

Time-Consistent Public Expenditures

Paul Klein,^a Per Krusell,^b José-Víctor Ríos-Rull^{c, *}

^a*University of Western Ontario*

^b*University of Rochester, Institute for International Economic Studies, and CEPR*

^c*University of Pennsylvania, CEPR, and NBER*

First draft: 10 August 2001

This version: February 18, 2002

PRELIMINARY

Abstract

In this paper, we study the optimal choice of public expenditures when there is no way of committing to future policy and “reputational” mechanisms are not operative. This amounts to confining our attention to Markov equilibria. The environment is a neoclassical growth model where consumers derive utility from a public good. This environment gives rise to a dynamic game between successive governments and the private sector and this game is made interesting by the presence of a state variable: the capital stock. We characterize equilibria in terms of an intertemporal first-order condition (a “Generalized Euler Equation”, or GEE) for the government and we use this condition both to gain insight into the nature of the equilibrium and as a basis for computation.

The GEE reveals how the government optimally trades off tax wedges over time. It also allows us to discuss in what sense a current government may be strategically influencing future governments in their taxation decisions. For a calibrated economy, we find that when the tax base available to the government is capital income—an inelastic source of funds at any moment in time—the government still refrains from taxing at high rates in order to smooth distortions over time. As a result, the economy is far from the mix of public and private goods that would be optimal in a static context, but the capital stock and production are high.

Keywords: Time-consistency, Markov-perfect equilibrium, optimal taxation.
JEL classification: E62, H21.

*We thank Harald Uhlig for valuable comments. Correspondence should be sent to Krusell at the Department of Economics, Harkness Hall, University of Rochester, Rochester, NY 14627. Krusell thanks the National Science Foundation for support and Ríos-Rull thanks the National Science Foundation (Grant SES-0079504) and the University of Pennsylvania Research Foundation for their support.

1 Introduction

How should public expenditures be determined over time? We study this question under the assumption that the government cannot commit to its future policy decisions. The framework is a canonical one: in an otherwise standard neoclassical growth model, public goods, which are perfect substitutes with private goods in production, directly yield utility every period, and the government has access to proportional taxation of a given tax base. In this paper, we also assume that the government cannot finance its expenditures with deficits: the government cannot even commit to paying back loans. (The possibility of deficit finance is an interesting one but we leave it for future work.) We insist on treating the government as rational and impose a time-consistency requirement: an allocation in our economy is an equilibrium of a dynamic game played between a sequence of governments, each foreseeing how its successors will behave. Our focus is on Markov equilibria: we assume that “reputational mechanisms”, for one reason or other, are not operative.¹

The paper has three main contributions. First, it shows how a Markov equilibrium can be found as the fixed point of a set of functional equations. These functional equations include those that typically characterize the competitive equilibria for given policies and, in addition, a functional equation that we denote the *Generalized Euler Equation* (GEE). The GEE is the intertemporal first-order, necessary condition for optimal government policy. The GEE summarizes the marginal costs and benefits of a change in current taxes and constitutes a convenient tool for forming intuition as well as gauging which costs and benefits may be quantitatively important. The GEE is a rather complicated expression, but we show that it amounts to setting a weighted average of “wedges”—distortions—equal to zero. Thus, there exists a dynamic extension of standard public economics concepts, and we derive its exact form. The GEE has terms that are non-standard since they are derivatives of the key unknown policy functions.

The second contribution of this paper is to show how to compute equilibria with

¹Our work, thus, is closer in spirit to the original Kydland and Prescott (1977) paper than some later papers, notably Chari and Kehoe (1990) and related work which, following Abreu, Pearce, and Stacchetti (1990), discuss conditions under which better (and worse) equilibria can be attained using triggers under the assumption that agents are sufficiently patient.

controlled accuracy. The presence of functional derivatives in the functional equations means that it is not possible to find a steady state without simultaneously characterizing dynamics. We show how a perturbation method can be used to easily compute the steady state. We provide a few examples with closed-form solutions that illustrate both the theoretical and the computational contents of the paper.

The third contribution of the paper is to characterize the time-consistent, Markov outcomes for a calibrated economy. One key finding here is that time-consistent equilibrium taxes can be much lower than those that would result in a static economy, even though the tax base is ex-post inelastic (as in the case of a capital-income tax). We compare the properties of the time-consistent Markov allocation both qualitatively and quantitatively with those that arise in an environment with lump-sum taxes (the Pareto allocation) and those of an environment where the government can commit to future policy (the Ramsey allocation).

Our focus on a Markov equilibrium is a selection device, and we see it as a way of singling out the infinite-horizon equilibrium which is a limit of finite-horizon equilibria (if the latter are unique and the limit exists). We believe that it is fruitful to pay separate attention to this “fundamental” equilibrium not only because it has a close connection with equilibria in with long but finite horizons, but also because it provides a benchmark against which the best reputation equilibrium, if one exists, can be contrasted.

Because we look at a growth model, the dynamic game has a state variable: the capital stock. This makes the game complicated, but it also makes it interesting. Capital is the reason why the Ramsey allocation is time-inconsistent: decisions about taxes at time $t > 0$ influence savings decisions in earlier periods. In a time-consistent Markov equilibrium without commitment, a current government would thus like the next government to set its tax at a lower level than it does, because the latter sees savings as inelastic. Because of the effect a current policy decision has on private budgets and prices, however, the accumulation of capital provides a channel through which the current government, despite its inability to directly affect future policy, can influence outcomes beyond the present period. We do not, however, wish to label this influence “strategic manipulation”, because given a level of savings, the current

and next-period governments have preferences that are perfectly aligned: an envelope theorem applies. In this sense, the present dynamic game contrasts sharply with the savings games between selves with conflicting discounting studied recently (see, e.g., Strotz (1956), Phelps and Pollak (1968), Pollak (1968), Laibson (1997), and Krusell, Kuruşçu, and A. (2000)). In those games, the objectives conflict, and manipulation is an active part of equilibrium behavior. However, even though the government in our economy does not try to manipulate its successors, it still wishes to influence capital accumulation. This is because current taxation can, via wealth effects, alter private decisions in a direction that alleviates future or intertemporal distortions: as pointed out above, the government trades these distortions off against one another, and the capital stock is the vehicle through which future distortions can be affected.

Markov equilibria may not be unique. In a similar context—a dynamic savings game between successive selves with conflicting time discounting—Krusell and Smith (2000) show that there is indeterminacy not only of equilibrium paths but of steady states as well. However, these Markov equilibria, which do not correspond to limits of finite-horizon equilibria, rely on discontinuous savings rules. Therefore, we insist here on *differentiable* policy rules. As already pointed out, our key equilibrium condition, the GEE, explicitly involves the *derivative* of the decision rule of the government decision maker. In contrast, in standard frameworks—where the commitment solution is time-consistent or where the government is assumed to be able to commit directly—the decision-rule derivatives never appear in the first-order conditions.

In our quantitative experiments it turns out that the properties of taxes and allocations in the time-consistent, Pareto, and Ramsey equilibria differ markedly, even in environments where there are no a-priori reasons to think that the existence of commitment matters, such as in economies with taxes only on labor income—a static tax. We also find that even though reputation is by definition ruled out, the mechanisms that are left—which involve the effects of current taxation on the capital stock bequeathed to the next decision maker—can be quite powerful and even qualitatively surprising. In an economy where labor supply is exogenous and the government taxes current capital income alone to finance the current provision of public goods, it does not provide an outcome with an optimal mix of private and public goods, even though the capital income tax is equivalent to a lump-sum tax. Foreseeing that the current

governments will tax capital, which retards capital accumulation earlier on, the current government wishes to increase that capital accumulation. It does so by taxing current capital income less: it lets the consumers keep some resources, thus using a wealth effect to increase current and future savings in the direction it desires.

The tools developed in this paper are, we think, entirely general and applicable to a wide variety of contexts. There is earlier work in this direction. First, Markov equilibria of the type that we are interested in have been studied in Cohen and Michel (1988) and Currie and Levine (1993), who explore linear-quadratic economies. In such economies, Markov equilibria can be characterized and computed explicitly, since the first-order conditions become linear in the state variable. In other words, the derivatives of decision rules here are constants, and although they play a role in the solution, higher-order derivatives of these rules are all identically zero. The drawback, of course, of linear-quadratic settings is that they only apply in extremely special settings. Thus, either one has to give up on quantitative analysis to apply them, or accept reduced-form objective functions and/or reduced-form private decision rules.

There is also a literature both in political economy (Krusell, Quadrini, and Ríos-Rull (1997), Krusell and Ríos-Rull (1999)) and in optimal policy with a benevolent government (Klein and Ríos-Rull (1999)) that has used computational methods to find quantitative implications of Markov equilibria for a variety of questions. This work is closely related to the present one, but it has two drawbacks. First, the methods used—essentially, numerical solution of value functions based on linear-quadratic approximations—are of the “black-box” type: they do not deliver interpretable conditions, such as first-order conditions for the key decision maker. The present paper fills this gap. Secondly, the numerical methods do not deliver controlled accuracy. In contrast, the methods used herein do.

In a related paper, Phelan and Stacchetti (2000) have looked at environments like those studied in this paper and have developed methods to find all equilibria. Their methods, however, do not allow Markov equilibria to be identified and explicitly interpreted. The only other closely related literature is that upon which the present work builds quite directly: the analysis of dynamic games between successive selves, as outlined in the economics and psychology literature by Strotz (1956), Phelps and

Pollak (1968), Pollak (1968), Laibson (1997), and others. This literature contains the derivation of a GEE, and Krusell, Kuruşçu, and A. (2000) show how to solve it numerically for a smooth decision rule equilibrium. As will be elaborated on below, the smooth rule can be difficult to find with standard methods, and Krusell, Kuruşçu, and A. (2000) resort to a perturbation method, which we also use here. This method relies on successive differentiation of the GEE. Here is one place where smoothness of the policy function becomes operationally important.

The broad outline of the paper is as follows. In Section 2 we describe our baseline environment, in which the only private economic decision is the consumption-savings choice (Section 2.1), define a Ramsey equilibrium (Section 2.2), and then define and discuss our time-consistent, Markov equilibrium (Section 2.3) step by step. The analysis of our equilibrium—which involves interpretations of the GEE and of government behavior as well as comparisons with alternative ways of stating the government problem/defining equilibrium—is contained in Section 3. Section 4 then discusses an extension to our baseline setup where leisure is valued and where there are different possibilities for what tax base might be used. In Section 5, we provide parametric examples for which closed-form solutions can be obtained. Section 6 discusses the properties of the policies that arise in an environment calibrated to U.S. data where governments do not have access to a commitment technology (Markov policies) and compares them to those that arise both in environments with commitment (Ramsey policies) and in environments where the government has access to lump sum taxation (Pareto policies). Section 7 concludes. The Appendix includes some auxiliary formal definitions and the description of the computational procedures we use.

2 The model

In this section, we describe the specific setup. We then define a benchmark “Ramsey equilibrium”—the solution to an optimal-policy problem where the government can commit to future policies. After that, we proceed toward a definition of a time-consistent equilibrium where the government does not have the ability to commit.

2.1 The environment

Our model is rather canonical. We consider a standard growth model with an infinitely-lived, representative household and a benevolent government with a period-by-period balanced budget and proportional taxation. In the first part of the paper, the tax base is total income and leisure is not valued.

