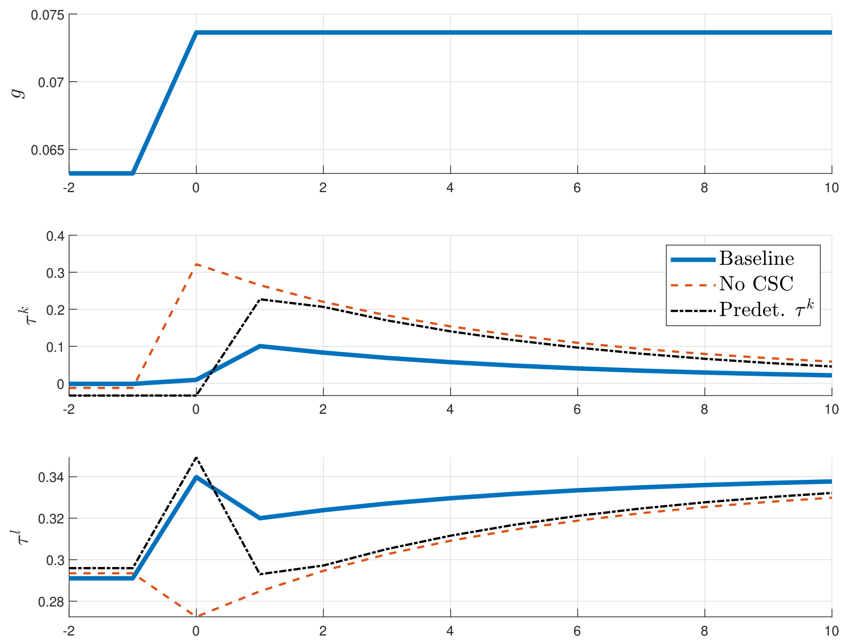

Notes: The figure displays the dynamics of fiscal variables around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. Top: government spending g_t; middle: capital income tax rate τ_t^k; bottom: labor income tax rate τ_t^l. Solid line: baseline model with costly state contingency ($\gamma^k = \gamma^l = 39$); dashed line: no costs of state contingency ($\gamma^k = \gamma^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma^k = \infty, \gamma^l = 0$).

To see why the model with costly state contingency behaves so differently, 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 (23). In turn, this level of labor dictates a level of consumption through equation (24), and a level of future capital through the resource constraint (2). Following an increase in government spending, in order to induce households to exert the required level of labor effort, the government must engineer a drop in current consumption, leveraging the effect of the marginal utility of consumption on labor supply.

In turn, to encourage consumption to fall, the government must ensure a sustained level of investment to satisfy the resource constraint. Encouraging higher investment to implement this allocation requires relatively low future capital taxes, driving a key dif-

from the impulse response functions, we find that costly state contingency leads to higher volatility of consumption than the models with more flexible taxes.

ference with respect to the comparison models, where taxes are easier to adjust, and the government is able to use other tools to balance the budget.

Furthermore, instead of relying only on high capital taxes to generate additional revenue, the government in our model increases labor taxes quickly and persistently. The response in the baseline model has remarkably different dynamics to the two comparison models, where labor taxes instead only gradually rise as capital taxes are withdrawn, and additionally feature either a positive or negative spike in labor taxes at the time of the shock. We show in Section 5.6 that the conditional dynamics of taxes in our model—especially once debt is included—are consistent with those in the data.

In particular, the response of both taxes on impact ($t = 0$) is relatively muted in the baseline model, with capital taxes moving little, and labor taxes rising slightly less than in the model with predetermined capital taxes.[15] Overall, to replace the lost increases in tax revenue relative to the more flexible models, the government with costly state contingency ensures that labor supply only falls by 1.4%, rather than the larger decrease seen in the model with predetermined capital tax. It does this by engineering a 3.8% fall in consumption in order to stimulate labor supply, which would have fallen by more than twice as much if consumption instead was held at its $t = -1$ value.

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, and, as we showed, announcements are unbiased under Full Commitment, realized tax rates are typically quite close to promised tax rates, except in the periods in which the value of government spending increases, when the government deviates from the noncontingent announcement to generate additional revenue.

In sum, by analyzing 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 period in which the shock hits. In contrast, calibrated costs of state contingency on *both* capital and labor taxes do not just induce otherwise optimal policies to be implemented with a lag. Instead, the government actively uses future policies to

[15] Notice that while adjusting either tax rate relative to the promise is equally costly in our parameterization ($\gamma^k = \gamma^l$), a one-percent increase in the labor tax generates more revenue than an equal increase in the capital tax, because the labor share of income is larger than the capital share. This fact contributes to explain why the government optimally adjusts labor taxes only slightly less than in the predetermined-capital-tax model when the government spending shock hits.

Figure 2: Full Commitment with Balanced Budget: Allocations

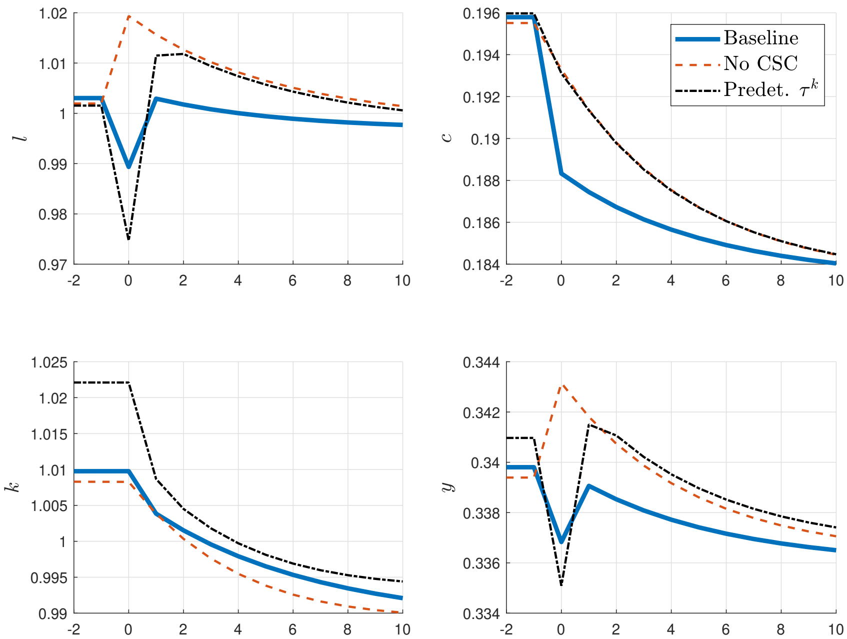

Notes: The figure displays the dynamics of allocations around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. Panels give labor l_t; consumption c_t; capital k_t; output y_t. Solid line: baseline model with costly state contingency ($\gamma^k = \gamma^l = 39$); dashed line: no costs of state contingency ($\gamma^k = \gamma^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma^k = \infty, \gamma^l = 0$).

ensure that the current budget constraint is satisfied, altering the optimal dynamic mix of capital and labor taxes and generating a major role of labor taxes in response to government spending shocks.

5.4 Decomposition: Cost on Individual Tax Instruments

To highlight the importance of considering costs of state contingency on both tax instruments, we now perform a decomposition of our results, by solving two counterfactual models where we impose our calibrated costs only on one tax at a time. Specifically, we solve a model with costs only on the capital tax ($\gamma^k = 39, \gamma^l = 0$) and a model with costs only on the labor tax ($\gamma^k = 0, \gamma^l = 39$), and compare these results with our baseline model.

We plot the results from these models in Figure 4. In panel (a) we display the dynamics of the model with costs on capital taxes only. The dynamics of tax rates in this counterfactual model are very similar to the ones that we obtained earlier when the capital tax

Figure 3: Full Commitment with Balanced Budget: Fiscal Announcements

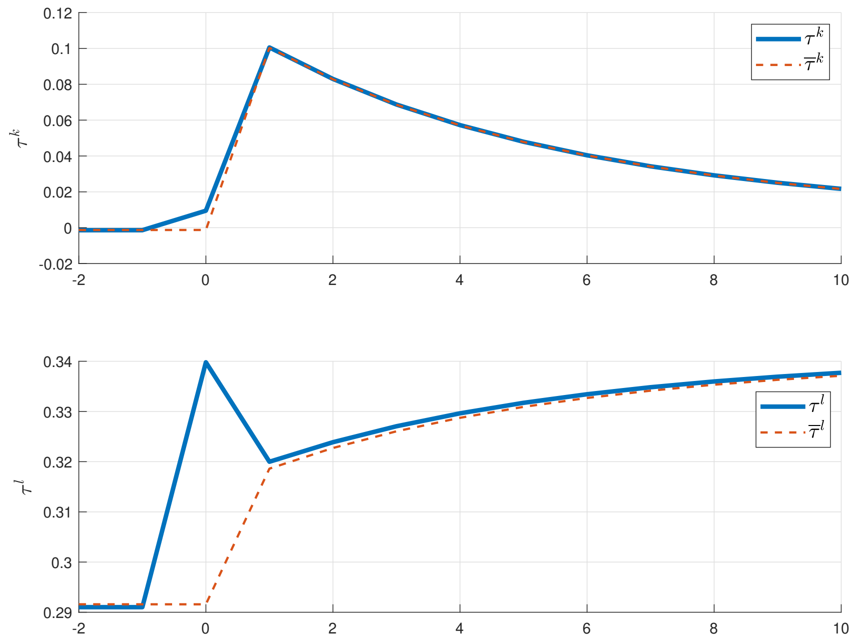

Notes: The figure compares the dynamics of realized tax rates (solid line) and announced tax rates (dashed line) around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. Top: capital income tax rate (τ_t^k and $\overline{\tau}_t^k$); bottom: labor income tax rate (τ_t^l and $\overline{\tau}_t^l$).

is fully predetermined. Thus, despite the cost of adjusting capital taxes being finite, the government prefers to avoid changing capital taxes if adjusting labor taxes is costless.

In panel (b) we display the dynamics of the model with costs on labor taxes only. Again, with costs on only one tax, the government prefers to treat this tax (in this case the labor tax) as essentially predetermined, and adjusts only the capital tax in the first period. However, from $t = 1$ onwards, the government lowers the labor tax, just as it does in the model with no costs on either tax.

This experiment clarifies the importance of considering frictions in adjusting both tax instruments. In particular, the interactions between the costs on adjusting either tax are highly nonlinear, and the optimal policy in the presence of both costs is not simply an average of the optimal policies in response to each cost in isolation. For example, capital taxes are more volatile in both counterfactual models than in the model with both costs, reaching peaks of over 20% in the former, and only 10% in the latter, during the transition. Similarly, the government with costs on both taxes significantly raises the labor tax from

Figure 4: Decomposition: Cost on Individual Tax Instruments

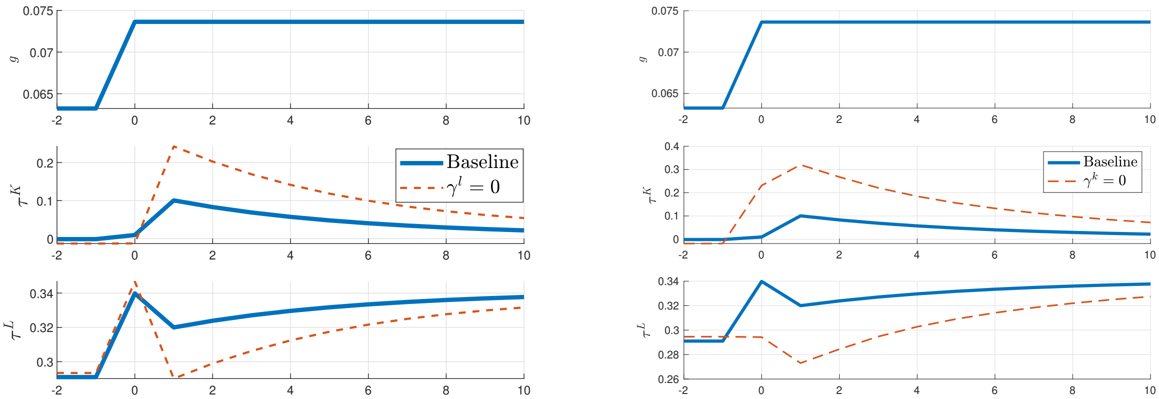

(a) Cost Only on Capital Tax

(b) Cost Only on Labor Tax

Notes: The figures display the dynamics of fiscal variables around a shock that increases government spending at time $t = 0$, under Full Commitment and balanced budget. Horizontal axes report time t. Top: government spending; middle: capital income tax rate; bottom: labor income tax rate. In panel (a), the solid line gives baseline model with costly state contingency ($\gamma^k = 39, \gamma^l = 39$) and the dashed line gives costs of state contingency only on capital tax ($\gamma^k = 39, \gamma^l = 0$). In panel (b) the dashed line instead gives costs of state contingency only on labor tax ($\gamma^k = 0, \gamma^l = 39$).

period one onwards, while this tax grows more slowly in both models with a cost only on one tax.

