Time-Consistent Policy

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, October 3, 2001

PRELIMINARY: DO NOT CIRCULATE W/O AUTHOR APPROVAL

Abstract

The goal of this paper is to study optimal government policy when there is no way of committing to future policy and this restriction is binding. The paper has a methodological part and an applied part. The applied part analyzes the optimal choice of government expenditures when the financing of these expenditures involves distortionary taxation and has dynamic consequences—we are studying a simple but canonical dynamic public economics problem here. The analysis is quantitative: we ask how big the expenditures ought to be, and how they ought to be financed, given that only proportional taxation (on labor or capital income) is feasible.

The methodological part develops methods that follow up on Kydland and Prescott's (1977) analysis: we show how to characterize, theoretically and numerically, *time-consistent* policy. By time-consistent policy we mean that made by a government rationally foreseeing how any of its future counterparts will respond to the initial conditions given to it by the past. Moreover, our focus is on policy outcomes where reputation has no role to play at all. That is, even though our environment may admit trigger mechanisms à la Chari and Kehoe (1990), allowing both better and worse equilibria than the one we look at, our goal is to establish precisely how to find the reputation-free, or "fundamental", equilibrium. Formally, the central object we search for is a smooth function mapping the directly payoff-relevant state variables into policy. We show that this function has to satisfy a key first-order condition for the government: a *generalized Euler equation*, which is a functional equation. This equation can be interpreted in several intuitive ways, reflecting both traditional public finance and macroeconomic analysis, and it shows how state variables such as the stock of capital actually does allow the current government to achieve dynamic goals also when reputation cannot be used. Finally, we show how to solve this equation numerically.

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

* Correspondence to Krusell at Department of Economics, University of Rochester. We thank comments of attendants to seminars in the Minneapolis and Richmond Feds, the IIES, and the SED 2001 meetings. Klein thanks the hospitality of the IIES, Krusell thanks the National Science Foundation, and Ríos-Rull thanks the National Science Foundation for Grant SES–0079504 and the University of Pennsylvania Research Foundation for their support.

1 Introduction

The following problem is pervasive in economics: a decision maker at time t cares about the future, disagrees with the decision maker at $t+1$, but has no direct influence over it. Examples of this kind of problem include those of optimal taxation, optimal monetary policy, savings under imperfect altruism, dynamic durable goods games, dynamic political economy and so on.

Kydland and Prescott (1977) and Calvo (1978) first pointed to the existence of time inconsistency in government policy choices. In addition to pointing out the time–inconsistency problem of optimal policy choice, the former authors also discuss how to find a Markov equilibrium. The idea was to think of what optimal policy would look like in a finite-period model, and then study the limit of finite-period models. Such a limit, if it exists and is unique at least, would by definition not involve reputation mechanisms, since no reputation is possible in finite-horizon contexts.[1] Later, the focus shifted toward studying equilibria where reputation is central, following methods developed in Abreu, Pearce, and Stacchetti (1990), as in the work of Chari and Kehoe (1990). The main results from that approach involve specifying circumstances under which the commitment solution is attainable with triggers, despite a lack of direct commitment.

The motivation underlying this work is a belief that reputation mechanisms do not necessarily work, and that one would at least like to know what happens in their absence. In particular, the infinite-horizon trigger-strategy equilibria emphasized by Chari and Kehoe are only some of many, many equilibria, and there is no consensus on which of all these equilibria is most "reasonable". For this reason, we believe that it is fruitful to pay separate attention to the "fundamental" equilibrium, i.e., that originally sought by Kydland and Prescott. Not only does that equilibrium have a close connection with any equilibrium one would find in a corresponding finite-horizon economy, but it seems to be a useful benchmark, and one with which the best reputation equilibrium can be contrasted.

[1] To be careful, reputation effects could occur in the finite-horizon games if there were multiple equilibria, thus allowing reputation play the role of selecting among equilibria. Our discussion thus presumes uniqueness in the finite-horizon contexts.

In this paper we explore the theoretical predictions of models with lack of commitment by deriving a set of functional first-order conditions that characterizes the smooth Markov-perfect equilibria of these environments: a *time-consistent policy equilibrium*. The key endogenous object is thus a smooth function mapping the directly payoff-relevant state variables—such as the capital stock in a macroeconomic application—into policy. Smoothness of our policy function is a key ingredient, as differentiability of this function is a central component in the derivation of the first-order condition characterizing optimal time-consistent policy. Smoothness is also key for numerical solution of the model, as will be explained below.