In a competitive equilibrium, households maximize

$$\sum_{t=0}^{\infty} \beta^t u(c_t, g_t)$$

subject to

$$c_t + k_{t+1} = k_t + (1 - \tau_t) [w_t + (r_t - \delta)k_t].$$

Firms maximize profits; using a constant-returns-to-scale production function $f(K, L)$, where f is concave, they employ inputs so that w_t and r_t are the marginal products of labor and capital, respectively.

Using capital letters to denote economy-wide levels, the resource constraint in this economy reads

$$C_t + K_{t+1} + G_t = f(K_t, 1) + (1 - \delta)K_t$$

The government's balanced-budget constraint, thus, reads

$$G_t = \tau_t [f(K_t, 1) - \delta K_t].$$

We will make use of the following functions:

$$\mathcal{G}(K, \tau) \equiv \tau [f(K, 1) - \delta K]$$

and

$$\mathcal{C}(K, K', \tau) \equiv f(K, 1) + (1 - \delta)K - K' - \mathcal{G}(K, \tau),$$

where 's denote next-period values. These functions— \mathcal{C} representing consumption as a function of current and next-period capital and the current tax rate and \mathcal{G} representing government expenditures as a function of current capital and the current tax rate—are exogenous and will economize on notation significantly.

2.2 Commitment: the Ramsey problem

If lump-sum taxes were available, the optimal allocation in this economy would involve two conditions: $u_c(C_t, G_t) = \beta(1 + f_k(K_{t+1}, 1) - \delta)u_c(C_{t+1}, G_{t+1})$ (optimal savings) and $u_c(C_t, G_t) = u_g(C_t, G_t)$ (optimal public expenditures). In our economy lump-sum taxes are assumed not to be available, and the the optimal allocation using a proportional income tax is more involved.

We will first assume that the government has the ability to commit to all its future policy choices at the beginning of time. The government's decision problem is therefore to choose a sequence of tax rates $\{\tau_t\}_{t=0}^{\infty}$ in order to maximize utility, taking into account how the private sector will respond to these taxes. A simple way to describe this problem formally is to choose $\{\tau_t, K_{t+1}\}_{t=0}^{\infty}$ to maximize

$$\sum_{t=0}^{\infty} \beta^t u(\mathcal{C}(K_t, K_{t+1}, \tau_t), \mathcal{G}(K_t, \tau_t))$$

subject to the private sector's first-order conditions for savings (which, together with a transversality condition which will be nonbinding, are necessary and sufficient conditions for consumer maximization)

$$u_c(\mathcal{C}(K_t, K_{t+1}, \tau_t), \mathcal{G}(K_t, \tau_t)) = \beta u_c(\mathcal{C}(K_{t+1}, K_{t+2}, \tau_{t+1}), \mathcal{G}(K_{t+1}, \tau_{t+1})) [1 + (1 - \tau_{t+1})(f_k(K_{t+1}, 1) - \delta)] \quad (1)$$

for all $t \geq 0$. We refer to the solution of the problem as the Ramsey allocation.

This problem has a noteworthy feature: its solution will, in general, not be time-consistent. That is, the optimal sequence of taxes and capital stocks will not be optimal ex post: if the government could reoptimize at a $t > 0$, they would choose to not follow the original sequence. For this reason, the assumption that the government can commit to future taxes is a binding one.

To see the source of time-inconsistency, note that both K_{t+2} and τ_{t+1} appear in the constraint at time t : for $t > 0$, the choices of K_{t+1} and τ_t influence the savings and tax choices in the previous period (K_t and τ_{t-1}). However, this is not true at $t = 0$, since K_0 is given (and τ_{-1} does not enter the problem at all). This means that the

first-order condition for the choice of K_1 has one term less than that for the choice of any other K_t (and similarly for τ_0 vs. any other τ_t). So if $\{\tau_t, K_{t+1}\}_{t=0}^\infty$ satisfies all the first-order conditions for the above problem, it will not satisfy all the first-order conditions for the problem of reoptimizing beginning at time s . Reoptimization would involve choosing $\{\tau_t, K_{t+1}\}_{t=s}^\infty$ subject to the private sector's first-order conditions for savings for periods $s, s+1, \dots$. Clearly, the original chosen sequence could not satisfy the first-order condition for the reoptimization problem in period s ; hence the time-inconsistency.

Intuitively, the tax rate chosen by the government for time $t > 0$ does influence—distort—the savings choice in period $t-1$ (and therefore also in any earlier periods), but if it were to reoptimize at time t , it would not recognize this distortion.

2.3 No commitment: time-consistent equilibrium

We now assume that the government is not able to commit to future tax rates. Conceptually, we need the government to choose the tax rate only in the current period, while figuring out all the effects of this choice, present and future. The effects in the future are present through the capital accumulation decision of private households. A current tax will influence this choice—via a wealth effect—and hence the state variable for next period changes. Through this channel, the current tax rate affects next period's savings, too, and so on. In addition, to the extent that the current tax choice depends on the current capital stock, a tax choice today will also influence future tax policy—through the effects on capital accumulation. Therefore, it is possible for the current government to directly—i.e., not through “reputation effects”—manipulate future governments. We will now proceed to study how this possibility of manipulation is exercised.

The formalization of a time-consistent equilibrium proceeds in three steps. We employ recursive methods. We first define a recursive competitive equilibrium for given policy (Section 2.3.1). This is done in such a way as to allow us to define the government's problem (Section 2.3.2); a definition of a Markov-perfect equilibrium then follows in Section 2.3.3. This definition is phrased in terms of value functions and policy

functions. Our main focus, however, is more narrow: we seek the Markov equilibrium which is a limit of finite-horizon equilibria. In our general treatment here, we will make two assumptions: (i) we assume such an equilibrium exists; and (ii) we assume that this equilibrium has differentiable policy functions. Based on these assumptions, we provide a definition of such equilibria in terms of two functional equations in policy function space—Section 2.3.4. Our characterization of such an equilibrium then continues in Section 3.

2.3.1 Recursive competitive equilibrium for given policy

We will now define a recursive competitive equilibrium given a policy function Ψ : we let τ be the current income tax rate and $\tau = \Psi(K)$ is a rule specifying what tax rate will apply for every value of the state variable of the economy (the stock of capital K). The function Ψ is the fundamental equilibrium object we seek to characterize—how the government taxes—but for now we will regard it as an arbitrarily given function. The gist of the definition is a function for the aggregate law of motion of capital that depends only on K . It arises from imposing the representative-agent condition on the individual decision rules of agents, who understand that the economy evolves according to those functions and to the government policy function Ψ . However, we depart from the fully recursive way of defining these equilibrium functions: a fully recursive way would give next period’s capital as a function of the current capital stock alone, given that the government behaves according to Ψ . Instead, we include a one-period deviation for government policy: we define equilibrium for a given Ψ function as a function that yields next period’s capital as a function of both the aggregate capital stock and the current tax rate. That is, the idea is that the current tax rate is set at an arbitrary value but that future tax rates follow Ψ , evaluated at future capital values. We thus write the equilibrium function capital accumulation as

$$K' = \mathcal{H}(K, \tau) \tag{2}$$

The explicit dependence of this function on the tax rate allows us to use it to pose the problem that the government faces in any given period: it has the freedom to choose the current policy, but has to take as given how future governments react. In particular, the future governments respond to the capital stock they inherit, through the function

Ψ . Moreover, the current government can affect next period's capital according to \mathcal{H} —through its second argument—and the capital stocks after that as well, through similar channels. In Appendix A we define the concept of recursive competitive equilibrium in detail; it involves a value function for the consumer, whose arguments are K and τ , and an accompanying decision rule for capital accumulation that, when the representative-agent assumption is invoked, reduces to \mathcal{H} —the function that each agent takes as given (along with Ψ) in her maximization problem.

The household's first-order condition for saving in recursive competitive equilibrium will be utilized extensively below. It can be represented as a functional equation, i.e., it has to be satisfied for all τ and all K . Thus, for (K, τ) we require

$$u_c(\mathcal{C}(K, K', \tau), \mathcal{G}(K, \tau)) = \beta u_c(\mathcal{C}(K', K'', \tau'), \mathcal{G}(K', \tau')) \cdot \{1 + [1 - \tau'] [f_K(K') - \delta]\}, \quad (3)$$

where K' , τ' , and K'' are all functions of (K, τ) :

$$\begin{aligned} K' &= \mathcal{H}(K, \tau) \\ \tau' &= \Psi(K') = \Psi(\mathcal{H}(K, \tau)) \\ K'' &= \mathcal{H}(K', \tau') = \mathcal{H}(\mathcal{H}(K, \tau), \Psi(\mathcal{H}(K, \tau))). \end{aligned}$$

This is the functional-equation version of equation (1) above; it defines \mathcal{H} . Notice how Ψ is a determinant of \mathcal{H} : the expectations of future government behavior influence how consumers work and save. Also, notice how any taxes and capital stocks further into the future can also be written, using \mathcal{H} and Ψ , to depend on the current variables K and τ : $\tau'' = \Psi(K'') = \Psi(\mathcal{H}(\mathcal{H}(K, \tau), \Psi(\mathcal{H}(K, \tau))))$ and $K''' = \mathcal{H}(K'', \tau'') = \mathcal{H}(\mathcal{H}(\mathcal{H}(K, \tau), \Psi(\mathcal{H}(K, \tau))), \Psi(\mathcal{H}(\mathcal{H}(K, \tau), \Psi(\mathcal{H}(K, \tau))))$), and so on.

2.3.2 The government's problem

We are now ready to state the government's problem in a time-consistent equilibrium. Note that the government **only** chooses this period's tax rate, and that it takes as given what future governments do. But it does not take the future governments' actions as

given; instead, it takes as given the policy function Ψ used by future governments. Before writing the problem of the agent, note that the current return for the government is given by

$$u(\mathcal{C}(K, K', \tau), \mathcal{G}(K, \tau)). \quad (4)$$

The following period, capital is given by $K' = \mathcal{H}(K, \tau)$. Note why this is the case: the private sector makes its choices of how much to work and how much to save (and this determines consumption and government expenditures) as a function of the state of the economy K , whatever policy rate the government makes today τ , and taking into account that future policies are given by function Ψ .

The government also needs a function for assessing the value of the future. Let this assessment be given by a certain function v . This function has as argument the state of the economy tomorrow, K' (and of course it depends on future behavior by the private households and the ensuing governments). Before describing how this function is determined, note that the problem of the government can be written as

$$\begin{aligned} \max_{K', \tau} \quad & u(\mathcal{C}(K, K', \tau), \mathcal{G}(K, \tau)) + \beta v(K') \\ & \text{subject to} \\ & K' = \mathcal{H}(K, \tau). \end{aligned} \quad (5)$$

The function v , whose role it is to add up all future utility streams in a standard way using discounting weights, can also be defined recursively with the functions that describe the actions of the households and of future governments. Thus,

$$v(K) \equiv u[\mathcal{C}(K, \mathcal{H}(K, \Psi(K)), \Psi(K)), \mathcal{G}(K, \Psi(K))] + \beta v[\mathcal{H}(K, \Psi(K))] \quad (6)$$

for all K defines v . One can therefore view the government's problem in two steps: (i) using \mathcal{H} and Ψ , define the value of any amount of capital left for the future, v ; and (ii) using v , solve a one-variable maximization problem: choose τ to maximize the objective.