Intuitively, when there is a cost only on a single tax, the government essentially just delays adjustment of that tax by one period, while leaving the rest of the plan similar to the optimal plan when there are no costs. If instead there are costs on both taxes, delaying adjustment of both taxes is not feasible without violating the government's budget constraint in the period of the shock, and the government is forced to radically alter its plans. In particular, if it is costly to adjust contemporaneous taxes to boost tax revenue in response to the shock, the government uses future promises more actively to boost the current tax base.

An implication of this result is that it is not an innocuous assumption to place timing restrictions on the adjustment of one tax only. As a further illustration of this finding, in Appendix C.2, we show numerically that when $\gamma^l = 0$, there is no value of γ^k that can generate a volatility of the capital tax as low as in the data. Thus, in order to match the lower empirical volatility of capital taxes it is necessary to place state-contingency costs on both capital *and* labor taxes.

5.5 Role of Government Debt

We now relax the assumption of a balanced budget and allow the government to respond to government spending shocks by running fiscal deficits. To this end, we recalibrate the model with noncontingent government debt to reproduce the empirical volatility of capital taxes. We obtain a parameter value $\gamma = 5$ for the costs of state contingency.

We compare the predictions of the baseline model with the special case of predetermined capital taxes, as in Farhi (2010), which also considers noncontingent government debt. As columns 4 and 5 of Table 2 show, the model with predetermined capital tax predicts a volatility of capital taxes that is approximately ten times as large as the empirical counterpart.[16] Moreover, the calibrated model replicates the high persistence of taxes, whereas the model with predetermined capital tax generates a low autocorrelation for the capital tax.

Figure 5 reveals the reason for these findings. When government spending increases, the government facing costly state contingency raises taxes moderately and persistently. In contrast, in the comparison model the government engineers a large and short-lived spike in the capital tax in the next period. This spike takes the capital-tax rate from approximately 0 to approximately 100% and accounts for the bulk of the policy response to the shock. This spike is larger than in the balanced-budget case because, as Farhi (2010) explains, the government uses future capital taxes to manipulate the interest rate.

Moreover, the figure uncovers two additional salient differences between the calibrated model and the one with predetermined capital taxes. First, with costly state contingency, capital and labor taxes comove positively in response to the shock, whereas the comparison model predicts a temporary labor-tax cut. Second, with costly state contingency the government uses debt to gradually absorb the shock. In contrast, in the comparison model an initial deficit is followed by a reduction in government debt as the government obtains the proceeds of the capital-tax spike.

All these important differences in the conditional dynamics allow us to discriminate across these models using empirical evidence on the responses of taxes and debt to an exogenous increase in government spending in the next section.

[16] Because the model with predetermined capital tax predicts very large positive and negative spikes in the capital tax rate, we also compute its standard deviation without the log transformation; the standard deviation of τ^k equals 0.115, thus approximately five times larger than in the baseline model.

Figure 5: Full Commitment with Debt: Taxes

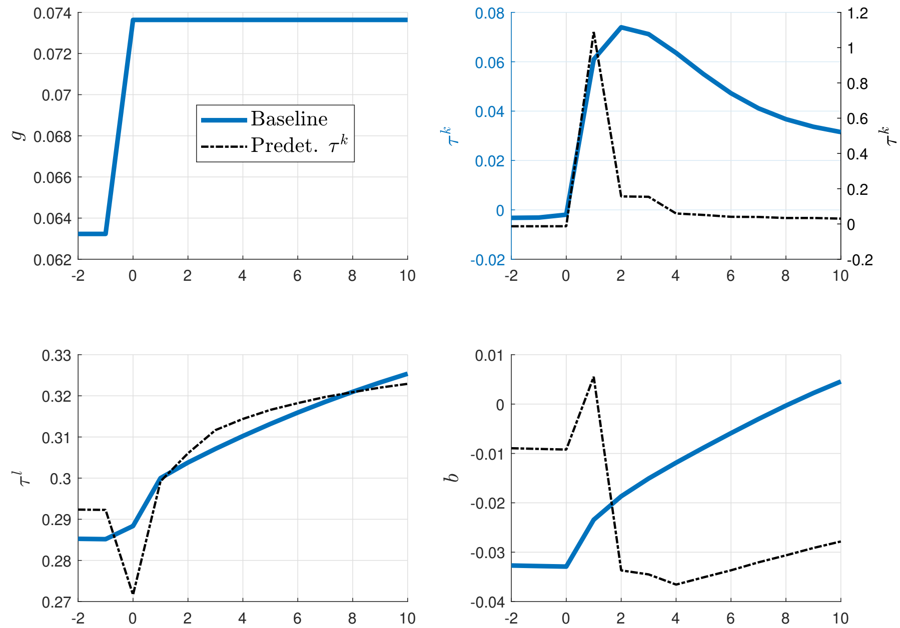

Notes: The figure displays the dynamics of fiscal variables around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. Top left: government spending g_t; top right: capital income tax rate τ_t^k; bottom left: labor income tax rate τ_t^l; bottom right: government debt. Solid line: baseline model with costly state contingency ($\gamma^k = \gamma^l = 5$); dashed-dotted line: predetermined capital tax ($\gamma^k = \infty, \gamma^l = 0$).

5.6 Comparison to Empirical Conditional Dynamics

In order to compare the conditional dynamics of our model with the data, we compute impulse responses of key variables to military spending shocks in US data, following the local-projection approach of Jordà (2005) as applied by Ramey and Zubairy (2018). They use a narrative approach to identify exogenous changes in government spending, measured as news about changes in military spending, which are likely to be exogenous with respect to the state of the economy.

We analyze the empirical response of capital and labor taxes, as well as government debt. By regressing changes in these variables on the military news shock, we uncover the impulse response of the variables to exogenous changes in government spending in the data. We provide a brief summary of the results here, with further details provided in Appendix D. To maximize power, given the demands of a full impulse response estimation using instrumented government spending, we extend our sample as much as possible using data from 1929 to 2015, which includes the large increases in spending seen during World War II and the 1950s Korean War. For this reason, it is important to caveat that the estimation strategy in this literature relies on early military buildups that precede our calibration sample and are thus significantly larger than the shocks in our model. Nevertheless, this analysis provides evidence on the qualitative response of different fiscal instruments to an exogenous increase in government spending. Moreover, in Appendix C.2 we consider larger and less persistent shocks in our model to verify that its main predictions are robust to this modification.

Figure 6: Empirical response to a military news shock

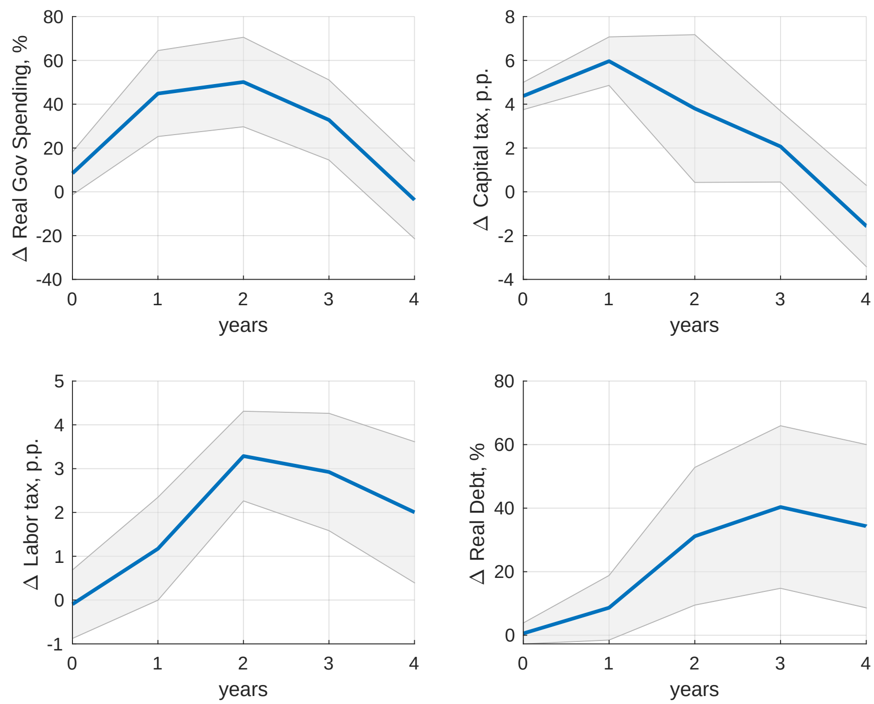

Notes: The figure displays the results of a local projection estimation on yearly US data from 1929 to 2015, giving the impulse responses to a military spending news shock. Solid lines give point estimates and the ranges are 95% confidence intervals.

We plot the results of this exercise in Figure 6. The size of the defense shock is normalized to create a 50% peak increase in total government spending, similar to the increase around the Korean War, and the impulse responses presented in Burnside, Eichenbaum, and Fisher (2004). The left panel shows that following an increase in military spending, total government spending increases for around four years. The remaining panels show how this is funded. Capital taxes tend to increase immediately and persistently and labor taxes increase over the following two years. These results are consistent with the findings of Burnside, Eichenbaum, and Fisher (2004) despite the different sample and estimation approach. Government debt also gradually increases, showing that not all of the increased spending is financed with taxes on average.

Comparing these empirical results with the conditional dynamics of our model in Figure 5 shows that costly state contingency appears to bring the qualitative dynamics of policy closer to the data. In particular, under costly state contingency capital and labor taxes and debt all gradually increase following a government spending shock, with capital taxes increasing faster than labor taxes. The model with predetermined capital taxes instead generates only a transitory spike in capital taxes, an initial decreasing in labor taxes, and government debt which falls (after a one period increase) in response to an increase in government spending.

This evidence suggests that restrictions on when and how governments can adjust their tax policies, such as costly state contingency, appear to be a relevant tool to improve the empirical fit of models of optimal taxation.

5.7 Total Factor Productivity Shocks

We now enrich our framework to analyze the role of business-cycle fluctuations in output for the volatility of taxes. To this end, we introduce aggregate productivity shocks beside the calibrated government spending shocks in our model with a government balanced budget. We assume that aggregate productivity z follows an AR(1) process in logs, $\log z_t = (1 - \rho_z)\mu_z + \rho_z \log z_{t-1} + u_t$, and calibrate the autocorrelation coefficient $\rho_z = .909$ and the standard deviation of innovations $\sigma_u = .014$. These values are in the standard range considered in the real-business-cycles literature to match the persistence and fluctuation of measured Total Factor Productivity in US data (e.g., Khan and Thomas, 2013).[17]

When we keep the costs of state contingency equal to our baseline value $\gamma = 39$, the volatility of the capital tax rate increases from 2% to approximately 4%, because the productivity shocks lead to an additional need to change tax policy over time in response to changes in the tax base. Accordingly, matching the empirical value of 2% requires us to increase the costs of state contingency moderately, to $\gamma = 53$. With this value, the model is broadly consistent with the standard deviation of both tax instruments as well as the autocorrelation of taxes. For comparison, we also introduce the productivity shocks in the model with predetermined capital tax. In this case, the volatility of the capital tax is

[17] We also set $\mu_z = .357$, consistent with our baseline normalizations.

significantly higher, at 12%, showing that TFP shocks further widen the distance between model and data in the absence of costly state contingency.

Table 3: Government Spending and TFP Shocks: Second Moments

Moment	$\gamma = 53$	$\gamma = 39$	Pred. τ^k
St. Dev. $\log(1+\tau^k)$	0.02	0.041	0.126
St. Dev. $\log(1+\tau^l)$	0.022	0.022	0.031
Autocorr. $\log(1+\tau^k)$	0.933	0.738	0.873
Autocorr. $\log(1+\tau^l)$	0.973	0.946	0.93

Notes: The table reports second moments of the tax rates on capital and on labor income in a model with both government spending shocks and productivity shocks, balanced budgets, and Full Commitment. The first column refers to the calibration with $\gamma^k = \gamma^l = 53$; the second column refers to the calibration with $\gamma^k = \gamma^l = 39$; the third column refers to the model with predetermined capital tax ($\gamma^k = \infty$,$\gamma^l = 0$).