As well as describing our approach in an abstract way, we study a canonical problem in public finance: that of the optimal provision of a public good in a neoclassical growth environment by a benevolent government which only has access to distortionary taxation. This problem not only is a simple application of the general ideas, but it allows us to perform a very natural quantitative analysis: how much public good should the government provide, and how should it be financed? To this end, we use our theoretical tools in this application and provide answers to these questions: we calculate optimal time-consistent equilibria under a period-by-period balanced-budget constraint in economies with parameterizations that are aligned with those in the macro literature. In this context, in order to interpretat our findings, we compare our results to those that would result in a world where the government has access to a commitment technology—a Ramsey allocation—and to the "first best", or Pareto, outcome: the allocation that would result when taxes are all lump-sum.

Our methodological findings focus on the government's first-order condition. This condition has not been explicitly derived before in the literature, and it is useful in several ways. First, of course, it summarizes and allows an interpretation of the nature of the tradeoffs facing the government. Second, the tradeoffs that are highlighted in this equation can be evaluated quantitatively, as in the literature on dynamic consumer savings and portfolio choice. Thus, it allows us to see not only what the tradeoffs are but to form a view on their relative importance.

Third, the first-order condition of the government emphasizes a point which makes time-consistent policy equilibria without commitment different in nature than those dynamic problems we are used to studying: this first-order condition involves *derivatives* of decision rules of the government decision maker. We call our government first-order condition a *generalized Euler equation* (a GEE) for this reason, thus borrowing a term used in the dynamic multiple-selves games of the recent economics and psychology literature, where this phenomenon also appears. 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. The reason is simple: in these standard models, the envelope theorem eliminates them. That is, it is to a first-order approximation not necessary for the current decision maker to think about what the future decision makers will do, since he agrees with them: they behave optimally from his perspective. In the present context, a disagreement between the current and the future decision makers is instead precisely what is central: when I make a decision today, I have a reason to try to influence the behavior of the future decision makers in a direction that is favorable to me, so I will therefore care about how their decisions vary with the state variables that they inherit from me, and that I can affect through my present choice.

In terms of finding steady states, our characterization means that it is not possible to find a steady state without simultaneously characterizing dynamics: to ensure that it is optimal to keep the tax at a specific constant level, we need to make sure that a deviation is not profitable, and such a deviation involves a dynamic response—as represented by the derivative of the decision rule at the steady state point—of a future decision maker, by the above argument. Can the derivatives of these decision rules be found in a simple way? One can differentiate the first-order conditions—which are functional equations—and obtain an expression (yet another functional equation) in the unknown derivatives. However, since the level version of this equation already contained a derivative, we now also obtain a second derivative of the unknown function, which is yet another unknown. Absent a functional-form short-cut to what obviously becomes a problem of infinite regress, one therefore cannot in general find steady states as in a standard macroeconomic model. Instead, one needs to rely on numerical methods for the computation of steady states. Fortunately, however, our computational methods allow us to find steady states quite easily.

The derivation of the GEE is the outcome of a central insight: the optimal government actions in a Markov equilibrium solve a certain dynamic optimization program that is actually time–consistent. This program both has a standard sequential representation and a recursive one. Given that this recursive formulation is possible, the GEE can be thought of as arising from a typical variational problem: given a value for the current state variable, fix the value of state variable two periods from now and vary the state variable in between—by varying the controls in the current and next period—in an optimal manner. This principle is standard, but it perhaps surprising that it appears here, since it is not literally possible to directly influence the state variable two periods hence—it is chosen by the next decision maker!

In our application to the optimal provision of public goods, we show that it is possible to state the GEE in two alternative ways, both of which provide key insights in terms of the solution of the problem faced by the government, and both of them equivalent up to algebraic manipulations. The first version of the GEE is the classic one in macroeconomics where the (total) effects of a tax hike today are compared with those that arise tomorrow. These effects are obviously posed in terms of marginal utilities. The alternative specification presents the GEE as a weighted average of wedges. In this case the government can be seen as choosing the optimal relative size of deviations from the first best, with the GEE providing a formula for how to weigh those wedges: if a certain wedge is narrowed today, some other wedge has to increase, either today or in the future, and the GEE tells us which and by how much. This approach follows a tradition in the public finance literature by describing optimal taxes as a compromise between different wedges.