2.3.3 Markov-perfect equilibrium

A *Markov-perfect equilibrium* now dictates that $\Psi(K)$ solves the above problem for all K :

$$\Psi(K) \in \arg \max_{\tau} \{u[\mathcal{C}(K, \mathcal{H}(K, \tau)), \tau], \mathcal{G}(K, \tau)] + \beta v(\mathcal{H}(K, \tau))\}. \quad (7)$$

This is our key fixed-point condition. It states that the rule that the governments follows ends up being the same as the one they perceive future governments to be using. This idea captures their rational expectations, or, in this context, the *time consistency* of the equilibrium.

By construction now, the problem of the government must satisfy

$$v(K) = \max_{\tau} u[\mathcal{C}(K, \mathcal{H}(K, \tau)), \tau], \mathcal{G}(K, \tau)] + \beta v[\mathcal{H}(K, \tau)]. \quad (8)$$

This is the *recursive problem* that a time-consistent policy has to solve. That is, unlike in the case of a Ramsey equilibrium, a time-consistent equilibrium has the government solve a fully recursive problem. The recursive problem, however, is not expressed in terms of primitives alone: it involves both Ψ and \mathcal{H} as determinants.

Formally, a Markov-perfect equilibrium is now a set of functions Ψ , v , and \mathcal{H} satisfying the functional equations (7), (6), and (3).

2.3.4 Differentiable Markov-perfect equilibrium

We now assume that our equilibrium policy functions Ψ and \mathcal{H} are differentiable and proceed to derive a functional-equation first-order condition for the government's choice. This equation, which we will refer to as the GEE (the Government's, or Generalized, Euler Equation), will be in focus in the analysis below. We shortly discuss, in Section 2.3.5, our reasons to consider differentiable equilibria.

There are several ways to derive the GEE. The most straightforward way at this point is to start by deriving a first-order condition from the government's recursive problem, which will contain the unknown function $\frac{\partial v}{\partial K}$, and then use an envelope condition to eliminate $(\frac{\partial v}{\partial K})'$. This procedure is nonstandard only in that it is somewhat more roundabout to eliminate $(\frac{\partial v}{\partial K})'$ than in the standard growth model.

The first-order condition for the government produces

$$u_c(\mathcal{C}_{K'}\mathcal{H}_\tau + \mathcal{C}_\tau) + u_g\mathcal{G}_\tau + \beta\left(\frac{\partial v}{\partial K}\right)'\mathcal{H}_\tau = 0.$$

Here, we are economizing on notation by suppressing the arguments of the functions. This equation, thus, pins down τ . To obtain an expression for $\frac{\partial v}{\partial K}$, we differentiate equation (6), which has to hold for all K ; we obtain

$$\frac{\partial v}{\partial K} = u_c(\mathcal{C}_K + \mathcal{C}_{K'}(\mathcal{H}_K + \mathcal{H}_\tau \frac{\partial \Psi}{\partial K}) + \mathcal{C}_\tau \frac{\partial \Psi}{\partial K}) + u_g(\mathcal{G}_K + \mathcal{G}_\tau \frac{\partial \Psi}{\partial K}) + \beta\left(\frac{\partial v}{\partial K}\right)'(\mathcal{H}_K + \mathcal{H}_\tau \frac{\partial \Psi}{\partial K}) = 0.$$

This equation contains *indirect effects*, via $\frac{\partial \Psi}{\partial K}$. Notice, too, that a grouping of the $\frac{\partial \Psi}{\partial K}$ terms leads to $\frac{\partial \Psi}{\partial K} \cdot 0$, where the “0” results from use of the first-order condition above: this is the envelope theorem. However, unlike in the setting of the standard growth model, the use of the envelope theorem does not suffice to make $(\frac{\partial v}{\partial K})'$ disappear here: $\frac{\partial v}{\partial K}$ still depends on $(\frac{\partial v}{\partial K})'$:

$$\frac{\partial v}{\partial K} = u_c(\mathcal{C}_K + \mathcal{C}_{K'}\mathcal{H}_K) + u_g\mathcal{G}_K + \beta\left(\frac{\partial v}{\partial K}\right)'\mathcal{H}_K = 0.$$

This is not a problem, however: $(\frac{\partial v}{\partial K})'$ can be expressed in terms of primitives and decision rules from the first-order condition above. This delivers

$$\beta\left(\frac{\partial v}{\partial K}\right)' = -\frac{1}{\mathcal{H}_\tau}(u_c(\mathcal{C}_{K'}\mathcal{H}_\tau + \mathcal{C}_\tau) + u_g\mathcal{G}_\tau).$$

Thus, the expression for $\frac{\partial v}{\partial K}$ in terms of primitives and decision rules reads

$$\frac{\partial v}{\partial K} = u_c(\mathcal{C}_K + \mathcal{C}_{K'}\mathcal{H}_K) + u_g\mathcal{G}_K - \frac{\mathcal{H}_K}{\mathcal{H}_\tau}(u_c(\mathcal{C}_{K'}\mathcal{H}_\tau + \mathcal{C}_\tau) + u_g\mathcal{G}_\tau).$$

We can now update this expression one period and substitute back into the original first-order condition to obtain our GEE:

$$\beta\mathcal{H}_\tau \left\{ u_c[-\mathcal{H}_\tau - \mathcal{G}_\tau] + u_g\mathcal{G}_\tau + u'_c[f'_K + 1 - \delta - \mathcal{H}'_K - \mathcal{G}'_K] + u'_g\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau}(u'_c[-\mathcal{H}'_\tau - \mathcal{G}'_\tau] + u'_g\mathcal{G}'_\tau) \right\} = 0, \quad (9)$$

where we have also used the definition of \mathcal{C} in terms of primitives. Equation (9), where arguments are still suppressed for readability and primes on functions indicate that

the function is evaluated in the next period, holds for all K . It is our fundamental functional equation determining $\Psi(K)$ given $\mathcal{H}(K, \tau)$. I.e., it defines the government policy rule Ψ as the optimal policy determination under the assumption that the private sector behaves according to an arbitrary \mathcal{H} .

We are now ready to define our time-consistent equilibrium as a differentiable Markov-perfect equilibrium.

Equilibrium: A *time-consistent policy equilibrium* is a set of differentiable functions Ψ and \mathcal{H} such that

- $\mathcal{H}(K, \tau)$ solves the functional first-order condition (3) of the private sector; and
- $\Psi(K)$ solves the functional first-order condition (9) of the government.

The definition presumes that the first-order conditions are also sufficient for maximization of the problems of representative agent and of the government. We will discuss the interpretation of these equations, especially the GEE, in detail below, and we will discuss how to find Ψ and \mathcal{H} for calibrated parameter values.

2.3.5 Connections to a finite-horizon model: qualifications

Our goal is to find the Markov equilibrium which is a limit of the corresponding finite-horizon equilibria. Is there such a limit, and does our definition of a time-consistent (differentiable) equilibrium above produce it when it exists? Can our definition also be satisfied by an equilibrium which is not the limit of finite-horizon equilibria?

First, we provide no general conditions under which the limit is well-defined. Existence of an equilibrium for any finite horizon T can be obtained in a straightforward manner, but it is possible that the sequence of policy rules so obtained, even if there is a unique such rule for a given T , does not converge. It is also possible, in general, that several equilibria exist in the finite-horizon game. In this case, it would also be possible to construct “reputational” equilibria, despite the finite horizon (see Benoit and

Krishna (1985)). In other words, our equilibrium definition will only be useful in cases where the finite-horizon equilibria are unique and deliver policy rules that converge as T goes to infinity.

Second, is such a limit equilibrium differentiable? It need not be; here again, our work presumes that the problem is well-behaved enough (of course, we assume that our primitive functions u and f are sufficiently differentiable). As an illustration, we display a closed-form solution to our public-expenditure problem in Section 5 below that meets all the desired properties, and it is possible to show, under some conditions, that this equilibrium is also the limit of finite-horizon equilibria.

Third, in cases where the limit equilibrium is well-defined and differentiable, is it possible that there exist other differentiable Markov-perfect equilibria satisfying our two equilibrium functional equations? We cannot rule this out. In cases where a closed-form solution to our problem can be derived, it should be straightforward to check whether a differentiable equilibrium is also the limit equilibrium. Indeed, in our closed-form example we show that, if certain conditions are met, there can be a solution to the functional equations that is not the limit equilibrium. When no closed-form solution is available, one has to rely on global numerical search to ensure that no other differentiable solution exists.² The numerical procedure that we propose below (in the Appendix) is designed to find steady states. It is general in nature and designed as a global search. Thus, it will deliver a numerical answer to this question.

The search for an equilibrium as a limit of finite-horizon equilibria is a difficult one in general, and few theoretical results are available. Our approach here is entirely applied; it relies on a well-behaved problem, and we suspect that the careful numerical work that has to go along with any applied problem will reveal any pathologies. In the case we study, we are guided by analytical results for a benchmark economy: we reproduce these results numerically and then move away from this economy gradually.

²It is also possible that a differentiable solution exists when no limit equilibrium exists.

3 Characterization

We now move to our discussion of the GEE—the government’s first-order condition. After providing various interpretations of this condition, we present two reformulations of our problem. The first one of these is more compact than the one above, although perhaps less transparent. It is particularly useful in the numerical work later. The second reformulation casts the government’s problem as a sequential one. This problem is useful in delivering an easier method for deriving the GEE than the one above.

3.1 Interpretation of the government’s Euler equation

We begin with two alternative interpretations of the costs and benefits of raising current taxes. Thereafter, we address the question of whether there is a sense in which the current government, through the effect its taxation decision has on capital accumulation, manipulates its successors.

3.1.1 The macroeconomist’s version

A first property of the GEE is that it has a finite number of terms. That is, even though the current tax rate choice in general has repercussions into the infinite future—recall that the present government cannot “keep future variables constant” because it cannot commit future governments—the marginal costs and benefits at an optimum can be summarized with terms involving only two consecutive periods. This, of course, is due to the recursive structure and the use of the envelope theorem. The envelope theorem in this context means that, *when K' is viewed as given*, the current government agrees with the next government on how to set τ' ; these two governments have identical ways of evaluating utility from tomorrow and on. That is, disagreement is only present if the effects of τ' on K' are taken into account.

The recursive structure of the government’s problem makes it equivalent to a sequential problem, which we will state and discuss in some detail in Section 3.1.3. Thus, one can also view the GEE as resulting from a *variational* (2-period) problem: keeping

the state variables K and K'' fixed, vary K' , through the control variables τ and τ' , in order to obtain the highest possible utility. Viewed this way, one observes that any change in K' , which is effectuated by a change in τ , requires an accompanying change in τ' so that K'' remains unchanged. Total differentiation of $K'' = \mathcal{H}(K', \tau')$ thus states that this change in τ' has to be $\frac{d\tau'}{dK'} = -\frac{\mathcal{H}'_K}{\mathcal{H}'_{\tau'}}$. Notice that this term appears in the GEE: it thus reflects a partial change in τ' coming about due to a change in K' . It is not, however, equal to Ψ' , which is the net change in τ' .