5.8 Asymmetric Costs of State Contingency

Our baseline specification for the costs of state contingency is a symmetric quadratic function. We now explore the role of asymmetries in costly state contingency. Specifically, we consider the case in which increasing tax rates relative to fiscal announcements is more costly than reducing them. To this end, we modify our functional form assumption as follows:

$$\Gamma(\tau_t^j, \overline{\tau}_{t-1}^j) = \begin{cases} \frac{\kappa\gamma}{2}(\tau^j - \overline{\tau}^j)^2 & \text{if} \quad \tau^j \geq \overline{\tau}^j \\ \frac{\gamma}{2}(\tau^j - \overline{\tau}^j)^2 & \text{if} \quad \tau^j < \overline{\tau}^j, \end{cases} \tag{47}$$

with $\kappa \geq 1$ denoting the relative cost of tax hikes, thus nesting our baseline model ($\kappa = 1$). In this case, the optimal fiscal announcements can be expressed as follows:

$$\overline{\tau}_t^j = \mathbb{E}_t \tau_{t+1}^j + (\kappa - 1)\, \mathbb{E}_t \mathcal{I}\left(\tau_{t+1}^j \geq \overline{\tau}_t^j\right)\left(\tau_{t+1}^j - \overline{\tau}_t^j\right). \tag{48}$$

The second term on the right-hand side represents a wedge between announcement and expected tax that arises because of precautionary behavior.

We solve this version of the model setting $\gamma = 39$ as in our baseline calibration and $\kappa = 2$. We report the results in Table 4. We find that precautionary announcements lead to a positive average capital tax (in this case approximately 2%), because the government wants to avoid large positive spikes. Moreover, we verify that ex post, in response to government spending shocks, both capital and labor taxes display slightly smaller (absolute) deviations from the announcements when government spending increases, consistent with the asymmetry in the cost of state contingency.

Table 4: Asymmetric Cost of State Contingency: First and Second Moments

Moment	Baseline ($\kappa = 1$)	Asymmetric ($\kappa = 2$)
$\mathrm{E}\tau^k$	0.001	0.018
$\mathrm{E}\tau^l$	0.316	0.314
St. Dev. $\log(1 + \tau^k)$	0.021	0.013
St. Dev. $\log(1 + \tau^l)$	0.018	0.021
Autocorr. $\log(1 + \tau^k))$	0.826	0.905
Autocorr. $\log(1 + \tau^l))$	0.969	0.981

Notes: The table reports first and second moments of the tax rates on capital and on labor income in a model with asymmetric state contingency costs ($\kappa = 2$), balanced budgets, and Full Commitment. $\gamma^k = \gamma^l = 39$ in both models.

6 Role of Commitment Frictions

We now discuss our quantitative results for the case in which the government lacks commitment to state-contingent policies and makes one-period ahead noncontingent announcements, as described in Section 4.2.

6.1 Dynamics of Taxes with Limited-Time Commitment

The key difference relative to the analysis in the previous section is that costly state contingency now generates an endogenous degree of partial commitment to fiscal promises, because governments face a cost of deviating from previous noncontingent fiscal announcements. Governments thus optimally choose to what extent they implement the past promise, and to what extent they instead follow their current unconstrained optimal policy, depending on the state of the economy.

In the absence of costly state contingency ($\gamma = 0$), our model nests the case of No Commitment, for which Martin (2010) shows that the temptation for each government to raise capital taxes is so powerful that the model may display no equilibrium. In contrast, for our previously calibrated value of costly state contingency ($\gamma = 39$, which induces empirically consistent second moments for tax rates) we find that not only is the equilibrium interior, but governments tend to implement policies that are relatively similar to the ones that arise under Full Commitment, because costly state contingency endogenously generates a significant degree of commitment to past promises.

The moments generated by this version of the model are reported in Table 5. The average capital-income tax rate is approximately equal to 8%, which is higher than under Full Commitment (approximately zero) because of time inconsistency, but not as high as in models without any commitment, where it is typically higher than the labor tax (Klein and Ríos-Rull, 2003).

Table 5: Limited-Time Commitment: First and Second Moments

Moment	LTC ($\gamma = 39$)	LTC ($\gamma = 2$)	FC ($\gamma = 39$)	Data
$E\tau^k$	0.079	0.357	0.001	0.355
$E\tau^l$	0.301	0.243	0.313	0.226
St. Dev. $\log(1+\tau^k)$	0.021	0.007	0.02	0.022
St. Dev. $\log(1+\tau^l)$	0.017	0.021	0.018	0.015
Autocorr. $\log(1+\tau^k)$	0.988	0.435	0.83	0.868
Autocorr. $\log(1+\tau^l)$	0.971	0.971	0.972	0.876

Notes: The table reports first and second moments of the tax rates on capital and on labor income in the balanced-budget models with Limited-Time Commitment, Full Commitment, and in the data.

Furthermore, the dynamic response of the economy to government spending shocks are similar to the ones associated with Full Commitment. In particular, the standard deviations of capital and labor taxes are almost identical, despite the higher average capital tax and lower labor tax. A noticeable difference is that under Limited-Time Commitment capital taxes are more persistent than under Full Commitment, inheriting the high persistence of the spending shock. This is because the Limited-Time Commitment model has effectively fewer state variables and thus cannot generate the more complex transition path for capital taxes seen under Full Commitment. In Figure 1, we see that capital taxes overshoot and then recover under Full Commitment, which is due to the effect of the Lagrange multiplier on the implementability constraint. The conditional dynamics of taxes and allocations under Limited-Time Commitment are displayed in Figures 7 and 8 respectively, where we see that capital taxes instead gradually transition post shock without any overshooting.

In Table C1 in Appendix C.2 we complement this analysis with the first and second

Figure 7: Limited-Time Commitment: Taxes

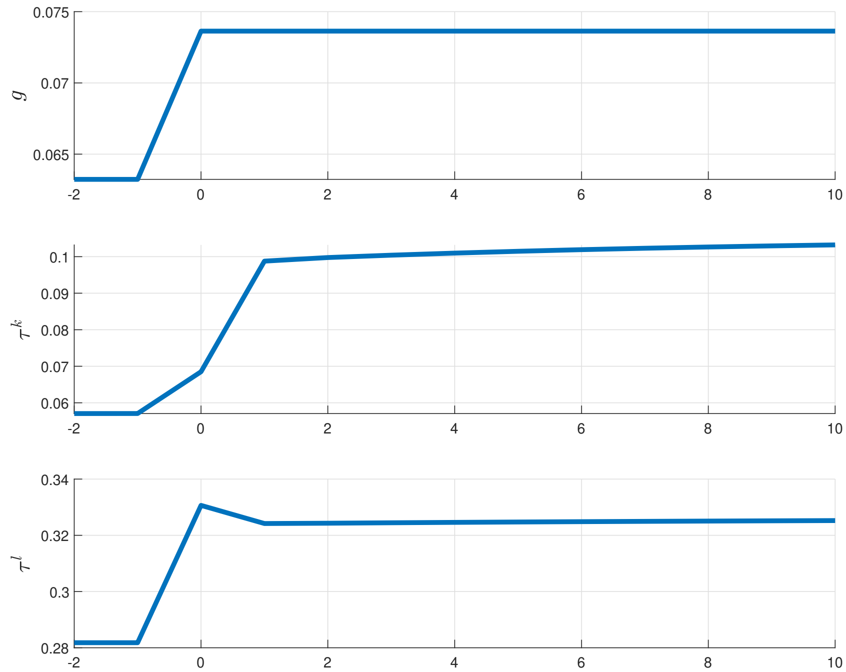

Notes: The figure displays the dynamics of fiscal variables around a shock that increases government spending, at $t = 0$, under Limited-Time Commitment with costly state contingency. Horizontal axes report time t. Top: government spending g_t; middle: capital income tax rate τ_t^k; bottom: labor income tax rate τ_t^l.

moments of allocations, comparing the Full-Commitment and Limited-Time Commitment regimes. Consistent with results for tax rates, we find that costly state contingency induces a significant degree of commitment; as a result, output and consumption are on average only marginally smaller in the presence of partial commitment.

In Figure 9, we compare realized taxes with noncontingent announcements in the Limited-Time Commitment 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. We also find small deviations in periods in which the level of government spending stays constant. These differences arise because of the strategic bias terms in equations (43) and (44): In formulating announcements, the government trades off a forecast of future taxes with an incentive to strategically manipulate future policies.

Next, we discuss two counterfactual analyses that highlight the role of costly state contingency in building partial fiscal commitment.

Figure 8: Limited-Time Commitment: Allocations

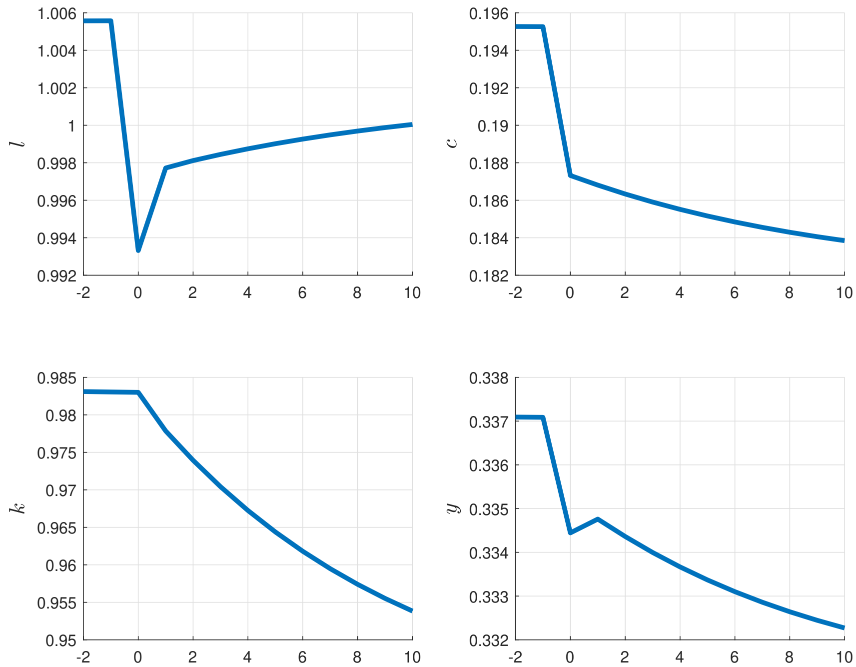

Notes: The figure displays the dynamics of allocations around a shock that increases government spending, at $t = 0$, under Limited-Time Commitment with costly state contingency. Horizontal axes report time t. From top to bottom: labor l_t; consumption c_t; capital k_t; output y_t.

6.2 Low Cost of State Contingency and Average Capital Tax

To further analyze the role of costly state contingency in building commitment, we now perform a counterfactual experiment. Specifically, we significantly reduce the level of the cost of state contingency relative to our baseline parameterization. We set $\gamma^k = \gamma^l = 2$.

We find that in this counterfactual economy with low costs of state contingency, the capital income tax is on average significantly higher—approximately 36%, and thus close to its empirical counterpart. This result is consistent with the fact that a lower cost of deviating from fiscal promises reduces the endogenous degree of commitment.

Furthermore, we find that the bias between fiscal announcements and realizations becomes substantially larger compared to our baseline Limited-Time Commitment economy. In Figure 10, we display the path of realized and preannounced tax rates. On average, governments promise relatively low taxes on capital; ex post, there is a positive deviation, toward higher capital taxation, because governments do not fully internalize the distor-

Figure 9: Limited-Time Commitment: Fiscal Announcements

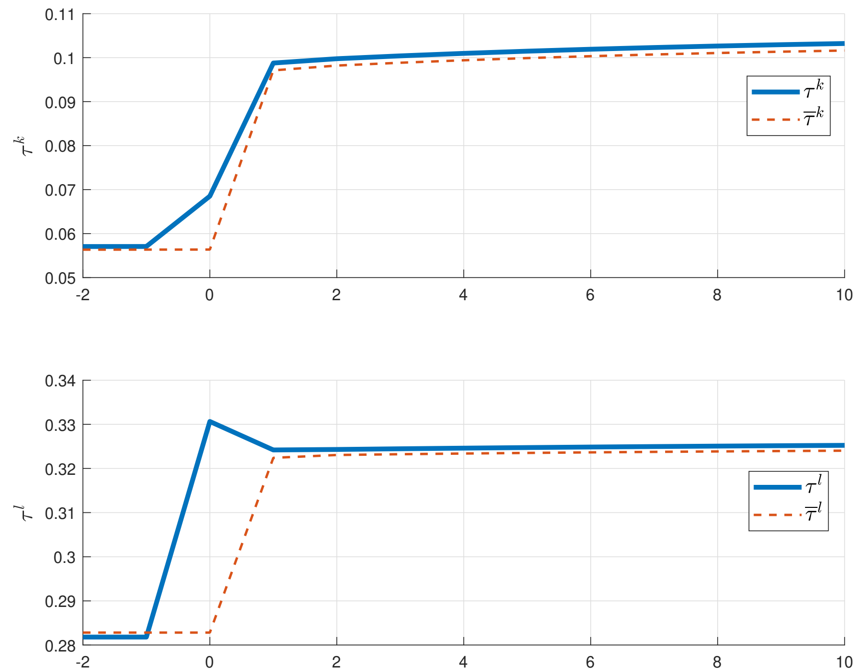

Notes: The figure compares the dynamics of realized tax rates (solid line) and announced tax rates (dashed line) around a shock that increases government spending, at $t = 0$, under Limited-Time Commitment with costly state contingency. Horizontal axes report time t. Top: capital income tax rate (τ_t^k and $\overline{\tau}_t^k$); bottom: labor income tax rate (τ_t^l and $\overline{\tau}_t^l$).

tionary effect of capital taxes on past investment. These dynamics are consistent with the analytical insights from our two-period model that we discuss in Section 4.2.1 and in more detail in Appendix B.