In our quantitative experiments it turns out that the properties of taxes and allocations are quite different among the time-consistent, Pareto, and Ramsey equilibria, 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 strategic manipulation of the capital stock bequeathed to the next decision maker—can be quite powerful and even qualitatively surprising. For example, in an economy where the government taxes current capital income alone to finance the current provision of public goods, it does not do so optimally, even in the first period. One would perhaps have expected an optimal outcome there, since the capital income tax is like a lump-sum tax. However, foreseeing that the current governments will tax capital, which retards capital accumulation earlier on, the current government wishes to increase that capital accumulation somehow. It does this by taxing current capital income less, thus using a wealth effect to increase current and future savings in the direction it desires. This is of course a strategic manipulation: it raises savings to counteract a view of the next government that it disagrees with: it views capital taxes next period as distortionary, whereas the next government does not. Thus, in this economy, the provision of public goods ends up being too low.

The tools developed in this paper are, we think, entirely general and applicable to a wide variety of contexts, as hinted to above. However, there is earlier work in this direction. First there is some study of Markov equilibria of the type that we are interested in 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 for identifying Markov equilibria nor to verify whether they exist.

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), Pollack (1971), Laibson (1997), and others. This literature contains the derivation of a GEE, and Krusell, Kuruşçu, and Smith (2000) show how to solve it numerically for a smooth decision rule equilibrium. As will be elaborated on below, the smooth rule is very difficult to find with standard methods, and the contribution of Krusell, Kuruşçu, and Smith (2000) which is applied here is to show how a perturbation method, which relies on successive differentiation of the GEE, resolve these problems. Here is where smoothness of the policy function becomes operationally important.

The outline of the paper is as follows. Section 2 describes a rather general case and our approach to dealing with it. In section 3 we consider a particular environement and describe the problem faced by the agent. Section 4 goes on to derive the GEE for this economy. Section 5 discusses the economic interpretation of the GEE both in terms of wedges as it is standard in Public Finance as in terms of a standard Euler equation where the effects today of a change in policy are weighted against the effects tomorrow. Section 6 discusses some non-standard computational problems that arise in this context. Section 7 discusses the properties of the actual policies that arise in environments where governments do not have access to a commitment technology (Markov) and compares them to those that arise both in environments with commitment (Ramsey) and in environments where the government has access to lump sum taxation (first best or Pareto). Section 8 concludes, while the Appendix includes some auxiliary formal definitions and the description of the computational procedures that we use to compute equilibria.

2 A general description of the problem and its solution

Suppose household decisions are characterized by the following set of Euler equations.

$$f(x_t, x_{t+1}, d_t, d_{t+1}, \tau_t, \tau_{t+1}) = \mathbf{0}$$

where the zero is of dimension $n_x + n_d$. In this context, x_t is an (aggregate) state vector, d_t is a vector of (aggregate) private decisions, and τ_t is a vector of government policies made in period t. These Euler equations include first-order conditions for households as well as resource constraints and other exogenously specified constraints faced by the government.

Suppose also that the government acts as a Stackelberg leader vis-à-vis the private sector in each period and chooses current policies in each period t so as to maximize

$$\sum_{s=t}^{\infty} \beta^{s-t} u(x_s, d_s, \tau_s)$$

subject to the constraints that (i) all future policies are determined by a given policy function Ψ, i.e. for all $s > t$, $\tau_s = \Psi(x_s)$, and (ii) that the private sector Euler equations will be satisfied whatever the government chooses to do; this is a perfection requirement.

We are now in a position to define a recursive Markov equilibrium for this environment.

Definition 1 *A recursive Markov equilibrium for the environment described above is a private decision rule $\mathcal{D}(x, \tau)$, an equilibrium law of motion $x' = \mathcal{H}(x, \tau)$, an equilibrium policy rule $\Psi(x)$, and a value function $v(x)$ such that*

1.
$$f(x, \mathcal{H}(x, \tau), \mathcal{D}(x, \tau), \mathcal{D}(\mathcal{H}(x, \tau), \Psi(\mathcal{H}(x, \tau))), \tau, \Psi(\mathcal{H}(x, \tau))) = \mathbf{0}$$
for all x and τ,

2.
$$\Psi(x) \in \underset{\tau}{\operatorname{argmax}} \{u(x, \mathcal{D}(x, \tau), \tau) + \beta v(\mathcal{H}(x, \tau))\}$$
for all x, and

3.

$$v(x) = u(x, \mathcal{D}(x, \Psi(x)), \Psi(x)) + \beta v(\mathcal{H}(x, \Psi(x)))$$

for all x.