Second, the GEE contains both primitive functions, such as marginal utility, and (endogenous) decision rules: \mathcal{H} and Ψ .³ Moreover, and this is why the term “generalized” Euler equation is appropriate, the equation contains derivatives of decision rules; it contains \mathcal{H}_K and \mathcal{H}_τ , both evaluated in the present and in the future. Thus, our equilibrium system of functional equations is actually a differential equation system.

As we shall see below, it is possible to eliminate the derivatives of \mathcal{H} by use of the consumer’s Euler equation for savings, equation (3). This equation defines $H(K, \tau)$, and by differentiation with respect to K and τ , respectively, these derivatives can be obtained. However, this differentiation will involve another unknown derivative— Ψ' —since Ψ is present in the forward-looking consumer’s Euler equation. That is, there is no way around the fact that our two key equilibrium functional equations are differential equations.

We will now interpret the GEE in terms of marginal benefits and marginal costs of changing τ . These benefits and costs will involve the unknowns \mathcal{H}_K and \mathcal{H}_τ . In our discussion, we will assume that the former of these is positive and the latter negative. These assumptions seem natural: under the normal goods assumption regarding c and c' (i.e., time-additive utility and a concave u), one would expect increased income to increase savings, and both an increase in K and a decrease in τ reflect increased income. However, notice that, to the extent K also changes tax rates (recall that K is not the individual’s state variable, but the economy-wide capital stock), which it does in our economy, other effects could be present. In our analytically solved example below as well as in our quantitative section, we confirm the assumed signs of \mathcal{H}_K and \mathcal{H}_τ .

³Due to the compact notation, Ψ is not visible in (9) but appears as an argument of \mathcal{G} .

For convenience, we restate the GEE.

$$u_c [-\mathcal{H}_\tau - \mathcal{G}_\tau] + u_g \mathcal{G}_\tau + \beta \mathcal{H}_\tau \left\{ u'_c [f'_K + 1 - \delta - \mathcal{H}'_K - \mathcal{G}'_K] + u'_g \mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} (u'_c [-\mathcal{H}'_\tau - \mathcal{G}'_\tau] + u'_g \mathcal{G}'_\tau) \right\} = 0,$$

We can thus describe our marginal benefits and costs as follows. In terms of effects on today's utility-relevant variables, a marginal increase in the current tax rate affects

1. current consumption, which
 - (a) goes up via lower savings, delivering a utility effect of $-u_c \mathcal{H}_\tau > 0$, and
 - (b) down via higher government spending, with an effect on utility of $-u_c \mathcal{G}_\tau < 0$; and on
2. current government spending, whose rise leads to a utility change of $u_g \mathcal{G}_\tau > 0$.

The effects of the tax hike on future utility-relevant variables occur via a decrease in savings ($H_\tau < 0$), leading to

3. effects on next period's consumption which
 - (a) goes down via a direct effect on production and undepreciated capital, affecting utility by $\beta \mathcal{H}_\tau u'_c (f'_K + 1 - \delta) < 0$;
 - (b) goes up via an indirect negative effect from lowered saving ($\frac{dK'}{d\tau} \frac{dK''}{dK'} = \mathcal{H}_\tau \mathcal{H}'_K < 0$), affecting utility by $\beta \mathcal{H}_\tau u'_c (-\mathcal{H}'_K) > 0$; and
 - (c) goes up via an indirect negative effect on government spending ($\frac{dK'}{d\tau} \frac{d\mathcal{G}'}{dK'} = \mathcal{H}_\tau \mathcal{G}'_K < 0$), affecting utility by $\beta \mathcal{H}_\tau u'_c (-\mathcal{G}'_K) > 0$;
4. a decrease in government spending, producing a utility change of $\beta \mathcal{H}_\tau u'_g \mathcal{G}'_K < 0$; and
5. two additional induced effects which occur via the above-mentioned decrease in next period's tax rate, $\frac{d\tau'}{d\tau} = -\mathcal{H}_\tau \frac{\mathcal{H}'_k}{\mathcal{H}'_\tau} < 0$; this effect
 - (a) raises next period's consumption, which results in a change in next period's utility by $\beta \mathcal{H}_\tau (-\frac{\mathcal{H}'_K}{\mathcal{H}'_\tau}) u'_c [-\mathcal{H}'_\tau - \mathcal{G}'_\tau] > 0$ (assuming $\mathcal{H}_\tau + \mathcal{G}_\tau > 0$), and

- (b) lowers next period's government spending, which leads to a utility change of the amount $\beta\mathcal{H}_\tau(-\frac{\mathcal{H}'_K}{\mathcal{H}'_\tau})u'_g\mathcal{G}'_\tau < 0$.

In our numerical work below, we derive (steady-state) values to these different terms, allowing us to determine the effects in the GEE are quantitatively important and which are not.

3.1.2 The public-finance version

The GEE can be rewritten so that it is a linear combination of wedges. Rearranging terms we obtain the following equation.

$$\mathcal{G}_\tau \left[u_g - u_c \right] + \mathcal{H}_\tau \left[-u_c + \beta u'_c (1 + f'_K - \delta) \right] + \beta \mathcal{H}_\tau \left(\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \mathcal{G}'_\tau \right) \left[u'_g - u'_c \right] = 0. \quad (11)$$

Three terms in brackets appear: these are the three different “wedges” that are affected by the change in the current tax rate. Note that only wedges in the current and in the next period appear, even though this intertemporal economy has wedges in every period: again, envelope theorems imply that future wedges are handled optimally and hence can be ignored in comparing marginal costs and benefits of a current tax increase.

How are the different distortions traded off against each other? First, an increase in the tax rate influences the gap between u_g and u_c . This gap, which would be zero with lump-sum taxes since private and public goods are perfect substitutes in production, must be positive since it is costlier to provide G than C here. That is, a tax increase, by increasing G , makes this gap smaller.

Second, since the tax increase leads to a decrease in savings, the intertemporal distortion is affected. The second bracket, $-u_c + \beta u'_c (1 + f'_K - \delta)$, actually equals $u'_c (f'_K - \delta) \tau'$ from the consumer's Euler equation: so long as the tax rate next period is positive, the marginal utility of consumption today is too low (because savings are

too low). Thus, the decrease in savings resulting from an increase in current taxes will be detrimental: it increases the intertemporal distortion further.

Third, the lowered savings will lead to changes in the provision of public goods next period and it will thus influence the gap between the marginal utilities of public and private goods in that period. The channels are two: lowered capital directly lowers G' and it also induces a decrease in next period's tax rate, which also lowers G' : $\frac{dG'}{d\tau} = \mathcal{H}_\tau(\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau}\mathcal{G}'_\tau) < 0$. Thus, this effect is a negative: next period's distortion is made larger.

In sum, we are weighing one positive effect of increasing the current tax rate—it increases the amount of public goods, thus decreasing the wedge between public and private goods—against two negative ones: it increases the same wedge next period, and it also increases the intertemporal wedge. Not all wedges can be zero, because the optimal provision of public goods— $u_c = u_g$ —demands a positive tax rate, which necessarily makes the intertemporal distortion nonzero. This result is perhaps surprising: the use of the income tax seems nondistortionary in this model from the perspective of the current government: it is like a lump-sum tax. Nevertheless, the government *does not* tax at the high rates that would be necessary to deliver (statically) optimal public-goods provision! This is because the government finds it in their interest to leave more resources than that in the hands of the private sector: some of those resources will be saved, and this will help alleviate the intertemporal distortion. At an optimum, an optimizing government makes sure that a marginally decreased current public-goods-provision wedge is exactly counterbalanced by increases in the other wedges.

3.1.3 A sequential formulation

By Bellman's principle, it follows that we can alternatively characterize the problem of the government as one where it chooses a policy sequence, $\{\tau_t\}_{t=0}^\infty$, to solve the

following sequential problem:

$$\begin{aligned} \max_{\{\tau_t, K_{t+1}\}_{t=0}^{\infty}} \quad & \sum_{t=0}^{\infty} \beta^t u(\mathcal{C}(K_t, K_{t+1}, \tau_t), \mathcal{G}(K_t, \tau_t)) \\ \text{subject to} \quad & \\ & K_{t+1} = \mathcal{H}(K_t, \tau_t). \end{aligned} \tag{12}$$

By definition, this sequential problem has a stationary, or recursive, structure: the period objective involves K_t , K_{t+1} , and τ_t , and the constraint expresses K_{t+1} as a function of τ_t and K_t . That is, it is of the same form as, say, the planning problem of the neoclassical growth model; there, the control variable is c_t , and here it is τ_t . The time-inconsistency problem is apparently not present: variables chosen by the government at $t + 1$ do not affect those variables chosen at t .

Problem (12), however, does not correspond to the decision problem the government is actually facing, because it features a different feasible set for taxes: it does allow the government full power to choose any $\{\tau_t\}_{t=0}^{\infty}$. In contrast, in our time-consistent equilibrium, the government at time t only has a one-dimensional way of affecting future taxes: by the choice of τ_t , which via \mathcal{H} influences K_{t+1} , which in turn via Ψ affects τ_{t+1} , and so on. But despite being formally different problems, the equilibrium sequence chosen by the government in our Markov-perfect equilibrium will also solve this problem, and vice versa!

How is it possible that, given the ability to choose any tax sequence, the government would choose tax rates satisfying $\tau_t = \Psi(K_t)$? Problem (12) does not allow the government to choose any sequence of capital stocks: the capital stocks are given by \mathcal{H} . In contrast, the Ramsey problem described in Section 2.2 leaves more freedom for the capital stocks. In particular, the constraint in the Ramsey problem, which restricts capital through the individual's first-order conditions for saving, allows K_1 and τ_0 to be selected independently. Here, in contrast, once τ_0 is pinned down, so is K_1 , from $K_1 = \mathcal{H}(K_0, \tau_0)$. The reason why the Ramsey problem allows more freedom, intuitively, is that once the current tax has been set, savings can still be influenced by future taxes, which you control. Problem (12) instead implicitly *restricts the private sector's expectations* about future events to be consistent with those in our time-consistent equilibrium.

Problem (12) is not autonomous: it is still part of a fixed-point problem in Ψ . How does Ψ enter in the problem? As just pointed out, it is incorporated into \mathcal{H} : the shape of this function reflects the expectations of future taxes (as described above in (3), the private sector's Euler equation.

Finally, the sequential formulation makes it evident that one can derive the GEE using standard variational methods: fix the state variable this period, K , and two periods hence, K'' , and vary the controls in the current and next period (τ and τ' , respectively) in order to attain the highest possible utility over these two periods. Straightforward differentiation of (12), as for the standard growth model, delivers the GEE immediately. This way of deriving the GEE is perhaps more direct than the route we followed above.

3.2 An alternative equilibrium definition

It is possible to provide a more compact definition of equilibrium. We state this definition because its compactness actually matters for computation (as we will comment on later in detail) and because it is more closely connected to the Ramsey problem of Section 2.2. The compactness is accomplished by not allowing a distinction between the government and the private sector. The government thus chooses both K' and τ directly, with associated equilibrium mappings $h(K)$ and $\psi(K)$.