To evaluate the welfare cost of commitment frictions in the presence of costly state contingency, we compute the welfare cost of going from Full Commitment to Limited-Time Commitment given our calibrated cost $\gamma = 39$. Consistent with the fact that costly state contingency generates a substantial degree of fiscal commitment, this welfare cost is relatively small, and approximately equal to 1% of permanent consumption. Welfare losses rise as the cost of state contingency is lowered and capital taxes rise: The LTC model with $\gamma = 2$ leads to a welfare loss of 2% of permanent consumption relative to the LTC model with $\gamma = 39$.

Figure 10: Limited-Time Commitment with Low γ: Taxes and Announcements

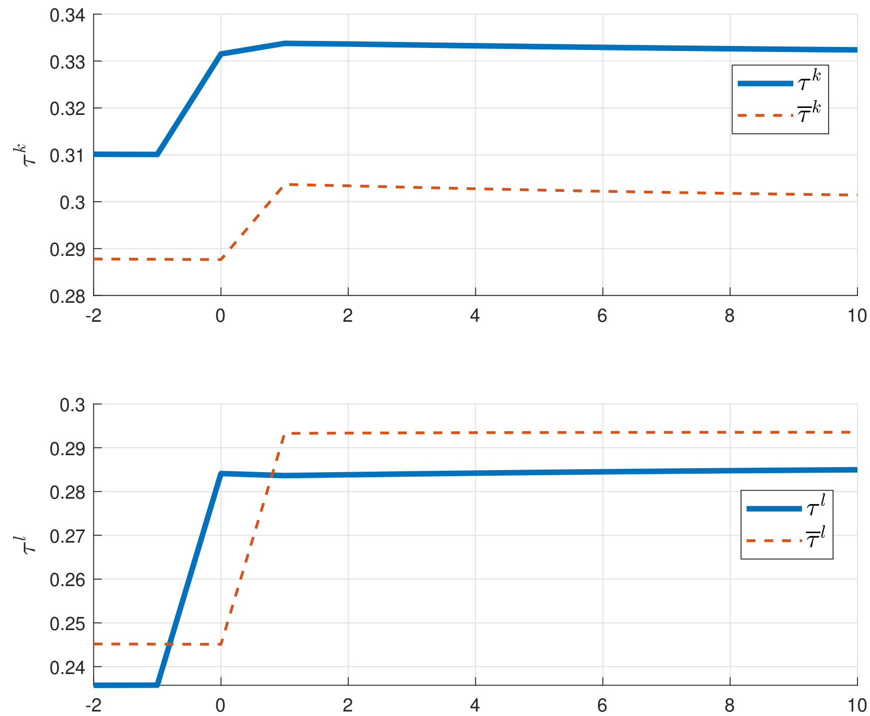

Notes: The figure compares the dynamics of realized tax rates (solid line) and announced tax rates (dashed line) around a shock that increases government spending, at $t = 0$, under Limited-Time Commitment with low costly state contingency, i.e., $\gamma = 2$. Horizontal axes report time t. Top: capital income tax rate (τ_t^k and $\overline{\tau}_t^k$); bottom: labor income tax rate (τ_t^l and $\overline{\tau}_t^l$).

6.3 Role of Uncertainty

In the Limited-Time Commitment economy, costs of state contingency play a dual role. First, they penalize the contemporaneous response of tax rates to shocks, as in the Full-Commitment case. Second, they build commitment by punishing deviations from fiscal promises.

The commitment-building role is clearly present even in the absence of any shocks. Thus, to decompose the two roles of costly state contingency with Limited-Time Commitment, we also consider a deterministic economy subject to the same costs of adjusting tax rates relative to previous promises, but with constant government spending. In so doing, we separately identify the effects of costly state contingency on the average level of allocations through their commitment-building role from its effects on the stochastic behavior of taxes and allocations.

We find that the allocation that we obtain in the deterministic economy is remarkably similar to the *average* allocation in our stochastic Limited-Time Commitment economy. For instance, consumption approximately equals 0.189 in the deterministic economy, the same value as the (approximate) mean of the stochastic economy, and only marginally lower than under Full Commitment.

This comparison suggests that, under Limited-Time Commitment, a crucial role of costly state contingency is to build an endogenous level of fiscal commitment that sustains an allocation with a high average level of output and consumption. At the same time, as we have seen, in the stochastic economy costly state contingency generates a similar degree of volatility of tax rates as under Full Commitment.

7 Conclusion

In this paper, we have explored the role of costly state contingency of 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. A key feature of our framework, which we believe captures a salient feature of reality, is that pre-announced fiscal plans are a state variable for the government, but they only partially constrain policy decisions. Importantly, announcements about future fiscal policy are used to give incentives to the private sector—for example, in order to generate changes in the current tax base—when changing current taxes is costly.

Under Full Commitment, when costs of state contingency apply to both capital and labor taxes symmetrically, they reduce the volatility of capital taxes in response to shocks, thus significantly improving the quantitative performance of models of optimal fiscal policy. 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. In a calibrated framework with taxes on capital and labor income as well as noncontingent government debt, our model successfully accounts for the empirical dynamics of all fiscal variables in response to an increase in government spending.

When the government lacks Full Commitment, fiscal announcements play a strategic role and allow the current government to affect future policies, by partially constraining future governments. As a consequence, we find that optimal announcements are biased forecasts of future policies and governments deviate from them in a systematic fashion by raising positive taxes on capital income.

Overall, by improving our understanding of the role of frictions that governments face in responding to shocks, this paper provides a step forward in the quantitative application of models of optimal fiscal policy.

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.

Bhandari, A., D. Evans, M. Golosov, and T.Sargent (2017): "Fiscal Policy and Debt Management with Incomplete Markets," *The Quarterly Journal of Economics*, 132(2), 617–663.

Bhandari, A., and E. R. McGrattan (2021): "Sweat Equity in U.S. Private Business," *The Quarterly Journal of Economics*, 136(2), 727–781.

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.

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.

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

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 (2021): "Optimal Policy with General Signal Extraction," *Journal of Monetary Economics*, 118, 54–86.

Jones, J. B. (2002): "Has Fiscal Policy Helped Stabilize the Postwar U.S. Economy?," *Journal of Monetary Economics*, 49(4), 709–746.

Jordà, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *The American Economic Review*, 95(1), 161–182.

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.

Khan, A., and J. K. Thomas (2013): "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," *Journal of Political Economy*, 121(6), 1055–1107.

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.

Krusell, P., F. M. Martin, and J.-V. Ríos-Rull (2004): "Time-Consistent Debt," .

Leeper, E. M., M. Plante, and N. Traum (2010): "Dynamics of Fiscal Financing in the United States," *Journal of Econometrics*, 156(2), 304–321.

Maliar, L., and S. Maliar (2003): "Parameterized Expectations Algorithm and the Moving Bounds," *Journal of Business and Economic Statistics*, 21(1).

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

Mendoza, E. G., A. Razin, and L. L. Tesar (1994): "Effective Tax Rates in Macroeconomics: Cross-country Estimates of Tax Rates on Factor Incomes and Consumption," *Journal of Monetary Economics*, 34(3), 297–323.

Mertens, K., and M. Ravn (2011): "Understanding the Effects of Anticipated and Unanticipated Tax Policy Shocks," *Review of Economic Dynamics*, 14(1), 27–54.

——— (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.

Mertens, K., and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *The American Economic Review*, 103(4), 1212–1247.

Ortigueira, S., and J. Pereira (2021): "Lack of Commitment, Retroactive Taxation, and Macroeconomic Instability," *Journal of the European Economic Association*, Forthcoming.

Ramey, V. A., and S. Zubairy (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126(2), 850–901.

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 (2023): "A Machine Learning Projection Method for Macro-Finance Models," Discussion paper.

APPENDICES

A Model Appendix

A.1 Optimality Conditions with Balanced Budget and Full Commitment

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

$$\lambda_t = u_c(c_t) - \mu_t u_{cc}(c_t)k_t + \mu_{t-1}\left[u_{cc}(c_t)(c_t+k_t) + u_c(c_t)\right] - \nu_t^c, \tag{A1}$$

$$v_l(l_t) = \lambda_t F_l(k_{t-1}, l_t) - \mu_{t-1}\left[v_{ll}(l_t)l_t + v_l(l_t)\right] + \nu_t^l. \tag{A2}$$

The first-order condition with respect to capital k_t is

$$\lambda_t = \beta \mathbb{E}_t \lambda_{t+1}\left[F_k(k_t, l_{t+1}) + 1 - \delta\right] - \mu_t u_c(c_t) + \mu_{t-1} u_c(c_t) - \beta \mathbb{E}_t\left[\nu_{t+1}^l h_k^l(k_t, g_{t+1}, \tau_{t+1}^k, \tau_{t+1}^l) - \nu_{t+1}^c h_k^c(k_t, g_{t+1}, \tau_{t+1}^k, \tau_{t+1}^l)\right]. \tag{A3}$$

Equations (A1), (A2), and (A3) 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. These act as wedges which drive the allocation away from the benchmark without costs of state contingency.

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

$$\nu_t^c h_{\tau^k}^c(k_{t-1}, g_t, \tau_t^k, \tau_t^l) - \nu_t^l h_{\tau^k}^l(k_{t-1}, g_t, \tau_t^k, \tau_t^l) = \Gamma_{\tau^k}^k(\tau_t^k, \overline{\tau}_{t-1}^k), \tag{A4}$$

$$\nu_t^c h_{\tau^l}^c(k_{t-1}, g_t, \tau_t^k, \tau_t^l) - \nu_t^l h_{\tau^l}^l(k_{t-1}, g_t, \tau_t^k, \tau_t^l) = \Gamma_{\tau^l}^l(\tau_t^l, \overline{\tau}_{t-1}^l). \tag{A5}$$

Equations (A4) and (A5) 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.[18]

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

$$\mathbb{E}_t \Gamma_{\overline{\tau}^k}^k(\tau_{t+1}^k, \overline{\tau}_t^k) = 0, \tag{A6}$$

$$\mathbb{E}_t \Gamma_{\overline{\tau}^l}^l(\tau_{t+1}^l, \overline{\tau}_t^l) = 0. \tag{A7}$$

[18]Notice that when costs of state contingency are removed, we have $\nu_t^c = \nu_t^l = 0$, and the FOCs reduce to those in Stockman (2001).

A.2 Discussion: Costly State Contingency of Taxes and 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

$$b_t(g^t) = \tau_t^k r_t k_{t-1} + \tau_t^l w_t l_t - g_t + \sum_{g^{t+1}} q_t(g^{t+1}|g^t) b_{t+1}(g^{t+1}), \tag{A8}$$

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 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 (A8) by recursively substituting out state-contingent debt, we obtain a single implementability constraint:

$$u_c(c_0)\left[b_{-1} + k_{-1} + \left(F_k(k_{-1}, l_0) - \delta\right)(1 - \tau_0^k) k_{-1}\right] = \mathbb{E}_0 \sum_{t=0}^{\infty} \left(u_c(c_t) c_t - v_l(l_t) l_t\right) \tag{A9}$$

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

Consider first the case of no costs of state contingency, i.e., $\gamma^j = 0$ 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 $\gamma^j > 0$ 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 $\gamma^j = 0$ imply variation in the capital tax rate across states. Thus, they are no longer optimal when $\gamma^j > 0$, 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 indeterminacy between the role of debt and capital taxes in absorbing fiscal shocks.

B 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. As in the main text, we distinguish between the case of Full Commitment and the case of Limited-Time Commitment.

B.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 at $t = 1$ with no subscripts, and we index $t = 0$ variables with subscript 0.

A representative household has utility function

$$c_0 + \beta \mathbb{E}\left(\log(c) - \chi \frac{l^{1+\eta}}{1+\eta}\right), \tag{B1}$$

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

$$c_0 + k = y_0, \tag{B2}$$

$$c + g = z k^{\alpha} l^{1-\alpha}, \tag{B3}$$

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 one 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

$$\left(\alpha\tau^k + (1-\alpha)\tau^l\right) zk^\alpha l^{1-\alpha} \geq g, \tag{B4}$$

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 (B4) 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

$$\chi l^\eta c = (1-\alpha)zk^\alpha l^{-\alpha}\left(1-\tau^l\right), \tag{B5}$$

$$1 = \beta\mathbb{E}\left[c^{-1}\left(1-\tau^k\right)\alpha z k^{\alpha-1} l^{1-\alpha}\right]. \tag{B6}$$

Equation (B6) is the standard Euler equation for capital, which implies the usual time-inconsistency for capital taxation. In particular, time-1 capital taxes, τ^k, appear in the Euler equation, which constrains the government at time 0. We can combine equations (B5) and (B6) with the government budget constraint (B4) to derive the following two implementability constraints. Firstly, a labor supply optimality condition,

$$\chi l^{\eta+\alpha}\left(zk^\alpha l^{1-\alpha} - g\right) = (1-\alpha)zk^\alpha\left(1-\tau^l\right), \tag{B7}$$

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

$$k \leq \beta\left[1 - \chi\mathbb{E}\left(h(k,g,\tau^l)\right)^{1+\eta}\right], \tag{B8}$$

which we express as an inequality because we allow the government to pay a non-negative lump-sum transfer. Given a choice of labor tax τ^l for each realization of g, private-sector allocations must satisfy constraints (B7) and (B8), and the associated capital tax can be then obtained using (B4).