The main contribution of this paper comes from the following characterization of the Markov equilibrium.

Proposition 1 *Let $\mathcal{H}$, $\mathcal{D}$ and Ψ be part of some recursive Markov equilibrium. Fix x_0 and define the sequence $\{\tau_t\}$ via the following recipe. Let $\{y_t\}$ be defined via $y_0 = x_0$ and $y_{t+1} = \mathcal{H}(y_t, \Psi(y_t))$ and let $\tau_t = \Psi(y_t)$. Then $\{\tau_t\}$ solves the following maximization problem.*

$$\max_{\{\tilde{\tau}_t\}} \sum_{t=0}^{\infty} \beta^t u(x_t, \mathcal{D}(x_t, \tilde{\tau}_t), \tilde{\tau}_t)$$

subject to x_0 given,

$$x_{t+1} = \mathcal{H}(x_t, \tilde{\tau}_t)$$

and some suitable (environment-specific) no-Ponzi scheme condition.

Proof. Bellman's principle of optimality. ■

Given this characterization, we can derive necessary conditions in the usual way. We call these necessary conditions the *generalized Euler equations* (GEE). In this very general setting, they can be written as

$$u_d \mathcal{D}_\tau - \lambda^T \mathcal{H}_\tau = 0$$

and

$$\lambda + \beta \left\{ u'_x + u'_d \mathcal{D}'_x - (\lambda')^T \mathcal{H}'_x \right\} = 0.$$

where λ is the Lagrange multiplier associated with the constraint $x' = \mathcal{H}(x, \tau)$. T indicates the transpose and $'$ indicates evaluation in the subsequent period, i.e. at $\mathcal{H}(x, \Psi(x))$. As an equilibrium object, λ is a function of x only.

In specific cases, especially when there are just as many elements in τ as there are in x, then one can substitute out the Lagrange multiplier, as indeed we have done in the examples in the subsequent section.

This looks like a perfectly standard Euler equation, but there is one problem with it: it features derivatives of functions that are unknown from the point of view of the researcher. Thus it is not possible to solve for the steady state without knowing something about the first-order dynamics. Regarding the private sector Euler equations as functional equations, one can of course characterize the derivatives $\mathcal{H}_\tau$ etc. by differentiating these functional equations. But this gives rise to more unknowns, in particular the derivative of the policy function Ψ. In order to characterize this derivative, we can differentiate the GEE (when formulated as a functional equation). But this gives rise to still more unknowns, in particular, to *second* derivatives of $\mathcal{H}$ and $\mathcal{D}$. And so on. See section 6 for how we break this infinite regress.

3 Optimal public goods provision under a period-by-period balanced-budget requirement

Consider a standard growth model with a large number of agents that care about discounted sums of a per period utility function of consumption, leisure and a public good. The technology is neoclassical: a constant-returns-to-scale production function that uses capital and labor services to produce output. The government finances the public good under a period-by-period balanced budget through income taxes (we also discuss the implications of using only labor or capital income taxation). We start in section 3 describing the problem of the household, we follow in in section 3.2 with production, in section 3.3 with a generic government policy, in section 3.4 we describe the equilibrium given future policies, and we finish in section 3.5 where we state the problem of the household.

3.1 The household's problem

The representative household chooses consumption, hours worked, and savings taking prices and taxes as given.

$$\max_{\{c_t,\ell_t,k_{t+1}\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \; u(c_t, 1-\ell_t, g_t) \tag{1}$$

subject to k_0 and

$$c_t + k_{t+1} = k_t + (1-\tau_t)\,[w_t\;\ell_t + r_t\;k_t] \tag{2}$$

where, as is standard, $\{c_t, \ell_t, k_{t+1}\}$ denotes current consumption, current hours worked, and the choice for next period's assets (capital), g_t is the per caput amount of public expenditures, and w_t and r_t are the rental prices of the production factors.

3.2 Firms' problem

The firms solve a standard, static problem that yields that factor prices equal marginal productivities in equilibrium. To avoid cumbersome notation from here on we impose directly the equilibrium conditions that competition imposes on prices. They are

$$r_t = f_K(K_t, L_t) - \delta \tag{3}$$
$$w_t = f_L(K_t, L_t), \tag{4}$$

where f is the neoclassical production function, K_t and L_t are the aggregate quantities of capital and labor, and where δ is the depreciation rate of capital.