The problem that the government solves—maximize utility subject to the Euler equation of the households—reads as follows:

$$\begin{aligned}
v(K) &= \max_{K', \tau} \{u(\mathcal{C}(K, K', \tau), \mathcal{G}(K, \tau)) + \beta v(K')\} \\
&\text{subject to} \\
&u_c(\mathcal{C}(K, K', \tau), \mathcal{G}(K, \tau)) = \\
&\beta u_c(\mathcal{C}(K', h(K'), \psi(K')), \mathcal{G}(K', \psi(K')))) \cdot \{1 + [1 - \psi(K')][f_K(K') - \delta]\}.
\end{aligned} \tag{13}$$

Thus, K' is left as a choice variable and the restriction that it be consistent with private-sector behavior is not captured through a function \mathcal{H} but instead by including the consumer's Euler equation explicitly. Moreover, this equation is not a functional equation: unlike (3), which had to hold for all K and τ , thus defining \mathcal{H} , here the Euler

equation is a restriction on K' and τ (given any K). Part of this restriction, which sets it apart from the restriction faced by the Ramsey planner, is that K'' as well as τ' , which enter on the right-hand side of the Euler equation, have to $h(K')$ and $\psi(K')$, respectively.

Problem (13) is, of course, also a fixed-point problem: taking as given expectations about future behavior as captured by h and ψ , current optimal behavior has to reproduce these functions. This formulation is obviously closer to the Ramsey formulation: the difference is that future savings and tax choices here are restricted by h and ψ whereas they are free in the Ramsey problem.

It is straightforward to see that our two equilibrium definitions are equivalent: the function \mathcal{H} is *defined* to solve the Euler equation above with the only difference that what appears on the right-hand side is $\mathcal{H}(K', \Psi(K'))$, not $h(K')$, and $\Psi(K')$, not $\psi(K)$. Thus, if \mathcal{H} and Ψ constitute an equilibrium according to our original definition, then h and ψ defined by $h(K) \equiv \mathcal{H}(K, \Psi(K))$ and $\psi \equiv \Psi$ are an equilibrium as defined in this section. Conversely, if h and ψ satisfy the above equilibrium definition, then one can define \mathcal{H} from the Euler equation above (solve for K' as a function of K and τ) and set $\Psi \equiv \psi$ and it is evident that \mathcal{H} and Ψ are an equilibrium according to our original definition.

It is also possible to derive the GEE from (12). However, as perhaps is evident, it will lead to an equation which is very long, because both τ and, especially, K' , appear in a large number of places. Among the many terms, both ψ' and h' —the derivatives of the policy rules—will appear in this first-order condition.

One can simplify matters by summarizing the Euler equation by

$$\eta(K, \tau, K') = 0, \quad (14)$$

where η is defined as the left-hand side of the restriction in (13) minus the right-hand side. The GEE, then, becomes (after deriving the first-order condition and utilizing the envelope theorem)

$$0 = -(u_c C_\tau + u_g G_\tau) \eta_{K'} - \eta_\tau \left[u_c C_{K'} + u'_c C'_K + u'_g G'_K - (u'_c C'_\tau + u'_g G'_\tau) \frac{\eta'_K}{\eta'_\tau} \right]. \quad (15)$$

Noting that \mathcal{H} is defined by

$$\eta(K, \tau, \mathcal{H}(K, \tau)) = 0$$

for all (K, τ) , we can differentiate with respect to K and τ and obtain, respectively,

$$\mathcal{H}_K = -\frac{\eta_K}{\eta_{K'}}$$

and

$$\mathcal{H}_\tau = -\frac{\eta_\tau}{\eta_{K'}}.$$

Dividing the GEE—equation (15)—by $\eta_{K'}$ and rearranging, we obtain our original GEE—equation (9). It, along with $\eta(K, \psi(K), h(K)) = 0$, has to hold for all K and constitute our definition of time-consistent equilibrium here.

3.2.1 Strategic policy: does the current government manipulate its successors?

The dynamic game played between governments involves a disagreement: the current government would like to see the next government choose a lower tax on income, τ' , than it ends up choosing. Does this mean that the current government attempts to “manipulate” the next government in its tax choice? It could influence τ' through its influence on saving, K' . Suppose, for example, that $\tau' = \Psi(K')$ is increasing. Then the current government might see a reason to increase τ a little extra, so as to decrease K' and thereby decrease τ' : it could influence the tax choice next period through savings.

Our GEEs, however, do not directly contain the derivative of the tax policy rule Ψ , as one might think it would. In fact, from our arguments earlier, and the very fact that the government’s problem can be written recursively, the successive governments actually agree in one important dimension: *given the value for current savings*, they agree on how to set next period’s taxes. That is why the derivative of Ψ does not appear directly in the government’s first-order conditions. It appears indirectly, as a determinant of \mathcal{H}_τ . But this appearance does not reflect strategic behavior; rather, it simply captures how the effects on private-sector savings of a current change in τ depends on how those extra savings will alter next period’s tax rate. That is, \mathcal{H}_τ reflects how a current tax change influences the *expectations* of private agents, and

therefore their savings. More precisely, if the tax rate today is changed, how much extra (or less) capital is saved— \mathcal{H}_τ —depends on how the determination of the future tax rate is perceived by the private sector.

To illustrate the role of Ψ in the determination of the savings response, let us compare the kind of government we model to a “myopic” alternative: a myopic government does not realize that their current taxation behavior influences future taxes. Suppose that the time-consistent equilibrium has Ψ as an increasing function: the higher the savings today, the higher the tax rate will be next period. In contrast, the myopic government perceives $\Psi(K)$ to be constant. How, then, would the myopic government’s first-order condition look? The answer is that it would look the same, with the one difference that \mathcal{H}_τ would be a different number: in terms of our compact equilibrium definition, we have $\mathcal{H}_\tau = -\frac{\eta_\tau}{\eta_{K'}}$, and here the denominator (but not the numerator) depends on the derivate of Ψ . Assuming that $u(c, g)$ is additively separable, that $\eta_\tau > 0$, and that $\eta_{K'} > 0$, one observes that if the change in the future tax is ignored, $\eta_{K'}$ could be too high—because of the lowered consumption, and therefore increased future marginal utility value of savings, implied by the higher future tax rate—or too low—because of the lower net-of-tax return from future savings. That is, a myopic government would misperceive \mathcal{H}_τ , but whether this leads to lower or higher equilibrium taxes is a quantitative question.

It is interesting to contrast the present finding—that the strategic manipulation of future players does not really take place here—with results from another dynamic game with disagreements: the individual savings problem under time-inconsistent (quasi-geometric) preferences (see, e.g., Strotz (1956), Laibson (1997), or Krusell, Kuruşçu, and A. (2000)). This problem is modeled as a game between successive selves who have conflicting discounting: the current self always places a higher relative weight on current (as opposed to future) consumption than did any of his previous selves. In a differentiable Markov-perfect equilibrium of this game, the derivative of the savings function appears directly in the Euler equation: when evaluating the marginal benefits of an additional unit saved, the current self sees a value in sending resources forward because any extra income will partly be saved by the next self and that next self, the current self thinks, saves too little. That is, because the envelope theorem does not apply—due to the disagreement over savings—the *response* of the next self matters in

the marginal evaluations made by the current self. In the present model, in contrast, since there is agreement on all decisions in the next period given a value for the capital stock next period, the corresponding response—the derivative of Ψ —is not directly present.

4 An extension: valued leisure

Suppose now that leisure is valued: we assume that utility is given by

$$\sum_{t=0}^{\infty} \beta^t u(c_t, 1 - \ell_t, g_t).$$

We continue assuming that the tax base is total income. Our equilibrium definition works as before, but one more element is required: we need to describe the equilibrium labor response to (K, τ) . The relevant mapping is $\mathcal{L}(K, \tau)$, which is obtained from the consumer's first-order condition for the labor-leisure choice. Thus,

$$u_c(\mathcal{C}(K, \mathcal{H}(K, \tau), \tau), 1 - \mathcal{L}(K, \tau), \mathcal{G}(K, \tau)) \cdot f_L(K, \mathcal{L}(K, \tau)) (1 - \tau) = \\ u_\ell(\mathcal{C}(K, \mathcal{H}(K, \tau), \tau), 1 - \mathcal{L}(K, \tau), \mathcal{G}(K, \tau)) \quad (16)$$

for all (K, τ) and the first-order condition for savings (which now contains a leisure argument, but which we will not restate) *jointly* define the functions $\mathcal{H}(K, \tau)$ and $\mathcal{L}(K, \tau)$.

The equilibrium conditions now include three functional equations: the private sector's first-order conditions for labor and savings and the government's first-order condition. We go straight to the latter—to the GEE—which can be derived with the same procedure as above. It reads

$$\mathcal{L}_\tau \left[u_c f_L - u_\ell \right] + \mathcal{G}_\tau \left[u_g - u_c \right] + \mathcal{H}_\tau \left[-u_c + \beta u'_c (1 + f'_K - \delta) \right] + \\ \beta \mathcal{H}_\tau \left\{ \mathcal{L}'_K \left[u'_c f'_L - u'_\ell \right] + \mathcal{G}'_K \left[u'_g - u'_c \right] - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \left(\mathcal{L}'_\tau \left[u'_c f'_L - u'_\ell \right] + \mathcal{G}'_\tau \left[u'_g - u'_c \right] \right) \right\} = 0 \quad (17)$$

for all K (again, the arguments of the functions are suppressed for readability). We see a new wedge appearing: $u_c f_L - u_\ell$, in the current period as well as in the next.

This wedge, which equals $u_c\tau$, must be positive so long as public goods are provided ($\tau > 0$). A current tax increase will now decrease labor supply (presumably) and thus increase this intratemporal distortion. Similarly, there will be repercussions through lowered savings on the same wedge in the future, in parallel with the induced effects on future savings.

This formulation follows the public finance tradition of characterizing optimal taxes as combinations of wedges. It is informative to consider other tax bases. Let us first look at the wedge version of the GEE that would have arisen in an economy with lump-sum taxes. It is given by

$$\mathcal{G}_\tau \left[u_g - u_c \right] + \beta \mathcal{H}_\tau \left(\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \mathcal{G}'_\tau \right) \left[u'_g - u'_c \right] = 0. \quad (18)$$

Here, a policy that sets marginal utility of government expenditures equal to that of private consumption satisfies equation (18) and hence is an equilibrium policy. Meanwhile, to satisfy the private Euler equations with no distortionary taxes, labor supply will be Pareto optimal as well in the Markov equilibrium. Thus the Markov equilibrium is Pareto optimal. When the first best can be achieved, time inconsistency is no longer a problem, and the GEE shows that.

On the other hand, consider the case where only (net) capital income can be taxed. Then the GEE becomes

$$\mathcal{G}_\tau \left[u_g - u_c \right] + \mathcal{H}_\tau \left[-u_c + \beta u'_c (1 + f'_K - \delta) \right] + \beta \mathcal{H}_\tau \left(\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \mathcal{G}'_\tau \right) \left[u'_g - u'_c \right] = 0. \quad (19)$$

Notice that this is *the same* GEE as in the model without leisure. This does not mean that the equilibrium tax rate is the same—the remaining equilibrium equation elements are different. As in the case without valued leisure, it will not be optimal to go all the way to (statically) optimal public-goods provision.