We maintain two assumptions on the environment, which are needed for the government's problem to be well behaved and to have a sensible interpretation.

Firstly, we focus on parameter configurations such that a weak condition is satisfied, namely that for given k and g, if the government raises labor taxes then the required capital tax to balance the budget decreases. This gives a natural sense in which the government in the second period must choose between either high labor taxes and low capital taxes, or vice versa. Specifically, define $\tau^k = h^{\tau^k}(k, g, \tau^l)$ as the required capital tax to balance the budget in (B4). We therefore consider parameters such that $h^{\tau^k}_{\tau^l} < 0$. Implicitly differentiating (B4) shows that this amounts to assuming that $h_{\tau^l}(\alpha\tau^k + (1-\alpha)\tau^l) > -l$, so that the negative labor supply effect of raising labor taxes does not outweigh the direct positive effect of labor taxes on the budget.

Secondly, we make assumptions so that the government's problem at time 1 in the LTC game is strictly concave in taxes. Specifically, when considering the special case of $\gamma^l > 0, \gamma^k = 0$, we require that the indirect utility function obtained from second-period utility $\log(c) - \chi\frac{l^{1+\eta}}{1+\eta}$ (with all equilibrium conditions plugged in) is concave in τ^l. When considering the special case of $\gamma^l > 0, \gamma^k = 0$, we require that the indirect utility function is concave in τ^k. This assumption ensures that the maximization problem features an interior solution when the maximization is done over taxes, and taking into account the additional (concave) state contingency costs. We discuss these conditions in detail during the proof.

B.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. 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

$$c_0 + \beta\mathbb{E}\left[\log(c) - \chi\frac{l^{1+\eta}}{1+\eta} - \frac{\gamma^k}{2}(\tau^k - \overline{\tau}^k)^2 - \frac{\gamma^l}{2}(\tau^l - \overline{\tau}^l)^2\right], \tag{B9}$$

where $\gamma^k \geq 0$ and $\gamma^l \geq 0$ are parameters that determine the costs of state contingency in taxes. We assume that these costs are quadratic functions of the distance between realized tax rates and noncontingent announcements.

The government maximization problem is subject to the resource constraints and the implementability constraints derived above. We denote by μ the multiplier on (B8), ν the multiplier on (B4), and directly substitute in $l = h(k, g, \tau^l)$ and the resource constraints, (B2) and (B3). 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

$$\nu \alpha z k^{\alpha} l^{1-\alpha} = \gamma^k (\tau^k - \overline{\tau}^k) \tag{B10}$$

$$\begin{aligned} \left[(1-\alpha) z k^{\alpha} l^{-\alpha} \left(c^{-1} + \nu(\alpha \tau^k + (1-\alpha)\tau^l)\right) - \chi l^{\eta}(1 + \mu(1+\eta))\right] h_{\tau^l}(k, g, \tau^l) \\ + \nu(1-\alpha) z k^{\alpha} l^{1-\alpha} = \gamma^l (\tau^l - \overline{\tau}^l), \end{aligned} \tag{B11}$$

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). 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 time 0.

The first-order conditions with respect to the optimal tax announcements are

$$\overline{\tau}^k = \mathbb{E}\tau^k, \tag{B12}$$

$$\overline{\tau}^l = \mathbb{E}\tau^l. \tag{B13}$$

Thus, 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.

We can gain further intuition on the role played by costly state contingency by taking expectations of the first order conditions (B10) and (B11) and using (B12) and (B13) to give

$$\mathbb{E} \nu \alpha z k^{\alpha} l^{1-\alpha} = 0 \tag{B14}$$

$$\begin{aligned} \mathbb{E}\Big\{ \left[(1-\alpha) z k^{\alpha} l^{-\alpha} \left(c^{-1} + \nu(\alpha \tau^k + (1-\alpha)\tau^l)\right) - \chi l^{\eta}(1 + \mu(1+\eta))\right] h_{\tau^l}(k, g, \tau^l) \\ + \nu(1-\alpha) z k^{\alpha} l^{1-\alpha} \Big\} = 0. \end{aligned} \tag{B15}$$

In a model with free state contingency, we would have $\gamma^k = \gamma^l = 0$ and the right-hand sides of (B10) and (B11) would equal 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.[19]

B.3 Optimal Policy with Limited-Time Commitment

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, a new government at $t = 1$ acts under discretion. This assumption implies that the 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

$$\log(c) - \chi\frac{l^{1+\eta}}{1+\eta} - \frac{\gamma^k}{2}(\tau^k - \overline{\tau}^k)^2 - \frac{\gamma^l}{2}(\tau^l - \overline{\tau}^l)^2, \tag{B16}$$

subject to the budget constraint (B4) with multiplier ν, the implementability constraint (B7), 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

$$\nu\alpha z k^\alpha l^{1-\alpha} = \gamma^k(\tau^k - \overline{\tau}^k), \tag{B17}$$

$$\begin{aligned}\left[(1-\alpha)zk^\alpha l^{-\alpha}\left(c^{-1} + \nu(\alpha\tau^k + (1-\alpha)\tau^l)\right) - \chi l^\eta\right] h_{\tau^l}(k, g, \tau^l)& \\ + \nu(1-\alpha)zk^\alpha l^{1-\alpha} = \gamma^l(\tau^l - \overline{\tau}^l),&\end{aligned} \tag{B18}$$

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 $\tau^l = \tilde{\tau}^l(k, g, \overline{\tau}^k, \overline{\tau}^l)$.

[19] 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.

The policy function for capital taxes is then given as the level of capital taxes required to balance the budget given the labor tax policy function. 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 $\nu = 0$. Notice therefore that, again, in the absence of state contingency costs for capital taxes ($\gamma^k = 0$) the optimal labor tax is noncontingent ($\tau^l = 0$); hence, costly state contingency for labor is irrelevant for the optimal solution.

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

$$\tilde{W}_{\overline{\tau}^k} = \gamma^k(\tau^k - \overline{\tau}^k), \tag{B19}$$

$$\tilde{W}_{\overline{\tau}^l} = \gamma^l(\tau^l - \overline{\tau}^l). \tag{B20}$$

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

$$c_0 + \beta \mathbb{E} \tilde{W}(k, g, \overline{\tau}^k, \overline{\tau}^l), \tag{B21}$$

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

$$k \leq \beta \left[1 - \chi \mathbb{E} \left(h(k, g, \tilde{\tau}^l) \right)^{1+\eta} \right], \tag{B22}$$

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:

$$\overline{\tau}^k = \mathbb{E}\tau^k - \frac{\chi(1+\eta)\mu}{\gamma^k} \mathbb{E} \left[l^\eta h_{\tau^l}(k, g, \tau^l) \tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) \right], \tag{B23}$$

$$\overline{\tau}^l = \mathbb{E}\tau^l - \frac{\chi(1+\eta)\mu}{\gamma^l} \mathbb{E} \left[l^\eta h_{\tau^l}(k, g, \tau^l) \tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) \right]. \tag{B24}$$

Comparing these to the optimal promises made under Full Commitment, (B12) and (B13), we see a crucial difference. Under Full Commitment, the government announces taxes equal to the average of ex-post realized taxes. In the Limited-Time Commitment 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 time inconsistency.

Specifically, the government at $t = 0$ understands that by marginally decreasing future labor it can relax its current implementability constraint (B22). Thus, intuitively, the biases in (B23) and (B24) 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 (B23): We have $\mu > 0$ and $h_{\tau^l}(k, g, \tau^l) < 0$, and so as long as raising the capital tax promise, $\bar{\tau}^k$ raises the implemented capital tax and lowers the implemented labor tax ($\tilde{\tau}^l_{\bar{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) < 0$) 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 (B24) for a symmetric reason.

In order to investigate the biases more formally, first note that the biases interact in non-trivial ways. For example, consider the limit of $\gamma^k \to \infty$ while holding γ^l finite. The second-period government is therefore forced to set $\tau^k = \overline{\tau}^k$ in all states to avoid the capital contingency cost. The government budget constraint then implies a certain value of the labor tax is required to balance the budget given the other state variables. This means that the labor tax is therefore completely pinned down by the capital tax promise, irrespective of the labour tax promise, and we would have $\tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) = 0$, and hence no bias in the labor tax promise.

Given this interaction, we focus on signing the bias for each tax while holding the cost of state contingency equal to zero for the other tax. In the following proposition, we prove that in this case the biases take the expected signs:

Proposition 1 *In the Limited-Time Commitment equilibrium, the time-0 government biases capital tax promises downwards and labor tax promises upwards as long as* $\tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) < 0$ *and* $\tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) > 0$ *respectively. Moreover, under our maintained assumptions, 1) holding* $\gamma^l = 0$ *it must be that* $\tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) < 0$*, and 2) holding* $\gamma^k = 0$ *it must be that* $\tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) > 0$.

Proof. Firstly, that the biases have the stated signs for given values of the derivatives $\tilde{\tau}^l_{\overline{\tau}^k}(k, g, \overline{\tau}^k, \overline{\tau}^l) < 0$ and $\tilde{\tau}^l_{\overline{\tau}^l}(k, g, \overline{\tau}^k, \overline{\tau}^l) > 0$ follows directly from (B23) and (B24).

Secondly, we must signs these derivatives in the stated cases where one cost is equal to zero. To do this, we rewrite the problem in an indirect utility form. Switching to a generic notation, consider the time-1 problem of choosing an instrument τ to maximize an indirect utility function $w(k, g, \tau)$ which is strictly concave in τ, subject to costs of state contingency $\frac{\gamma}{2}(\tau - \bar{\tau})^2$. The problem is therefore to solve:

$$\max_{\tau} w(k, g, \tau) - \frac{\gamma}{2}(\tau - \bar{\tau})^2. \tag{B25}$$

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

$$\tilde{\tau}(k, g, \bar{\tau}) = \bar{\tau} + \frac{1}{\gamma} w_{\tau}(k, g, \tilde{\tau}(k, g, \bar{\tau})). \tag{B26}$$

It is easy to see that the chosen policy is always increasing in the inherited promise. Implicitly differentiating (B26) to solve for the relevant derivative yields:

$$\tilde{\tau}_{\bar{\tau}}(k, \bar{\tau}, g) = \frac{1}{1 - \frac{1}{\gamma} w_{\tau,\tau}(.)}. \tag{B27}$$

Since the indirect utility function is strictly concave ($w_{\tau,\tau}(.) < 0$) by assumption, the denominator is positive, and raising the promise therefore raises the optimal policy.

The aforementioned special cases of $\gamma^l > 0, \gamma^k = 0$ and $\gamma^l = 0, \gamma^k > 0$ fit into the generic framework above, and hence the proof applies to those models. To see this, in both cases one of the costs is zero, so the problem reduces to choosing a single policy τ. The indirect utility functions in either case are strictly convex in each tax according to our maintained assumptions. The indirect utility functions are derived in the following section, along with a detailed discussion of the maintained assumptions required for strict concavity.

B.3.1 Assumptions to Ensure Concavity of Indirect Utility

In this section we derive the indirect utility functions used in the proof above, and discuss the conditions needed to ensure that the time-1 problem is concave in taxes.

First, consider the optimization over labor taxes only. In this case, we back out the required capital tax from the budget constraint, and analyze the effects of varying the level of τ^l on the equilibrium. For the indirect utility function to be concave in τ^l, a sufficient condition is that the function $l = h(k, g, \tau^l)$ is concave in labor taxes. This condition is satisfied in any equilibrium with positive consumption, and so requires no additional assumptions beyond the existence of an equilibrium at the given parameters. To see this, note that household utility is concave in l, once the resource constraint (B3) has been used to substitute out for c. Therefore it is also concave in τ^l as long as l is concave in τ^l.