3.3 Government policy

Suppose, in line with our Markov equilibrium concept, that the government tax policy obeys a rule determines the current tax rate of income as a function of the aggregate stock of capital: $\tau = \Psi(K)$.

The government is subject to a period-by-period balanced-budget restriction that effectively implies that public expenditures are equal to government revenues:

$$G_t \;=\; \tau_t\,[f(K_t, L_t) - \delta K_t] \;=\; \Psi(K_t)\,[f(K_t, L_t) - \delta K_t]. \tag{5}$$

Feasibility is implied by the government and the private sector decisions together; it is given by

$$C_t + K_{t+1} + G_t = f(K_t, L_t) + (1 - \delta)K_t. \tag{6}$$

3.4 Recursive equilibrium given future policies

With these elements in place, the definition of a recursive competitive equilibrium given a policy function Ψ of the government is standard. The gist of the definition are functions for the aggregate law of motion of capital and for aggregate labor that depend only on the fundamental state variable of the economy (the stock of capital K). They arise 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 and labor as functions 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 as a pair of functions that yield aggregate labor and next period 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 functions for labor supply and capital accumulation as

$$L = \mathcal{L}(K, \tau) \tag{7}$$

$$K' = \mathcal{H}(K, \tau) \tag{8}$$

The explicit dependence of these functions on the tax rate allows us to use them 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. And the current government can affect next period's capital according to H—through its second argument—and the capital stocks after that as well, through similar channels. Note also that sometimes the equilibrium functions are written as $\mathcal{L}(K, \tau; \Psi)$, making explicit how current equilibrium responses depend on the expectations of how future policy is set. We do not use this notation here to avoid clutter. In Appendix A we define in detail the concept of recursive competitive equilibrium.

For convenience, we also define two auxiliary functions that are helpful in reducing clutter; they are completely determined by the other objects.

$$C = \mathcal{C}(K, \tau) \tag{9}$$

$$G = \mathcal{G}(K, \tau) \tag{10}$$

and satisfy

$$\begin{aligned} \mathcal{G}(K, \tau) &= \tau \left\{ f[K, \mathcal{L}(K, \tau)] - \delta K \right\} & (11) \\ \mathcal{C}(K, \tau) &= f\left[K, \mathcal{L}(K, \tau)\right] + (1 - \delta)K - \mathcal{H}(K, \tau) - \mathcal{G}(K, \tau). & (12) \end{aligned}$$

3.4.1 The household's first-order conditions in recursive competitive equilibrium

We just now state the necessary conditions that these functions have to satisfy: those implied by the first-order conditions of maximization problem of the representative household of this economy. These are functional equations, i.e., they have to be satisfied for all τ and all K. We write them as the first-order condition for labor,

$$\begin{aligned} u_c \left[\mathcal{C}(K, \tau), 1 - \mathcal{L}(K, \tau), \mathcal{G}(K, \tau)\right] \cdot f_L \left[K, \mathcal{L}(K, \tau)\right] (1 - \tau) = \\ u_\ell \left[\mathcal{C}(K, \tau), 1 - \mathcal{L}(K, \tau), \mathcal{G}(K, \tau)\right], \end{aligned} \tag{13}$$

for all (K, τ), and the first-order condition for saving, also for all (K, τ):

$$\begin{aligned} u_c \left[\mathcal{C}(K, \tau), 1 - \mathcal{L}(K, \tau), \mathcal{G}(K, \tau)\right] = \\ \beta u_c \left(\mathcal{C}\{\mathcal{H}(K, \tau), \Psi[\mathcal{H}(K, \tau)]\}, 1 - \mathcal{L}\{\mathcal{H}(K, \tau), \Psi[\mathcal{H}(K, \tau)]\}, \mathcal{G}\{\mathcal{H}(K, \tau), \Psi[\mathcal{H}(K, \tau)]\}\right) \cdot \\ \left[1 + \{1 - \Psi[\mathcal{H}(K, \tau)]\{(f_K\{\mathcal{H}(K, \tau), \mathcal{L}[\mathcal{H}(K, \tau)]\} - \delta)\right], \end{aligned} \tag{14}$$

Notice how Ψ is a determinant of $\mathcal{H}$ and $\mathcal{L}$: the expectations of future government behavior influence how consumers work and save.