We now look at the GEE that results when there are only labor taxes. The GEE becomes

$$\mathcal{L}_\tau \left[u_c f_L - u_\ell \right] + \mathcal{G}_\tau \left[u_g - u_c \right] + \beta \mathcal{H}_\tau \left\{ \left(\mathcal{L}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \mathcal{L}'_\tau \right) \left[u'_c f'_L - u'_\ell \right] + \left(\mathcal{G}'_K - \frac{\mathcal{H}'_K}{\mathcal{H}'_\tau} \mathcal{G}'_\tau \right) \left[u'_g - u'_c \right] \right\} = 0. \quad (20)$$

Is the lack of commitment binding in the economy where labor income is taxed: is the corresponding Ramsey equilibrium time-consistent? Even though the labor decision is a static one, the answer is no. The reason is that Ramsey policy maker takes into account the fact that a tax increase at t not only lowers labor supply at t but raises it at $t - 1$, due to an income effect. A Markov policy maker treats the latter as a bygone and does not take it into account. We do not display the first-order condition from the Ramsey problem here, but it is straightforward to verify it does not coincide with equation (20).

5 An example with a closed-form solution

We now use a parametric example to illustrate some of the results in our model. We first consider a model without leisure and finite horizon; this economy allows us, among other things, to discuss the relation between any time-consistent equilibrium in the infinite-horizon economy and the limit of finite-horizon equilibria. We then look at the case of leisure, but here we restrict attention to the infinite-horizon case. This model serves as a benchmark case for our quantitative section below: it has all the main ingredients we consider there, but one cannot use it for quantitative purposes as there is full depreciation. For the first economy, we look at capital income taxation, and for the second we use a general income tax. With any form of proportional taxation in these economies we obtain closed-form solutions.

5.1 A finite-horizon model without leisure

Suppose the representative agent has preferences represented by

$$\sum_{t=0}^T \beta^t [\ln C_t + \gamma \ln G_t].$$

The resource constraint is given by

$$C_t + K_{t+1} + G_t = K_t^\theta,$$

reflecting the assumption that there is full depreciation, and the government's budget constraint is

$$\tau_t \theta K_t^\theta = G_t,$$

which assumes that only capital income is taxed, that the government cannot issue debt, and that depreciation is not tax deductible. Initial capital $K_0 > 0$ is given and $K_{T+1} = 0$.

For this economy, we will show that tax rates do not depend on the state variable—the Ψ function is constant. To establish benchmarks, consider first the extreme cases $T = 0$ and $T = \infty$. If $T = 0$ it is easy to see that

$$\tau = \frac{\gamma}{\theta(1 + \gamma)};$$

this follows from the condition $u_g = u_c$. On the other hand, if $T = \infty$, then the tax is lower. To find what it is, we go through the following steps. First, suppose taxes in every period are τ . Then the savings function is

$$K' = \beta\theta(1 - \tau)K^\theta.$$

The value of being born into such an economy endowed with initial capital K is

$$v(K; \tau) = A \ln K + B(\tau).$$

where

$$A = \frac{\theta(1 + \gamma)}{1 - \beta\theta}.$$

Now suppose the current tax rate is τ and all subsequent tax rates are equal to τ' . Then the equilibrium savings function— \mathcal{H} —is given by

$$K' = \beta\theta \frac{(1 - \tau')(1 - \theta\tau)}{1 - \theta\tau'} K^\theta.$$

We are now in a position to state the current government's maximization problem.

$$\max_{\tau} \left\{ \ln C + \gamma \ln G + \beta \frac{\theta(1 + \gamma)}{1 - \beta\theta} \ln K' + \beta B(\tau') \right\}.$$

Substituting in the values for C , G and K' and setting the derivative with respect to τ equal to zero, it turns out that the optimal current tax rate τ is independent of the future tax rate(s) τ' . We obtain

$$\tau = \frac{\gamma(1 - \beta\theta)}{\theta(1 + \gamma)}.$$

So with $\Psi(K) = \frac{\gamma(1-\beta\theta)}{\theta(1+\gamma)}$ for all K and $\mathcal{H}(K, \tau) = \beta\theta^{\frac{\theta(1+\gamma)-\gamma+\gamma\beta\theta}{1+\gamma\beta\theta}}(1-\theta\tau)K^\theta$ for all (K, τ) , we satisfy our conditions for a time-consistent equilibrium. These functions are smooth; it is straightforward to obtain their derivatives and show that the GEE is satisfied.

Now consider the intermediate case $0 < T < \infty$. In this case, the assumption that tomorrow's tax rate is the same as subsequent tax rates will not be appropriate. However, it turns out that the slope of the value function does not depend on tax rates, only on the time left until T . This is because private consumption, public consumption and savings in period t are all proportional to K_t^θ .

Thus we may write

$$v_t(K) = A_t \ln K + B(\tau_t, \tau_{t+1}, \dots, \tau_T)$$

where A_t satisfies the difference equation

$$\begin{cases} A_{t+1} &= (\beta\theta)^{-1}A_t - \frac{(1+\gamma)}{\beta} \\ A_{T+1} &= 0. \end{cases}$$

Meanwhile, savings, and hence also consumption, is proportional to $1 - \theta\tau_t$. On the other hand, public consumption is proportional to τ_t itself. This means that the first-order condition for the optimal current tax becomes

$$\frac{-\theta}{1-\theta\tau_t} + \frac{\gamma}{\tau_t} - \frac{\beta\theta A_{t+1}}{1-\theta\tau_t} = 0.$$

The solution is

$$\tau_t = \frac{\gamma}{\theta(1+\gamma) + \beta\theta A_{t+1}}.$$

Now we solve the difference equation for A_t . We get

$$A_t = \frac{\theta(1+\gamma)}{1-\beta\theta} [1 - (\beta\theta)^{T-t+1}]$$

and hence

$$\tau_t = \frac{\gamma}{\theta(1+\gamma)} \cdot \frac{1-\beta\theta}{1 - (\beta\theta)^{T-t+1}}.$$

It now seems clear that the time-consistent equilibrium tax rate we find above for the infinite-horizon economy is the limit of finite-horizon equilibria. But wait a minute!

What if the expression for $1 - \tau_t$ is less than or equal to zero for any t ? This occurs whenever $\frac{\gamma}{\theta(1 + \gamma)}$, the tax rate in the economy with only one period—the last-period, or period- T , economy, is greater than or equal to 1, since the tax sequence is decreasing in t . If the optimal tax rate in the last period is above 100%, then savings in the period before that will be zero. This means that period $T - 1$ is like period T : it is as if there is no future, since K_T will be set to zero. Again, taxes in that period will be set to be confiscatory, and so on. The end result is that no matter how long the horizon is in such an economy, the limit of finite-horizon equilibria is that the first period looks like a static economy, and every period after that is utter misery with $C = K = G = 0$.

One lesson we learn from this example is that for some parameter values there are two time-consistent (differentiable) equilibria for our infinite-horizon economy. One of these is a collapsed economy, with $\mathcal{H}(K, \tau) = 0$ for all (K, τ) and $\Psi(K) = \frac{\gamma}{\theta(1 + \gamma)} > 1$ for all K ; it is the limit of finite-horizon equilibria. The other one, which we computed above, is a more healthy economy with a constant tax rate $\frac{\gamma(1 - \beta\theta)}{\theta(1 + \gamma)} < 1$ and economic activity in every period (this economy converges to a steady state with positive capital and consumption).⁴ This second equilibrium relies on “optimistic” expectations.

5.2 Infinite horizon, valued leisure, and a general income tax

The environment is as follows. A representative agent has preferences represented by

$$\sum_{t=0}^{\infty} \beta^t [\alpha \ln C_t + (1 - \alpha) \ln \ell_t + \gamma \ln G_t]$$

where ℓ_t is the fraction of time spent not working and G_t is the quantity of government-provided goods. The fraction of time spent working is denoted by $H_t = 1 - \ell_t$.

The aggregate resource constraint is given by

$$K_{t+1} + C_t + G_t = K_t^\theta H_t^{1-\theta}$$

which means that the depreciation rate is 100 percent. The government cannot issue debt and taxes income at a proportional rate τ . Depreciation is not tax deductible. Its

⁴Clearly, if β or θ are close enough to zero also this economy will be a collapsed one.

budget constraint is

$$G_t = \tau_t K_t^\theta H_t^{1-\theta}.$$

Regardless of current and future tax rates and the current capital stock, labor supply is

$$H = \frac{\alpha}{\alpha + (1 - \alpha)(1 - \beta\theta)}.$$

We also have the following competitive equilibrium law of motion for capital.

$$K' = \beta\theta(1 - \tau)K^\theta H^{1-\theta}.$$

By the government budget constraint, we have of course

$$G = \tau K^\theta H^{1-\theta}.$$

The value of being born in an economy with initial capital stock K and a constant tax rate τ is given by

$$v(K) = \frac{(\alpha + \gamma)\theta}{1 - \beta\theta} \ln K + B(\tau)$$

where $B(\tau)$ is some function whose exact properties turn out not to be important from our point of view. The current government's maximization problem, which turns out to be independent of future tax rates, can be written as

$$\max_{\tau} \left\{ \alpha \ln C(K, \tau) + (1 - \alpha) \ln(1 - H) + \gamma \ln G(K, \tau) + \beta \frac{(\alpha + \gamma)\theta}{1 - \beta\theta} \ln K'(K, \tau) + \beta B(\tau') \right\} \quad (21)$$

where the function C can be derived from the function K' in an obvious way.

The resulting tax outcome for this economy is

$$\tau = \Psi(K) = \frac{\gamma(1 - \beta\theta)}{\alpha + \gamma}.$$

This tax function is accompanied by a savings function that takes the form $\mathcal{H}(K, \tau) = \beta\theta(1 - \tau)K^\theta \left(\frac{\alpha}{\alpha + (1 - \alpha)(1 - \beta\theta)} \right)^{1-\theta}$ and by a labor supply function of the form

$\mathcal{L}(K, \tau) = \frac{\alpha}{\alpha + (1 - \alpha)(1 - \beta\theta)}$. Again, it is straightforward to verify that these three functions satisfy the GEE.

An interesting property of this equilibrium is that it coincides with the Ramsey allocation. That could be observed already above when it was noted that neither the savings nor the labor supply decisions depend on future tax rates. This result, which of course depends on our functional-form assumptions, also relies on the assumption of a general income tax; with another tax base—a pure capital income tax or a pure labor income tax—there are income effects of future taxation that influence current decisions, and the Ramsey equilibrium is not time-consistent.

6 Optimal policy for actual economies

We proceed next to look at numerical solutions for a selected set of economies with some aggregate statistics that resemble those of the United States postwar economy. For the sake of comparison we also provide the optimal policy under the first best (lump-sum taxation) allocation and those implied by a benevolent government that has access to commitment but not to a technology to save resources, that is, the Ramsey equilibrium given a period-by-period balanced budget constraint.⁵

We specify the per-period utility function of the CES class as

$$u(c, \ell, g) = \frac{\left[(1 - \alpha_p) (\alpha_c c^\rho + (1 - \alpha_c) \ell^\rho)^{\Psi/\rho} + \alpha_p g^\Psi \right]^{\frac{1-\sigma}{\Psi}} - 1}{1 - \sigma}. \quad (22)$$

This function reduces to a separable function with constant expenditure shares when $\sigma \rightarrow 1$, $\rho \rightarrow 0$, and $\Psi \rightarrow 0$, yielding

$$u(c, \ell, g) = (1 - \alpha_p) \alpha_c \ln c + (1 - \alpha_p)(1 - \alpha_c) \ln \ell + \alpha_p \ln g \quad (23)$$

Meanwhile, the production function is a standard Cobb-Douglas function with capital share θ : $f(K, L) = A \cdot K^\theta L^{1-\theta}$.