Defining the indirect utility function in this case, we have:

$$w(k,g,\tau^l) \equiv \log\left(zk^\alpha h(k,g,\tau^l)^{1-\alpha} - g\right) - \chi \frac{h(k,g,\tau^l)^{1+\eta}}{1+\eta}. \tag{B28}$$

To see that $h(k,g,\tau^l)$ is concave in τ^l, implicitly differentiate (B5) to yield

$$h_{\tau^l}(k,g,\tau^l) = \frac{-(1-\alpha)zk^\alpha}{(\alpha+\eta)\chi h(k,g,\tau^l)^{\alpha+\eta-1}c + \frac{\partial c}{\partial l}\chi h(k,g,\tau^l)^{\alpha+\eta}}, \tag{B29}$$

where c comes from the resource constraint $c = zk^\alpha h(k,g,\tau^l)^{1-\alpha} - g$, and hence $\frac{\partial c}{\partial l} = (1-\alpha)zk^\alpha l^{-\alpha}$. Plugging this into (B29) yields

$$h_{\tau^l}(k,g,\tau^l) = \frac{-(1-\alpha)zk^\alpha}{(\alpha+\eta)\chi h(k,g,\tau^l)^{\alpha+\eta-1}c + (1-\alpha)zk^\alpha \chi h(k,g,\tau^l)^{\eta}}. \tag{B30}$$

Since this expression is negative, as long as the denominator is decreasing in τ^l, $h_{\tau^l}(k,g,\tau^l)$ is decreasing in τ^l and hence $h(k,g,\tau^l)$ is concave. This turns out to be the case. To see this, first note that the second term of the denominator, $(1-\alpha)zk^\alpha\chi h(k,g,\tau^l)^\eta$ is decreasing in τ^l. The first term, $(\alpha+\eta)\chi h(k,g,\tau^l)^{\alpha+\eta-1}c$ initially appears to be ambiguous but is also decrease in τ^l. Firstly note that if $\alpha+\eta-1 \geq 0$ then this term is decreasing in τ^l because both c and $h(k,g,\tau^l)$ are decreasing in τ^l. Secondly, if $\alpha+\eta-1 < 0$ then replace c using the resource constraint to yield $(\alpha+\eta)\chi h(k,g,\tau^l)^{\alpha+\eta-1}c = (\alpha+\eta)\chi h(k,g,\tau^l)^{\alpha+\eta-1}(zk^\alpha h(k,g,\tau^l)^{1-\alpha} - g) = (\alpha+\eta)\chi\left(zk^\alpha h(k,g,\tau^l)^\eta - gh(k,g,\tau^l)^{\alpha+\eta-1}\right)$. Since $\alpha+\eta-1 < 0$ the $-gh(k,g,\tau^l)^{\alpha+\eta-1}$ term is decreasing in τ^l, and hence so is the whole term. Putting this together, the denominator is decreasing in τ^l and hence $h(k,g,\tau^l)$ is concave.

Next, consider the optimization over capital taxes only. In this case, we back out the required labor tax from the budget constraint, and analyze the effect on equilibrium of varying the level of τ^k. As with labor taxes, concavity of the indirect utility in τ^k requires that labor is concave in τ^k. To investigate this, we first need to define a function giving equilibrium labor for any value of the capital tax, which we call $l = \hat{h}(k,g,\tau^k)$. This function is found by combining the labor supply condition (B5) with the government budget constraint (B3). This yields the equation

$$\tau^k = \frac{1}{\alpha}\left(\alpha - 1 + \frac{g}{zk^\alpha l^{1-\alpha}} + \chi l^{1+\eta}\left(1 - \frac{g}{zk^\alpha l^{1-\alpha}}\right)\right). \tag{B31}$$

Note that the term $\left(1 - \frac{g}{zk^\alpha l^{1-\alpha}}\right)$ must be positive in any equilibrium with positive consumption. This equation reveals that the relationship between labor and the capital tax contains both concave and convex components. If τ^k is convex in l then l is concave in τ^k. But τ^k is composed of both convex functions of l ($\frac{g}{zk^\alpha l^{1-\alpha}}$ and $\chi l^{1+\eta}$) and concave ones ($-\frac{g}{zk^\alpha l^{1-\alpha}}$). Hence for l to be concave in τ^k we need the convex components to dominate. Intuitively, this is true as long as government spending is a sufficiently small share of output, so that the concave term $-\frac{g}{zk^\alpha l^{1-\alpha}}$ is dominated by the convex terms. For example, in the limit of government spending going to zero we can solve for labor explicitly as $l = (\alpha\tau^k + 1 - \alpha)^{\frac{1}{1+\eta}}$, which gives that labor is concave in the capital tax, as required. For the purposes of the proposition, we simply maintain that g is small enough that $l = \hat{h}(k, g, \tau^k)$ is concave in τ^k.

With this function in hand, the indirect utility function when considering τ^k as the choice variable is defined as

$$\hat{w}(k, g, \tau^k) \equiv \log\left(zk^\alpha h(k, g, \tau^k)^{1-\alpha} - g\right) - \chi \frac{h(k, g, \tau^k)^{1+\eta}}{1+\eta}. \tag{B32}$$

B.4 Numerical Example

To illustrate the results from our two-period model, we now provide a numerical example, where for simplicity we impose costs of state contingency only on capital taxes.

As the exercise is illustrative, we choose parameter values to deliver sensible results, but do not carry out a full calibration exercise as we did in our infinite horizon model. For the labor elasticity we choose $\eta = 2$ as in our full model. We set $\beta = 0.95$ and $\alpha = 1/3$, and normalize $z = 1$. We choose $\chi = 0.7208$ so that $l = 1$ in the first best solution. We choose an average value of government spending of 0.0525, which corresponds to around 7.5% of first-best output, and implies that capital in the absence of any costs of state contingency would equal around 75% of first-best capital. For our uncertainty process, we assume that government spending takes two equally likely values, low and high, equal to 25% below and above average respectively. We verify numerically that the indirect utility function is concave in taxes at these parameter values.

With parameter values in hand, we solve our model under the assumptions of Full Commitment and Limited-Time Commitment for a range of values of the cost of state contingency, γ, between 0 and 2. We plot our results for the tax policy functions in Figure B1.

Panel (a) plots the results when the government possesses Full Commitment. In all plots, the values where $\gamma = 0$ correspond to the Full Commitment solution with no costs of contingency. As can be seen in the left panel, the optimal solution in this case features constant labor taxes across the two states, and a higher capital tax (middle panel) when government spending is high. Moving to the right, we see what happens to the variables of interest when the cost of state contingency in capital taxes increases. Raising this cost naturally reduces the state contingency of capital taxes, and the values in the two government spending states become more similar, the higher the cost is.

Accordingly, the labor tax must become more state contingent to fund the variation in government spending. This is inefficient in the two-period model, because it discourages increased labor supply in response an increased government spending. The right panel shows the capital tax promise, $\bar{\tau}$, and the average capital tax across the two states. As we derived above, when the government has Full Commitment it simply sets the promise equal to the average of the policies it expects to set in order to minimize the cost function.

Panel (b) plots the results with Limited-Time Commitment. The results are starkly different, highlighting how the tradeoffs from reducing the ability of governments to set state contingent polices are very different when governments do not possess commitment.

The key difference between the results under Full Commitment and Limited-Time Commitment is that under Full Commitment raising γ tends to reduce the amount of state contingency while having small effects on averages, while under Limited-Time Commitment raising γ also affects average taxes and allocations, consistent with the findings of our infinite-horizon model.

We can see this very clearly in the behavior of capital taxes, shown in the middle plot. When $\gamma = 0$, we have high capital taxes under Limited-Time Commitment. The level of uncertainty in this example is high, but note that it is not high enough that the capital tax is negative in either state. Thus, capital taxes in both states are positive, in contrast to the Full Commitment solution, which has zero capital taxes on average, positive when realized spending is high and negative when low.

In the Limited-Time Commitment solution we see that there is a wedge between the promised and realized average capital tax, as shown in the right panel. For low costs of state contingency, this gap is large, and the time 0 government has to promise a very negative capital tax in order to try and reduce the average tax, which nonetheless remains high. As the cost rises, this pulls down the realized average capital tax, and brings the promise closer to the expected realization. Of course, this comes at the cost of lower state contingency, and for very high values of the cost the governments achieve low capital taxes

Figure B1: Two-Period Model Results: Full Commitment vs. Limited-Time Commitment

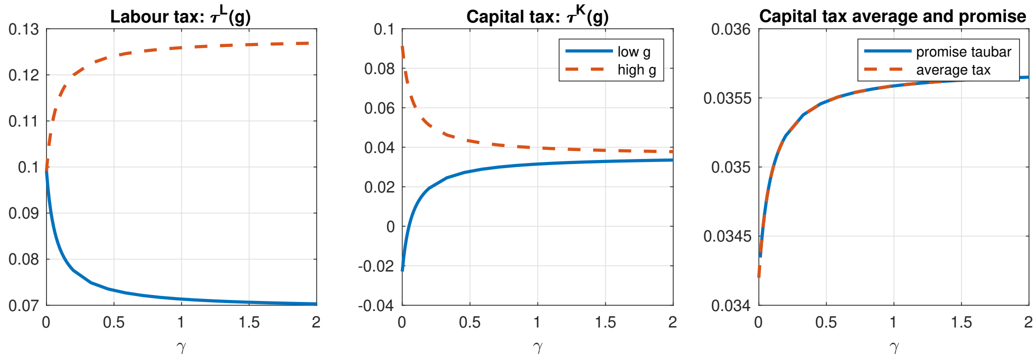

(a) Results under Full Commitment

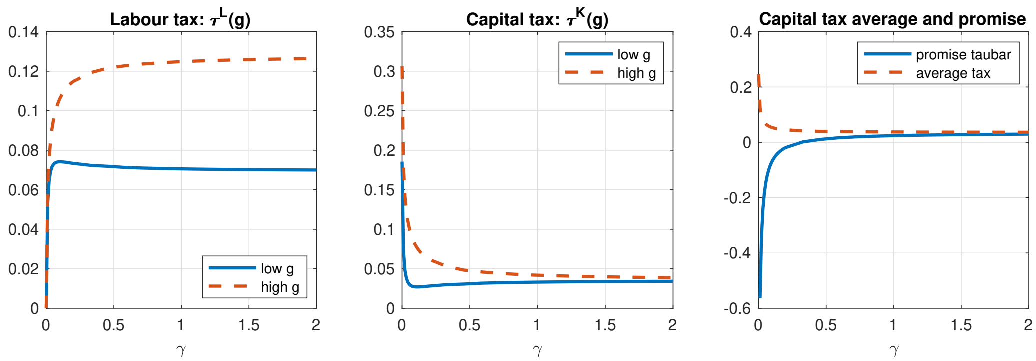

(b) Results under Limited-Time Commitment

The figure displays the values of policy functions across different values of the parameter γ, which determines the cost of state contingency of adjusting capital taxes relative their pre-announced level. Panel (a) gives results when the government acts with Full Commitment, and panel (b) Limited-Time Commitment.

on average, but at the cost of them being essentially unresponsive to shocks.

C Quantitative Appendix

C.1 Solution Method

We solve the model using a generalization of the Parameterized Expectations Algorithm (PEA) (den Haan and Marcet, 1990) that relies on neural networks instead of polynomials. The method and its advantages relative to standard PEA are described in detail in Valaitis and Villa (2023). In our model, this method is particularly suitable because of the large number of state variables and decision rules to approximate. Moreover, in the version of the model with debt, the method can accommodate the nonlinearities associated with the occasionally binding borrowing constraint.

We begin by describing the case of Full Commitment. The state variables of the government problem are $x^{FC} \equiv (k, g, \overline{\tau}^k, \overline{\tau}^l, \mu)$ under balanced budget and $(k, g, b, \overline{\tau}^k, \overline{\tau}^l, \mu, \nu)$ with noncontingent debt. 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 initially 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 to obtain a new iteration of our approximating neural network. We proceed up to convergence of our approximation.

In the version of the model with noncontingent debt, we use the method of Maliar and Maliar (2003) with gradually adjusted bounds for debt, starting from the solution of the special case with balanced budget.

In the case of the Limited-Time Commitment 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 compute the derivatives S_x for state variables x. We perform this step by numerically approximating these derivatives with finite differences.

C.2 Additional Numerical Results

Balanced Budget Models: In Figure C1 we show that in our model the volatility of capital tax is monotonically decreasing in the parameter γ. This moment identifies our calibration. We also highlight two further results. First, the volatility of the labor tax in our model is close to its empirical counterpart for a wide range of values for γ, including our calibration. Second, in a model with costs of state contingency only on the capital tax ($\gamma^l = 0$) it is not possible to match the volatility of the capital tax; this moment is robustly larger in the model than in the data. This result highlights the importance of assuming that costly state contingency applies to both taxes.

In Table C1 we report the first and second moments of allocations, comparing the Full-Commitment and Limited-Time Commitment regimes. We find that allocations are quite similar in all the scenarios we consider. Consistent with the discussion in Section 6, we find costly state contingency induces a significant degree of commitment; as a result, output and consumption are on average only marginally smaller in the presence of partial commitment. By comparing the baseline Full-Commitment model with its counterparts without costly state contingency on both tax instruments, we find that costly state contingency leads to a slightly higher unconditional volatility of consumption, consistent with our analysis of Section 5.2.