3.5 The problem of the government

We are now in a position of writing the problem of the government. 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 future governments rules. Before writing the problem of the agent, note that the current return for the government is given by

$$u[\mathcal{C}(K,\tau), 1-\mathcal{L}(K,\tau), \mathcal{G}(K,\tau)]. \tag{15}$$

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 these 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

$$\max_{K_{t+1},\tau_t} \; u[\mathcal{C}(K,\tau), 1-\mathcal{L}(K,\tau), \mathcal{G}(K,\tau)] + \; \beta\, v\left[\mathcal{H}\left(K,\tau\right)\right] \tag{16}$$

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,

$$v(K) \equiv u\left(\{\mathcal{C}[K,\Psi(K)], 1-\mathcal{L}[K,\Psi(K)], \mathcal{G}[K,\Psi(K)]\} + \; \beta\, v\{\mathcal{H}[K,\Psi(K)]\}\right. \tag{17}$$

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

$$v(K) \;\; = \;\; \max_{\tau} \;\; u\left[\mathcal{C}(K,\tau), 1-\mathcal{L}(K,\tau), \mathcal{G}(K,\tau)\right] + \; \beta\, v[\mathcal{H}(K,\tau)], \tag{18}$$

because by definition Ψ solves (16) at all values of K. A fundamental property of the problem of the government so written, therefore, is that it is *recursive*!

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:

$$\max_{\{K_{t+1},\tau_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \quad \beta^t \quad u[\mathcal{C}(K_t,\tau_t), 1-\mathcal{L}(K_t,\tau_t), \mathcal{G}(k_t,\tau_t)] \tag{19}$$

subject to

$$K_{t+1} = \mathcal{H}(K_t,\tau_t), \quad \text{and } K_0 \text{ given.} \tag{20}$$

This problem has an important property: since it is recursive, the time-inconsistency problem does not appear. Variables chosen by the government at $t+1$ do not affect those variables chosen at t. Of course, this does not mean that the time-inconsistency features are not there, but instead that the objects that enter the problem of the government incorporate the elements required to have a time-consistent version of the problem.

We turn now to the derivation of the GEE. We derive it from the sequential version of the problem of the government because it is less cumbersome than using the recursive problem.

4 Derivation of the generalized Euler equation

We have already discussed the functions that are needed to characterize an equilibrium for this environment. They consist of functions $\mathcal{H}$and $\mathcal{L}$ that describe the behavior of the private sector, and function Ψ that describes the behavior of the government.[2,3] We have also discussed two necessary conditions, that these functions have to satisfy in order to be an equilibrium and that determine a functional equation. These are the static first-order condition of the household described in equation (13) and the

[2]Functions $\mathcal{C}$ and $\mathcal{G}$ giving aggregate consumption and public expenditures are trivially derived from the other 3 functions, and do not need to be described in any detail.

[3]The function v, giving the value for the government as a function of the state, is not needed for characterizing the equilibrium first-order conditions and, therefore, we omit it.

dynamic first-order condition described in equation (14). The additional functional equation that we need should characterize the behavior of the government. But the first-order condition of equation (19) should do the job.[4] This functional equation is the one that we call the generalized Euler equation, or with the acronym GEE.

To derive the GEE is straightforward given our sequential formulation: it is done as in any optimal growth model, with τ playing the role that consumption plays in the growth model. The differentiation yields

$$\begin{aligned} u_c\,[L_\tau\; f_l - H_\tau - G_\tau] - u_l\; L_\tau + u_g\; G_\tau +\; \beta\; H_\tau \cdot \\ \Big\{ u_c'\,[f_K' + f_l'\; L_K' + 1 - \delta - H_K' - G_K'] - u_l'\; L_K' + u_g'\; G_K' \\ - \frac{H_K'}{H_\tau'}\,(u_c'[\; f_l'\; L_\tau' - H_\tau' - G_\tau'] - u_l'\; L_\tau' + u_\tau'\; G_\tau')\Big\} = 0, \quad (21) \end{aligned}$$

where we use primes instead of t's to maintain the notation of functional equations and where the arguments of the functions are omitted for readability. We now have a system of 3 functional equations and three unknown functions. Its solutions are Markov equilibria of the problem with a benevolent government that cannot commit.

Recall that the derivation of the first-order conditions of standard recursive problems uses mainly the the envelope condition. In this particular problem, the differentiation of the expression for $v(k)$ results in a number of terms, some of which multiply $\Psi'(k)$; all those terms cancel, using the first-order condition just derived (the envelope theorem being at work). However there is another term remaining in this expression: $\beta v'(H(k,\Psi(k)))H_k(k,\Psi(k))$. It can be solved for from the first-order condition, delivering an expression for $v'(k)$ not involving any further value function derivative. Finally, we would need to evaluate the expression for v' at $H(k,\Psi(k))$ and substitute it into the first-order condition before we obtain the GEE.