⁵Stockman (1998) for a Ramsey government and Klein and Ríos-Rull (1999) for a government without access to a commitment technology perform a quantitative analysis of optimal taxation (labor and capital income taxes) for exogenous public expenditures under a period by period balanced budget constraint.

Our parameterization of the baseline economy is also standard. We calibrate the baseline model economy, which is the one with only labor taxes, to have some statistics within the range of U.S. data in the lack-of-commitment economy. So we set the share of GDP that is spent by the government to be slightly under 20%, the capital share to 36%, the investment-to-output ratio to a little over 20%, hours worked to about one fourth of total time, and the capital-to-output ratio to about 3. These choices are common in the macroeconomic literature.

We choose the baseline economy to have logarithmic utility which makes preference separable (making cross derivatives zero).⁶ We report the values of the parameters that implement our choices in Table 1.

Parameter Values		
$\theta = 0.36$	$\alpha_c = 0.30$	$\alpha_p = 0.13$
$\beta = 0.96$	$\delta = 0.08$	$\rho = 0$
$\Psi = 0$	$\sigma = 1.0$	

Table 1: Parameterization of the Baseline Model Economy

6.1 Labor income taxes

We now look at the steady states of the baseline economy under three different benevolent governments that we label Pareto, Ramsey, and Markov. These labels, respectively, refer to: a government with commitment and access to lump-sum taxation (Pareto); a government restricted by a period-by-period balanced-budget constraint and to the use of labor income taxation, both one with access to a commitment technology (Ramsey) and one which does not have access to such commitment technology (Markov, because

⁶When, in addition, the depreciation rate is 100%, this economy allows a closed-form solution to all equations, including the GEE. The key functional-form insight here is that all the first-order conditions become log-linear.

we look at the Markov equilibrium). Table 2 reports the steady-state allocations of these three economies.

Labor taxes, Endogenous g			
Steady State Statistic	Type of Government		
	Pareto	Ramsey	Markov
Y	1.000	0.700	0.719
K/Y	2.959	2.959	2.959
C/Y	0.509	0.509	0.573
G/Y	0.254	0.254	0.190
C/G	2.005	2.005	3.017
L	0.350	0.245	0.252
τ	–	0.397	0.297

Table 2: Baseline Model Economy

The absence of capital income taxes in all economies ensures that the steady-state interest rate is equated to the rate of time preference, yielding an equal capital-to-output ratio in all economies. Comparing the Pareto and the Ramsey economy, we get a glimpse of the role of distortionary labor taxation. The Pareto economy delivers the optimal allocation while the Ramsey economy has a distortionary tax that discriminates against produced goods and in favor of leisure. As a result, leisure is significantly higher in the Ramsey economy than in the Pareto economy, and because of this and the equal rate of return, the steady-state stock of capital and output are much lower in the Ramsey economy. However, the ratio between private and public consumption is the same in both economies given that this margin is undistorted. This latter feature is a special implication of the functional form that we have chosen and it relies on preferences being separable in all three goods and on being of the CRRA class with respect to consumption.⁷

⁷This is a simple implication of the first-order conditions of the Ramsey problem when written in the primal form.

When we look at the behavior of the Markov economy, we see two things: first, qualitatively, the distortion introduced by the tax on labor is also present in this economy, inducing more leisure and less consumption (both private and public) than in the Pareto economy; and second, the ratio between private and public consumption is not the same as in the other economies (where it was equal to the relative share parameter in preferences). Recall that from equation (20) the optimal policy of the Markov case amounted to striking a balance between achieving the first best in terms of equating the marginal utility of the private and public good and the distortion that the labor tax induces on the leisure–private consumption margin. This balance does not imply setting the margin between the public and the private good to zero. Indeed, the term $u_g - u_c$ is positive in the Markov case, making the second term of equation (20) positive and the first negative. The difference with the Ramsey case can perhaps be best described by the fact that the Ramsey policy maker takes into account the fact that a tax hike at t not only lowers labor supply at t but raises it at $t - 1$. In contrast, a Markov policy-maker treats the latter as a bygone.

6.2 Capital income taxation

Table 3 shows the steady state when the only available tax is the capital income tax.

This tax is in general very distortionary. The Ramsey government understands this and, therefore, reduces future taxes so as to mitigate the distortionary effect. However, since no other tax base is available here, the result is that the ratio of private to public consumption is much lower than in the unconditional first best. The Markov government, however, does not see the current tax as distortionary at all, as capital is already installed when the government chooses the tax rate: capital is inelastically supplied.

The Markov government, however, understands that the government that follows one period later will distort the allocation significantly, and is therefore willing to attempt to transfer resources into the future to increase future consumption. For this reason, it does not tax capital so as to set the private-to-public consumption ratio at the first-best level. The ability of the Markov government to influence the future choices

Capital Income taxes, Endogenous g			
Steady State Statistic	Type of Government		
	Pareto	Ramsey	Markov
Y	1.000	0.588	0.488
K/Y	2.959	1.734	1.193
C/Y	0.509	0.712	0.690
G/Y	0.254	0.149	0.215
C/G	2.005	4.779	3.211
L	0.350	0.278	0.255
τ	–	0.673	0.812

Table 3: Baseline Model Economy; Separable utility in logs.

is of course smaller than that of the Ramsey government, and as a result its capital tax rate is higher and capital and output are lower.

Another feature of this case that we find interesting is that leisure is the lowest in the Pareto case, even when there is no tax on leisure. Our understanding of the reasons for this goes as follows. With the preferences of this model economy, in any market implementation, the household’s choice of leisure can be decomposed into two parts. One part is what it would choose if all income were labor income—it equals $(1 - \alpha_c)$ exactly, independently of the wage (that in this case is 0.7). The other part comes from the amount of additional income that the household has, so that leisure is increasing in that additional income. In the Pareto economy, the lump-sum tax levied is larger than the amount of capital income, inducing the household to enjoy less leisure than 0.7, while in all the other economies, the after-tax capital income is always positive, which accounts for why workers enjoy leisure of more than 0.7 in those economies.

6.3 Taxes on total income

With respect to the case of a tax on total income, a couple of points are worth stressing.

Total Income Taxes, Endogenous g			
Steady State Statistic	Type of Government		
	Pareto	Ramsey	Markov
Y	1.000	0.669	0.693
K/Y	2.959	2.527	2.649
C/Y	0.509	0.532	0.587
G/Y	0.254	0.265	0.201
C/G	2.005	2.005	2.928
L	0.350	0.256	0.258
τ	–	0.334	0.255

Table 4: Baseline Model Economy; Separable utility in logs.

First, the Ramsey government can set the ratio of private to public consumption to its unconditionally optimal level. Due partly to the special nature of the preferences used in this model economy, the distortions that affect the intertemporal margin and the consumption leisure margin do not affect the private-to-public-consumption margin. From the point of view of the Markov government, however, this is not the case. An uncommitted policy maker does not take into account that today's taxes increase yesterday's incentives to work, and in addition it wishes to increase savings by taxing less today, and these effects induce a smaller government sector. This result is perhaps surprising because one might have guessed that a Markov government, which views its taxes as less distortionary than does the Ramsey government, would tax more.

In addition to the comparisons that we have performed between the three taxing technologies that the government may have access to (and that yield the Pareto, Ramsey, and Markov cases), for each of the tax tools, we should also compare the allocations for the Markov case across tax instruments.

From the point of view of the Markov government, taxing capital is not distortionary since it is already installed and hence is like a lump-sum tax. On the other hand, the tax base is quite small, as capital income is much smaller than labor income.⁸ On the other hand, labor taxes are distortionary but its base is larger. Finally, total income taxes have the highest tax base and they are as distortionary as the labor income tax rate for the same tax rate, or less distortionary for the same revenue.

With respect to tax outcomes, first, as should have been expected, the larger the role of capital income taxes (which implies an ordering with capital income first, followed by total income and last labor income), the lower the stock of capital, and hence the lower will output be. The differences are large. Second, hours worked are actually varying very little across environments. Third, perhaps the most surprising feature that we obtain is that the ratio of private consumption to public consumption is the highest in the capital-tax economy. This is very surprising, since we should expect that the government, since it considers taxes to be non-distortionary, would allocate current resources optimally across these goods, thus equating the marginal utility of public and private consumption (which is what the Pareto government does). The reason why this does not occur is that the government in the capital-income economy understands that the next government will tax capital heavily (more heavily, indeed, than what this government would like), and in an effort to move resources into the future it thus sacrifices current public consumption. Note also that this effect is non-linear in that the private-to-public consumption ratio closest to the first best is that of the total income-tax economy.

7 Conclusion

In this paper we have characterized the set of functional equations that are required for characterizing Markov equilibria in an environment where a benevolent government that does not have access to commitment sets tax rates to finance a public good. We have shown how the problem of the government has a sequential structure that can be written as if it had access to commitment by posing the behavior of the private sector—

⁸Note that because the tax base excludes depreciation, the tax base of a capital income tax is not a constant fraction of GDP.

and future governments—in a specific way. This leads to a natural characterization of government behavior in terms of the first-order conditions of such a problem that we have called the GEE.

We have discussed some issues pertaining to the computation of such equilibria and we have found the solutions to a variety of parameterized model economies. We have compared those solutions to those that result from governments that have access to lump-sum taxation or to commitment and we have found that the implied taxes and allocations are very different depending on the environment in which the government lives. We leave for future research the characterization for an environment where there are explicit intertemporal links in the behavior of the government, i.e., debt.

We believe that the methods that we have developed for solving for our Markov equilibrium are quite general and can be applied to a larger set of environments than that of optimal fiscal policy studied here. Such environments may include optimal monetary policy, dynamic political economy, dynamic industrial organization issues (e.g., the durable goods monopoly, dynamic oligopoly), models with impure intergenerational altruism, and so on.

References

- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.
- BENOIT, J.-P., AND V. KRISHNA (1985): "Finitely Repeated Games," *Econometrica*, pp. 905–922.
- CHARI, V. V., AND P. J. KEHOE (1990): "Sustainable Plans," *Journal of Political Economy*, 98(4), 784–802.
- COHEN, D., AND P. MICHEL (1988): "How Should Control Theory Be Used to Calculate a Time-Consistent Government Policy?," *Review of Economic Studies*, 55, 263–274.
- CURRIE, D. A., AND P. L. LEVINE (1993): *Rules, Reputation and Macroeconomic Contracts*. Cambridge University Press, Cambridge.
- JUDD, K. L. (1998): *Numerical Methods in Economics*. MIT Press.
- KLEIN, P., AND J.-V. RÍOS-RULL (1999): "Time-Consistent Optimal Fiscal Policy," Mimeo, University of Pennsylvania.
- KRUSELL, P., B. KURUŞÇU, AND S. A. A. (2000): "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting," forthcoming, *Journal of Economic Theory*.
- KRUSELL, P., V. QUADRINI, AND J.-V. RÍOS-RULL (1997): "Politico-Economic Equilibrium and Economic Growth," *Journal of Economic Dynamics and Control*, 21(1), 243–272.
- KRUSELL, P., AND J.-V. RÍOS-RULL (1999): "On the Size of U.S. Government: Political Economy in the Neoclassical Growth Model," *American Economic Review*, 89(5), 1156–1181.
- KRUSELL, P., AND A. SMITH (2000): "Consumption-Savings Decisions with Quasi-Geometric Discounting," Manuscript.
- KYDLAND, F. E., AND E. C. PRESCOTT (1977): "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85(3), 473–91.
- LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, pp. 443–477.
- PHELAN, C., AND E. STACCHETTI (2000): "Sequential Equilibria in a Ramsey Tax Model," forthcoming in *Econometrica*.