Model with Low Persistence of g: For robustness, we solve a model with lower persistence of g ($\rho_g = 0.89$) and higher volatility of innovations ($\sigma_g = 0.0319$), as in Farhi (2010). The moments for this version of the model are given in Table C2, and the path for taxes and allocations following a change in government spending in Figures C2 and C3 respectively. The mechanism that we describe in the main text is robust to this alternative parameterization and costly state contingency provides a comparable reduction in the volatility of the capital tax.

Figure C1: Identification of Cost of State Contingency: Balanced Budget

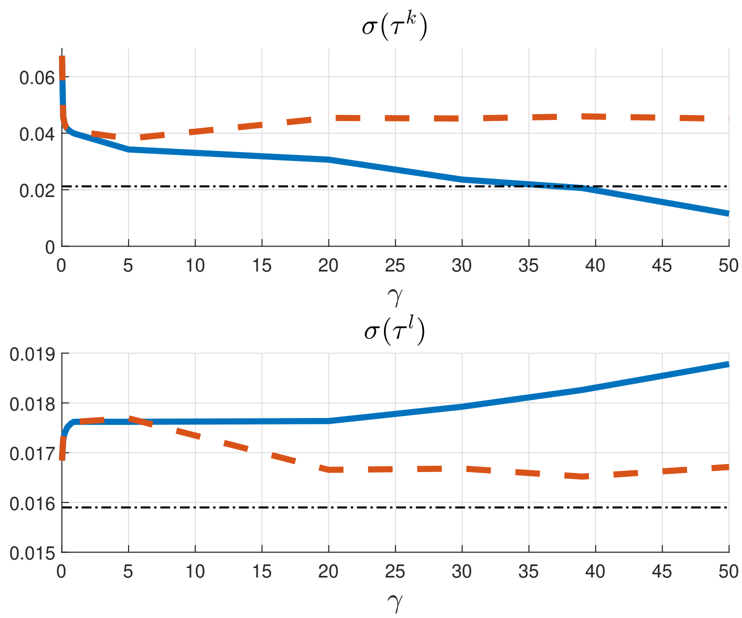

Notes: The figure displays the standard deviation of the capital tax (top panel) and labor tax (bottom panel). The solid line refers to our baseline model ($\gamma^k = \gamma^l = \gamma$) with Full Commitment and balanced budgets. The dashed line refers to the model with cost of state contingency only on the capital tax ($\gamma^k = \gamma$, $\gamma^l = 0$). The dashed-dotted line reports the empirical moment in US data 1971-2013 at annual frequency.

Table C1: Balanced Budget Models: First and Second Moments of Allocations

Moment	FC	FC no CSC	FC pred. τ^k	LTC
Ey	0.338	0.338	0.339	0.334
Ec	0.19	0.19	0.19	0.189
El	1	1	1	1.004
Ei	0.08	0.08	0.081	0.077
St. Dev. log(y)	0.002	0.002	0.002	0.003
St. Dev. log(c)	0.032	0.03	0.031	0.033
St. Dev. log(l)	0.003	0.004	0.004	0.003
St. Dev. log(i)	0.014	0.012	0.021	0.023
Autocorr. log(y)	0.972	0.957	0.917	0.995
Autocorr. log(c)	0.991	0.998	0.996	0.99
Autocorr. log(l)	0.757	0.873	0.353	0.867
Autocorr. log(i)	0.763	0.873	0.595	0.953

Notes: The table reports first and second moments of allocations. The first four rows report the means of output y, consumption c, hours l, and investment i respectively; the next four rows report the standard deviations; the bottom four row report autocorrelations. The first column refers to the baseline calibration under Full Commitment ($\gamma^k = \gamma^l = 39$); the second column refers to the Full-Commitment model without costs of state contingency ($\gamma^k = \gamma^l = 0$); the third column refers to the Full-Commitment model with predetermined capital tax ($\gamma^k = \infty$,$\gamma^l = 0$); the fourth column refers to the model under Limited-Time Commitment with $\gamma^k = \gamma^l = 39$.

Table C2: Low Persistence of g: Second Moments

Moment	CSC	No CSC	Pred. τ^k
St. Dev. $\log(1+\tau^k)$	0.046	0.108	0.066
St. Dev. $\log(1+\tau^l)$	0.015	0.01	0.014
Autocorr. $\log(1+\tau^k)$	0.77	0.748	0.741
Autocorr. $\log(1+\tau^l)$	0.783	0.941	0.611

Notes: The table reports second moments of the tax rates on capital and on labor income in a model with lower persistence of g, balanced budgets, and Full Commitment. The first column refers to the calibration with $\gamma^k = \gamma^l = 39$; the second column refers to the model without costly state contingency ($\gamma^k = \gamma^l = 0$); the third column refers to the model with predetermined capital tax ($\gamma^k = \infty$,$\gamma^l = 0$).

Figure C2: Low Persistence of g: Taxes

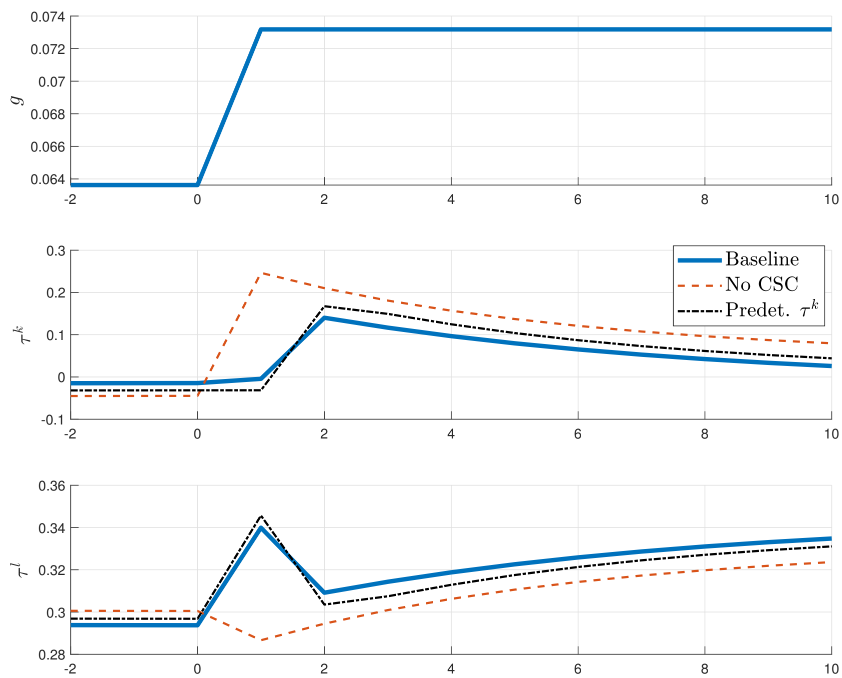

Notes: The figure displays the dynamics of fiscal variables around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. Top: government spending g_t; middle: capital income tax rate τ_t^k; bottom: labor income tax rate τ_t^l. Solid line: baseline model with costly state contingency ($\gamma^k = \gamma^l = 39$); dashed line: no costs of state contingency ($\gamma^k = \gamma^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma^k = \infty, \gamma^l = 0$).

Figure C3: Low Persistence of g: Allocations

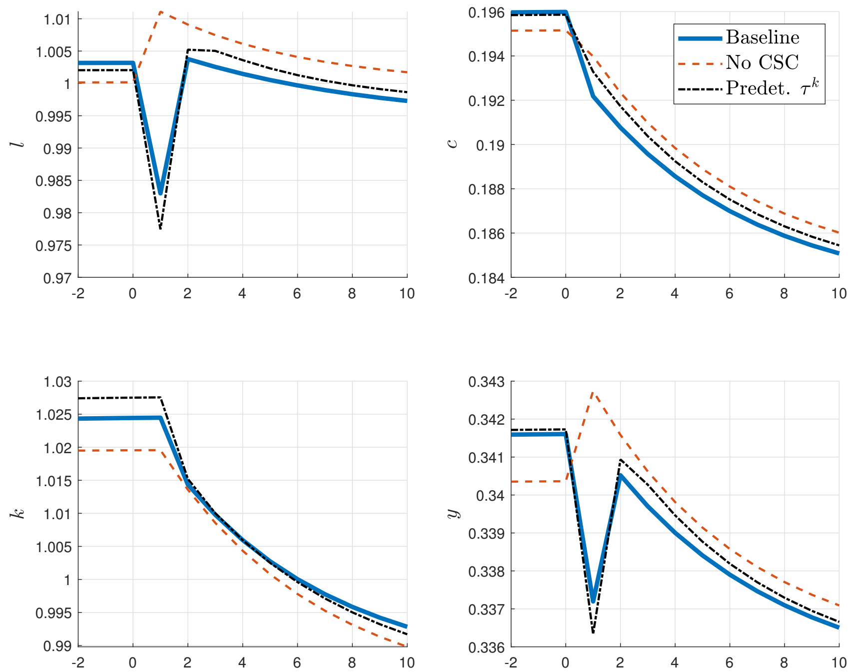

Notes: The figure displays the dynamics of allocations around a shock that increases government spending, at $t = 0$. Horizontal axes report time t. From top to bottom: labor l_t; consumption c_t; capital k_t; output y_t. Solid line: baseline model with costly state contingency ($\gamma^k = \gamma^l = 39$); dashed line: no costs of state contingency ($\gamma^k = \gamma^l = 0$); dashed-dotted line: predetermined capital tax ($\gamma^k = \infty, \gamma^l = 0$).

D Data

In this section discuss our data sources and construction, used for moment construction and our empirical impulse responses.

D.1 Measuring Capital and Labor Taxes

To measure capital and labor taxes in the data we follow the approach of measuring average tax rates using national accounts data. This approach is very common, and initially proposed by Mendoza, Razin, and Tesar (1994), and re-applied to US data by Jones (2002). Since then it has been used repeatedly in empirical work on taxation, such as Burnside, Eichenbaum, and Fisher (2004), Mertens and Ravn (2013), Leeper, Plante, and Traum (2010), and Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015).

An alternative approach accounts for non-linear tax schedules, with varying marginal tax rates. See, for example, Bhandari and McGrattan (2021), who construct a full marginal tax schedule for personal income taxes from US data, and model how business profits are taxed either as personal or corporate taxes. This approach has the advantage of being closer to the specifics of a given country's tax system (in this case the US) but maps less directly onto our model which features a representative household and the typically assumed system of linear capital and labor taxes. In practice, however, Jones (2002) finds that the two approaches yield taxes of a similar magnitude (to the extent they can be compared) and with a high correlation between the series.[20]

D.2 Data

For consistency with our model, we use data at the yearly frequency. We directly use yearly data where possible, and aggregate any quarterly data up to the yearly frequency.

Data for constructing tax variables We follow Fernández-Villaverde, Guerrón-Quintana, Kuester, and Rubio-Ramírez (2015) and include both federal and state taxes when constructing our data, so that our taxes capture all taxes paid domestically by households in the US. The data come from the National Accounts (NIPA) tables provided by the Bureau of Economic Analysis (BEA). We name variables for use in the formulas below:

[20] Jones finds a correlation between average tax rates and marginal tax rates of 0.9 and 0.8 for labor and capital taxes respectively, and with similar average levels. See Jones (2002) for details.

- Output and spending
  - DEF: output deflator (NIPA Table 1.1.4, line 1)
  - $NGDP$: nominal GDP (NIPA Table 1.1.5, line 1)
  - PCE: personal consumption expenditures (NIPA Table 1.1.5, line 2)
  - GSP: government spending and investment (NIPA Table 1.1.5, line 22)
- Incomes
  - CEM: compensation of employees (NIPA Table 1.12, line 2)
  - WSA: wage and salary accruals (NIPA Table 1.12, line 3)
  - PRI: proprietor's income (NIPA Table 1.12, line 9)
  - RI: rental income (NIPA Table 1.12, line 12),
  - CP: corporate profits (NIPA Table 1.12, line 13)
  - NI: interest income (NIPA Table 1.12, line 18)
- Taxes
  - TPI: taxes on production and imports (NIPA Table 3.1, line 4)
  - CT: taxes on corporate income (NIPA Table 3.1, line 5)
  - CSI: contributions to Social Security (NIPA Table 3.1, line 7)
  - PIT: federal, state, and local taxes on personal income (NIPA Table 3.2, line 3 plus NIPA Table 3.3, line 4)
  - PRT: state and local property taxes (NIPA Table 3.3, line 9)

Data for, and construction of, other basic variables Real GDP is nominal GDP over the price deflator (NGDP / DEF). Real government spending is nominal government consumption and investment over the price deflator (GSP / DEF).