We signalled in the introduction that the derivative of the policy rule, Ψ', would appear in the first-order condition for the government. However, Ψ' does not appear

[4] We assume that the first-order condition, along with a transversality condition we do not state—it will be automatically satisfied in any steady state, for example—is sufficient. This assumption is hard to defend analytically for a general model, just like in the optimal taxation literature exploring Ramsey problems. In practice, however, one can verify concavity with numerical methods in any specific application.

directly in the GEE above. However, it appears indirectly, through the derivatives of the private sector decision rules. For in order to find these derivatives, one would need to differentiate the private-sector first-order conditions, and in this process Ψ' would appear, since it is a determinant of those equations.

Finally, notice that the sequential formulation actually allows us to think of deriving 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. This result is surprising in the sense that the current government literally has no direct control over the state variable two periods hence (because it has no direct influence on τ'). The key behind our result thus is showing that there exists a problem for which variational analysis is possible. This problem is not the one where the government has full power over future policy (such a problem would by definition be a problem with commitment—the Ramsey formulation), but rather one where certain aspects of future policy are taken as given. These aspects, precisely, are the *shapes* of the private-sector functions, which fundamentally depend on the future Ψ functions and which the current government cannot affect.

We next discuss some characteristics of the GEE.

5 Discussion of the GEE

The formula depicted in equation (21) is only one of various alternative representations that we could have chosen. We will focus on two particular versions of the GEE. One of these draws on the tradition in macroeconomics, and the other is inspired by the public finance literature.

5.1 The GEE in terms of direct marginal utilities (for macroeconomists)

Perhaps the most obvious way of thinking about the GEE is that its various terms capture marginal effects on current and future utility of a change in policy, as in (21).

The effects of the tax rate over today's variables are be traced to effects on

1. Current consumption,
   (a) via higher labor supply ($L_\tau f_l$);
   (b) via lower savings (H_τ); and
   (c) via higher government spending (G_τ).
2. Current leisure, via higher labor supply (L_τ).
3. Current government spending, which goes up (G_τ).

While the effects of the tax hike on future utility-relevant variables occurs via a decrease in savings (H_τ), leading to:

1. A change in next period's consumption via a direct effect on production and undepreciated capital ($f'_k + 1 - \delta$), an indirect effect on labor supply ($f'_l L'_k$), an indirect effect on saving ($-H'_k$), and an indirect on government spending (G'_k);
2. A change in next period's leisure, through labor supply ($-L'_k$); and
3. A change in next period's government spending directly (G'_k).
4. An induced change in taxes ($-H'_k/H'_\tau$), which affects (as it did today)
   (a) Next period's consumption, via the same channels as in the initial period ($f'_l L'_\tau - H'_\tau - G'_\tau$);
   (b) Next period's leisure (L'_τ); and
   (c) Next period's government spending (G'_τ).

Let us again emphasize that it is by no means obvious that the effects that we list here suffice in characterizing the first-order condition. The above derivations—relying on recursive methods—behind the GEE are therefore a fundamental input into understanding this equation.

The variational aspect of the GEE can be illustrated through the $-H_k'/H_\tau'$ factor: this ratio equals the change in future taxes necessary to keep K'' constant, given that a change in K' has been induced by the hike in τ.

5.2 The GEE in terms of wedges (for the public economics crowd)

The GEE in (21) can be rewritten so that it is a linear combination of wedges. We obtain the following equation.

$$\begin{gathered}\mathcal{L}_\tau \left[u_c f_L - u_\ell\right] + \mathcal{G}_\tau \left[u_g - u_c\right] + \beta \mathcal{H}_\tau \cdot \\ \left\{\mathcal{L}_K' \left[u_c' f_L' - u_\ell'\right] + \mathcal{G}_K' \left[u_g' - u_c'\right] - \tfrac{\mathcal{H}_K'}{\mathcal{H}_\tau'} \left(\mathcal{L}_\tau' \left[u_c' f_L' - u_\ell'\right] + \mathcal{G}_\tau' \left[u_g' - u_c'\right]\right)\right\} + \\ \mathcal{H}_\tau \left\{-u_c + \beta u_c'(1 + f_K' - \delta)\right\} = 0\end{gathered} \tag{22}$$

for all K (again, the arguments of the functions are suppressed for readability).