- PHELPS, E. S., AND R. A. POLLAK (1968): "On Second-best National Saving and Game-equilibrium Growth," *Review of Economic Studies*.
- POLLAK, R. (1968): "Consistent Planning," *Review of Economic Studies*, 35, 201–208.
- STOCKMAN, D. R. (1998): "Balanced-Budget Rules: Welfare Loss and Optimal Policies," Mimeo, University of Delaware.
- STROTZ, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180.

Appendix

A Definition of the recursive competitive equilibrium

Here we present a formal definition of equilibrium given a policy.

Definition 1 A recursive competitive equilibrium given a government policy $\tau = \Psi(K)$, is a pair of functions for aggregate labor and aggregate next period capital stock, $L = \mathcal{L}(K, \tau)$ and $K' = \mathcal{H}(K, \tau)$, a value function for the representative household $\Omega(k, K, \tau)$, decision rules for labor and for savings for the representative household $l = \ell(k, K, \tau)$ and $k' = h(k, K, \tau)$, functions for factor prices $w(K, \tau)$ and $r(K, \tau)$ and a function for public consumption $\mathcal{G}(K, \tau)$ such that

1. Ω , h and ℓ solve the agents problem: for all (k, K, τ) ,

$$\Omega(k, K, \tau) = \max_{c, l, k'} u(c, l, g) + \beta \Omega(k', K', \tau') \quad (24)$$

$$\{\ell(k, K, \tau), h(k, K, \tau)\} \in \operatorname{argmax}_{c, l, k'} u(c, l, g) + \beta \Omega(k', K', \tau') \quad (25)$$

subject to

$$c + k' = k + [w(K, \tau) l + r(K, \tau) k] (1 - \tau) \quad (26)$$

$$g = \mathcal{G}(K, \tau) \quad (27)$$

$$K' = \mathcal{H}(K, \tau) \quad (28)$$

$$\tau' = \Psi(K); \quad (29)$$

2. the agent is representative: for all (K, τ) ,

$$\ell(K, K, \tau) = \mathcal{L}(K, \tau) \quad (30)$$

$$h(K, K, \tau) = \mathcal{H}(K, \tau); \quad (31)$$

3. factor prices are marginal productivities: for all (K, τ) ,

$$r(K, \tau) = f_K[K, \mathcal{L}(K, \tau)] - \delta \quad (32)$$

$$w(K, \tau) = f_L[K, \mathcal{L}(K, \tau)]; \quad (33)$$

4. and the government satisfies its budget constraint: for all (K, τ) ,

$$\mathcal{G}(K, \tau) = \tau \{f[K, \mathcal{L}(K, \tau)] - \delta K\}. \quad (34)$$

B Computation

The numerical computation of a time-consistent equilibrium involves a method for finding the equilibrium decision rule functions. Here we propose a procedure for finding these functions, and we specialize the discussion to our baseline model. That is, we need to find a savings function, $\mathcal{H}(K, \tau)$ and a tax function $\Psi(K)$. However, it is advantageous to use our alternative equilibrium definition here and instead look for savings as $h(K)$ (and taxes as $\psi(K) \equiv \Psi(K)$), because a function of one variable is a smaller object to find.

Conceptually, we have a system of two functional equations in h and ψ : the first-order conditions of the private sector and of the government. The former reads

$$\eta(K, \psi(K), h(K)) = 0 \quad (35)$$

and the latter, with all its arguments explicit now,

$$\begin{aligned} & \eta_{K'}(K, \tau, K') \left\{ u_c[\mathcal{C}(K, \tau, K'), \mathcal{G}(K, \tau, K')] \mathcal{C}_\tau(K, \tau, K') + \right. \\ & \quad \left. u_g[\mathcal{C}(K, \tau, K'), \mathcal{G}(K, \tau, K')] \mathcal{G}_\tau(K, \tau, K') \right\} \\ & - \eta_\tau(K, \tau, K') \left\{ u_c[\mathcal{C}(K, \tau, K'), \mathcal{G}(K, \tau, K')] \mathcal{C}_{K'}(K, \tau, K') \right. \\ & \quad + u_c[\mathcal{C}(K', \tau', K''), \mathcal{G}(K', \tau', K'')] \mathcal{C}_{K'}(K', \tau', K'') \\ & \quad + u_g[\mathcal{C}(K', \tau', K''), \mathcal{G}(K', \tau', K'')] \mathcal{G}'_{K'}(K', \tau', K'') \\ & \quad - \left(u_c[\mathcal{C}(K', \tau', K''), \mathcal{G}(K', \tau', K'')] \mathcal{C}_{\tau'}(K', \tau', K'') \right. \\ & \quad \left. + u_g[\mathcal{C}(K', \tau', K''), \mathcal{G}(K', \tau', K'')] \mathcal{G}_{\tau'}(K', \tau', K'') \right) \frac{\eta_K(K', \tau', K'')}{\eta_\tau(K', \tau', K'')} \left. \right\} = 0, \end{aligned} \quad (36)$$

where $K' \equiv h(K)$, $\tau \equiv \psi(K)$, $K'' \equiv h(h(K))$, and $\tau' \equiv \psi(h(K))$.

One procedure for solving these two functional equations for the two unknown functions is to specify parametric functional forms for h and ψ and pin down the parameters by imposing that the two equations hold on a set of grid point. This amounts to a set of “moment conditions” and boil down to choosing the parameters to minimize a measure of the errors in the equations. We did not follow this procedure, mainly because our main focus is on steady states and because the procedure we do follow delivers, beside information about the steady state levels of K and τ , information about local dynamics as well. Most importantly, however, our procedure for finding a steady state is simple and involves only a very small amount of nonlinear-equation solving (or minimization). It builds on in Krusell, Kuruşçu, and A. (2000), who use a

version of a ‘‘perturbation’’ method (for a discussion of the perturbation method, see Judd (1998)).

We use *s to denote steady-state values. In steady state we obtain

$$\eta(K^*, \psi(K^*), K^*) = 0 \quad (37)$$

and

$$u_c^* \mathcal{C}_\tau^* + u_g^* \mathcal{G}_\tau^* - \frac{\eta_\tau^* h^{*(1)}}{\eta_K^* + \eta_\tau^* \psi^{*(1)}} \left[u_c^* \mathcal{C}_{k'}^* + u_c^* \mathcal{C}_K^* + u_g^* \mathcal{G}_K^* - (u_c^* \mathcal{C}_\tau^* + u_g^* \mathcal{G}_\tau^*) \frac{\eta_K^*}{\eta_\tau^*} \right] = 0, \quad (38)$$

where a superindex $^{(i)}$ on a function refers to its i th derivative and we have used $\eta_K + \eta_\tau \psi^{(1)} + \eta_{K'} h^{(1)} = 0$ (which follows from (35) and h and ψ being differentiable).

Unfortunately, equations (37) and (38) cannot be solved directly for the steady-state levels because in addition to $K^* = h(K^*)$ and to $\tau^* = \psi(K^*)$ the terms $h^{(1)}(K^*)$ and a $\psi^{(1)}(K^*)$ appear yielding a system with two equations and four unknowns. To get around this problem we rewrite the steady state versions of the EE and the GEE compactly as

$$0 = EE^{(0)}(h^{(0)}, \psi^{(0)}) \quad (39)$$

$$0 = GEE^{(0)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}). \quad (40)$$

Note that the way we have written these equations allows us to take derivatives n times (under the assumption that the h and ψ functions are sufficiently differentiable) in a simple way which yields

$$0 = EE^{(n)}(h^{(0)}, \psi^{(0)}, \dots, h^{(n)}, \psi^{(n)}) \quad (41)$$

$$0 = GEE^{(n)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}, \dots, h^{(n+1)}, \psi^{(n+1)}). \quad (42)$$

We next propose an iterative procedure to obtain an approximate solution to this problem based on successively assuming that higher-order derivatives are zero. This procedure turns out to be quite easy to implement since it never involves more than the computation of derivatives and solving a nonlinear system of two equations and two unknowns.

To explain how our procedure works, consider the system

$$0 = EE^{(0)}(h^{(0)}, \psi^{(0)}) \quad (43)$$

$$0 = GEE^{(0)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}) \quad (44)$$

$$0 = EE^{(1)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}) \quad (45)$$

$$0 = GEE^{(1)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}, h^{(2)}, \psi^{(2)},) \quad (46)$$

...

$$0 = EE^{(n)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}, \dots, h^{(n)}, \psi^{(n)}) \quad (47)$$

$$0 = GEE^{(n)}(h^{(0)}, \psi^{(0)}, h^{(1)}, \psi^{(1)}, \dots, h^{(n)}, \psi^{(n)}, 0, 0) \quad (48)$$

where the solutions (if they exist) are denoted $\{h_n^{(i)}, \psi_n^{(i)}\}_{i=0, \dots, n}$.

This problem has a recursive structure. For any $h^{(0)}$ equation (43) determines $\psi^{(0)}$. Next, equations (44) and (45) determine $\{h^{(1)}, \psi^{(1)}\}$. The next two equations determine $\{h^{(2)}, \psi^{(2)}\}$, and so on. Finally the last equation may or may not be satisfied. So implicitly the set of equations (43)–(48) can be seen as one equation (equation (48)) in one unknown: $h^{(0)}$. To arrive at and evaluate the final equation, we thus have to solve one single equation (equation (43)), which has a closed-form solution for $\psi^{(0)}$, and n systems of two equations and two unknowns; the first of these n systems amounts to a quadratic equation, and the remaining $n - 1$ systems are all linear, given the values already solved for. Therefore, since all these steps can be accomplished in closed form, the only nonlinear routine necessary is one which handles one equation and one unknown, independently of n ! Notice also that the need to take successively higher derivatives is not practical to do by hand, but it can be automatized with symbolic software, such as MAPLE, and integrated with the main program.

Note now that the solutions $\{h_n^{(0)}, \psi_n^{(0)}\}$ define a sequence of solutions for the steady state. Each of these is based on the assumption that the $n + 1$ -order (and higher) derivatives of functions h and ψ are 0. If this sequence converges, its limit is a steady state. Our numerical calculations of the first few elements of this sequence show that after 3 or 4 iterations, the values of $\{h_n^{(0)}, \psi_n^{(0)}\}$, are very close to each other. We have also verified that our algorithm reproduces the closed-form solutions for our parametric examples in Section 6.