For government debt we start with data on nominal "Debt held by the public", which excludes debt held by other government departments. This is available from FRED (series FYGFDPUB) from 1939 onwards, and we extend back to earlier years using data from the

CBO.[21] The data is annual year-end debt. We convert this to real debt by dividing by the GDP deflator.

Finally, for our impulse response exercises we use the defence news shock series from Ramey and Zubairy (2018). Data are taken from their online appendix and converted from quarterly to annual by dividing the sum of news within the year by the total potential GDP of that year.

The time spans of our raw data series are as follows. Our NIPA data run from 1929 to 2021, as does our debt data once extended with the CBO data. The news shock data runs from 1890 until 2016. We thus have complete coverage for 1929 to 2016, and we discuss the different samples we use for our exercises in the text below.

D.3 Constructing our tax series

First step: Personal Income Tax The personal income tax is a key tax in the US, which applies to income which will be classified as either labor or capital income in our model. Hence, a first step is to measure the average personal income tax in the data. In the data we directly measure before-tax personal income, PI_t, and personal income taxes paid, PIT_t, so measuring the personal income tax rate is simply done as $\tau_{p,t} = PIT_t/PI_t$. The personal income tax in the data is

$$\tau_{p,t} = \frac{PIT_t}{LI_t + CI_t} \tag{D1}$$

where total taxable personal income, PI_t, is split into personal labor income, PLI_t, and personal capital income, PCI_t. These are defined as

$$PLI_t = WSA_t + PRI_t/2 \tag{D2}$$

$$PCI_t = PRI_t/2 + RI_t + CP_t + NI_t \tag{D3}$$

Labor income is wages and salaries plus half of proprietors income. The half split is arbitrary and from Jones (2002), who finds results are robust to how proprietors income is split. Capital income is made up of four components: half of proprietors income, and then rental

[21]Specifically, there is historical CBO data which we choose not to use as our main data since it is measured as debt to GDP and only given to the first decimal place, and only up to the year 2000. The two data series are almost identical in their overlapping years, and we splice them together. The CBO data is as the Economic and Budget Issue Brief "Historical Data on Federal Debt Held by the Public", available at https://www.cbo.gov/publication/21728.

income, corporate profits, and interest income.

Capital Tax In the data we measure capital taxes paid, KIT_t, and total taxable capital income, TCI_t. So measuring the tax is simply done as $\tau_{k,t} = KIT_t/TCI_t$:

$$\tau_{k,t} = \frac{\tau_{p,t} PCI_t + CT_t + PRT_t}{PCI_t + PRT_t} \tag{D4}$$

The numerator is capital taxes paid. This is capital taxes paid out of personal income, plus taxes paid on corporate income and property taxes. The denominator measures total capital income which adds property taxes, PRT_t, back to personal capital income ($TCI_t = PCI_t + PRT_t$). Property taxes are subtracted from profits and hence missing from personal capital income, and so are added back to the denominator to properly measure total capital income.

Labor Tax In the data we measure labor taxes paid, LIT_t, and total taxable labor income, TLI_t. So measuring the tax is simply done as $\tau_{l,t} = LIT_t/TLI_t$:

$$\tau_{l,t} = \frac{\tau_{p,t} PLI_t + CSI_t}{CEM_t + PRI_t/2} \tag{D5}$$

The numerator is total income taxes. These come from two sources. Firstly, personal labor income is taxed at rate $\tau_{p,t}$. Secondly, there are additional contributions to social security, CSI_t, which are not taxed as personal income. The denominator is total labor income, which is total labor compensation, CEM_t, plus half of proprietor income ($TLI_t = CEM_t + PRI_t/2$).

Consumption Taxes We do not use consumption taxes in our model or empirical exercises, but include details on how to construct consumption taxes using the Jones (2002) method here for reference. In a model, we would think of total consumption expenditure as $CS_t = (1 + \tau_{c,t})C_t$, where C_t is the amount of real good that is bought and CS_t is the total spending. The data gives CS_t and the tax bill, $CTAX_t = \tau_{c,t}C_t$, giving $\tau_{c,t} = CTAX_t/C_t = CTAX_t/(CS_t - CTAX_t)$:

$$\tau_{c,t} = \frac{TPI_t - PRT_t}{PCE_t - (TPI_t - PRT_t)} \tag{D6}$$

The numerator is taxes on production and imports (TPI_t) less state and local property taxes (PRT_t). Production taxes are equivalent to consumption taxes in the standard model. Property taxes are included in production taxes in the data, but are better thought of as capital taxes, so are subtracted and counted in the capital tax instead. The denominator is consumption spending before taxes: PCE_t is personal consumption expenditure, which includes taxes, so the tax is subtracted.

D.4 Real-World Composition of τ^L and τ^K

In this section we briefly discuss the real-world composition of capital and labor taxes, and argue that this does not introduce artificial correlations between the taxes which are driving our results. In particular, both taxes in the data are constructed from private income taxes, τ_p, plus other sources. Could it be that the data properties of both series are being driven by private income taxes, and so in the real world it is simply not possible for measured capital and labor taxes to behave differently? We find that the answer to this question is no.

First consider the average make up of each measured tax in the data, as per the formulas (D4) and (D5). On average, the measured capital tax income (KIT) comes 38% from personal income tax ($\tau_P PCI$), 27% from corporate taxes (CT) and 36% from property taxes (PRT). For labor tax income (LIT) on average 51% comes from the personal income tax ($\tau_P PLI$) and 49% from social security (CSI). We see that a large, but not much more than a half, fraction of each tax comes from personal income taxes, so a significant remainder comes from other sources, especially for capital taxes.

More important is to check how personal income taxes affect the time series properties of each tax, which we do in the following way. We construct hypothetical taxes, holding the personal income tax rate at its time series average, $\bar{\tau}_p$:

$$\hat{\tau}_{k,t} = \frac{\bar{\tau}_p PCI_t + CT_t + PRT_t}{PCI_t + PRT_t} \qquad \hat{\tau}_{l,t} = \frac{\bar{\tau}_p PLI_t + CSI_t}{CEM_t + PRI_t/2} \tag{D7}$$

We then plot these hypothetical taxes versus the true taxes in Figure D1. We see a stark difference between capital and labor taxes: Holding the personal income tax rate fixed (red line), the capital tax series looks very similar to the true capital tax series (blue line). In particular all of the large swings in the capital tax in the data are still there even holding the personal income tax fixed. Hence, most of capital tax variation in the data is not due to the personal income tax, and must instead be driven by changes in corporate or property taxes. For labor taxes the story is very different: Holding the personal income tax rate fixed (red line), the labor tax series looks very different from the true series (blue line). Holding the personal income tax rate fixed misses both the trends and rises and falls in the labor tax. Hence, the dynamics of the labor tax are actually driven mostly by the dynamics of the personal income tax, while the other component, social security contributions, does not move as much.

Figure D1: Role of Personal Income Taxes

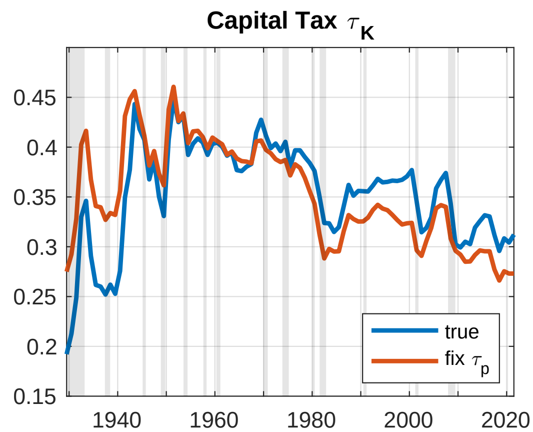

Notes: Actual tax rates (blue line) versus counterfactual tax rates with the personal income tax rate held constant (red line) in the data. See text for details.

In summary, capital and labor taxes in the US data have very different driving forces, even accounting for the US tax system. Labor tax movements seem to be dominated by changes in the personal income tax, while capital tax movements are driven more by changes in corporate income and property taxes. This reflects both that 1) personal income taxes make up a slightly larger share of labor taxes than capital taxes, and 2) the additional components of capital taxes (corporate and property taxes) are more volatile than the additional component of labor taxes (social security contributions).

D.5 Impulse Response of Capital and Labor Taxes to Military Spending Shocks

In this section we provide details of the impulse responses presented in Section 5.6. We are interested in the dynamics of taxes and debt to an exogenous increase in government spending, which we identify as the defence news shocks of Ramey and Zubairy (2018). They use a narrative approach to measure announced planned changes in government defence spending, as a fraction of potential GDP. In their work, they study the response of GDP to this shock to measure government spending multipliers, using local projections (Jordà, 2005). We adapt their approach to instead measure the response of fiscal instruments. We thus build on the work of Burnside, Eichenbaum, and Fisher (2004) who compute such impulse responses using a Vector Autoregression (VAR) approach. We differ from their study mainly in that 1) we use updated data, and also consider the response of debt, 2) we use the actual values of the defence news shock as an instrument, while they use the dates of major shocks as an instrument, 3) we use the local projection approach instead of a VAR.

Our baseline specification, adapted from Ramey and Zubairy (2018), is as follows:

$$x_{t+h} - x_{t-1} = \alpha_h + A_h z_t + \beta_h Z_t + \phi trend_t + \epsilon_{t+h} \qquad \text{for } h = 0, 1, 2, ..., H \tag{D8}$$

For any left-hand-side variable x, we are regressing the forward difference h periods ahead, $x_{t+h} - x_{t-1}$, on the the military spending shock, z_t, and a set of controls, Z_t. We consider as left-hand-side variables i) log real GDP, ii) log real government spending, iii) the level of capital taxes, iv) the level of labor taxes, and v) the log of real government debt. Each is regressed separately, giving $5 \times H$ regressions with associated coefficients. We control for a trend which is a fourth-order polynomial of time. Z_t are the controls used in a typical local projection set up. In particular, Z_t consists of lags of all of the five left hand side variables and the shock z_t. Since we are using yearly data, we use two lags. We use robust standard errors.[22] Our impulse responses plot the coefficients A_h for each variable, multiplied by a scaling factor (common to each variable) chosen to create a defense shock which leads to a 50% peak increase in total government spending.

Notice that the regression above is a simple OLS regression, but has an instrumental variables flavor since the dependent variable of interest, z_t, is considered as an exogenous shock due to the narrative identification. Since government spending itself is one of the variables regressed on this variable we are automatically performing an effective "first stage" regression to check that total government spending does rise in response to the military spending shock. The remaining variables then investigate how the rest of the economy responds to this military-spending-induced rise in government spending.

Our baseline results, as presented in the main text, use all available data to maximize power. Accordingly, we use the full dataset from 1929 to 2016 for our regressions. In Figure 6 in the main text we presented results for the fiscal variables, and in Figure D2 we

[22] Specifically, the estimations are run in Stata using the ivreg2 package with options robust and bw(auto).

Figure D2: Empirical response of output to a military news shock

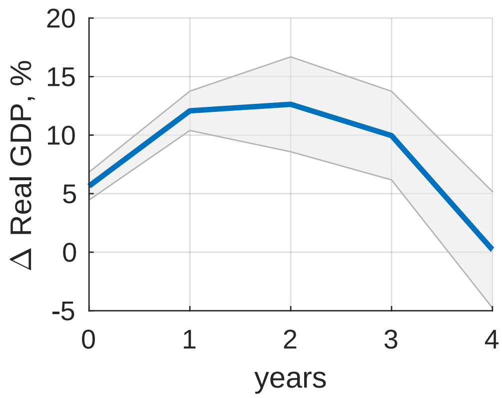

Notes: The figure displays the results of a local projection estimation on yearly US data from 1929 to 2015, giving the impulse responses to a Ramey and Zubairy (2018) military spending news shock. Solid lines give point estimates and the ranges are 95% confidence intervals. See text for further details.

present the remaining impulse response for real GDP. The results all appear sensible, with government spending and output rising after the shock, being paid for by a rapid rise in the capital tax, and a slower rise in the labor tax and stock of debt.

Robustness: The results are robust to changing the lag structure, and do not meaningfully quantitatively change when using either 1 or 3 years of lags. We also checked robustness to changing the estimation sample. We find that the results are driven mostly by the large positive defense spending shocks at the beginning of the sample in the 1940s and 1950s, which makes sense given that they are the largest shocks in the series. Thus changing the end date of the sample has limited effect on the results.

The results for the capital and labor tax responses are robust to using the same sample dates as Burnside, Eichenbaum, and Fisher (2004), who use data from 1947 to 1995, and hence drop the 1940s defense run up from the sample. We find similar results to our baseline (and hence their results) in this case, but only find significance for the labour tax when using one year of lags. This could be due to the reduced power from the smaller sample and our use of yearly data. Burnside, Eichenbaum, and Fisher (2004) use quarterly data with six quarters of lags. Overall, with the caveat that the early shocks are important for identification, our empirical results appear to be robust.