This formulation follows the public finance tradition of characterizing optimal taxes as combinations of wedges. To illustrate the power of this approach, let's look at the wedge version of the GEE that would have arisen in an economy with lump-sum taxes. It is given by

$$\mathcal{G}_\tau \left[u_g - u_c\right] + \beta \mathcal{H}_\tau \cdot \left[\mathcal{G}_K' - \frac{\mathcal{H}_K'}{\mathcal{H}_\tau'} \mathcal{G}_\tau'\right] \left[u_g' - u_c'\right] = 0. \tag{23}$$

where the sub-indices denote the derivatives of the functions. Note that a policy that sets marginal utility of government expenditures equal to that of private consumption satisfies equation (23) 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

$$\mathcal{G}_\tau \left[u_g - u_c\right] + \beta \mathcal{H}_\tau \left[\mathcal{G}_K' - \tfrac{\mathcal{H}_K'}{\mathcal{H}_\tau'} \mathcal{G}_\tau'\right] \left[u_g' - u_c'\right] + \mathcal{H}_\tau \left\{-u_c + \beta u_c'(1 + f_K' - \delta)\right\} = 0. \tag{24}$$

Here the intertemporal distortion introduced by the capital income taxes enters the GEE. However, it does so in a very specific way. The current government takes into account how its choice of taxes affects tomorrow's capital which in turn affects tomorrow's government's choice of taxes. Clearly, in this case, the Pareto optimum is not attainable, because of last term, which cannot be zero unless taxes are zero. Zero taxes are feasible, but they imply zero government expenditures, so that the first two terms become non-zero.

The tradeoff between wedges occurs as follows. Assuming a starting point with a less-than-optimal amount of government expenditures, so that $u_g > u_c$. Then a tax increase would increase government expenditures (assuming we are on the left side of the Laffer curve), leading to a benefit—an increase in the right-hand side of the equation. However, savings would slow—assuming, as seems natural due to a wealth effect, that $\mathcal{H}_\tau < 0$. This would create a cost, since it would decrease the left-hand side through the last term: the wedge from the savings choice, which is positive. Finally, the lowering of the capital stock next period will affect the public-private expenditure wedge in the next period. First, it will have a direct effect by lowering government spending next period, thus inducing a loss, since that wedge is positive. Second, it will force a compensating decrease in next period's tax rate—since K'' has to be held fixed. That tax decrease constitutes another loss, since it induces a second lowering of public spending next period. Thus, all in all an increase in the current tax rate makes the current g vs. c wedge smaller, but it increases all other wedges, through three separate channels.

Note also that the government is completely unconcerned with honoring any form of the intertemporal savings condition of the **previous** period as a Ramsey government would. This is what we should expect from the fact that the policy implied by the GEE has to be time-consistent.

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

$$\begin{gathered}\mathcal{L}_\tau\left[u_c f_L - u_\ell\right] + \mathcal{G}_\tau\left[u_g - u_c\right] + \\ \beta\, \mathcal{H}_\tau \left\{ \left[\mathcal{L}'_K - \tfrac{\mathcal{H}'_K}{\mathcal{H}'_\tau}\mathcal{L}'_\tau\right]\left[u'_c f'_L - u'_\ell\right] + \left[\mathcal{G}'_K - \tfrac{\mathcal{H}'_K}{\mathcal{H}'_\tau}\mathcal{G}'_\tau\right]\left[u'_g - u'_c\right] \right\} = 0.\end{gathered} \tag{25}$$

In this case the government strikes 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 in the leisure–private consumption margin in both the current and the next period. The way it does so is by weighting those wedges both today and tomorrow by how much they move labor and government expenditures.

Notice that in the general income as well as in the only capital or only labor income cases, the Markov equilibrium is not equal to the Ramsey solution. The reason is 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 $t-1$. A Markov policy maker treats the latter as a bygone.

We now turn to discuss some issues regarding the computation of solutions to these functional equations, but before we get into those issues, note that regardless of the form that we give to the GEE, this functional equation includes derivatives of the unknown functions. This means that these derivatives should also be solved for.

6 Issues on Computation

So far we have pointed out that the Markov equilibria of an environment where a benevolent government subject to a period-by-period balanced-budget constraint that does not have access to commitment can be characterized as the solution to a system of three functional equations and three unknown functions. Sometimes systems of three functional equations can be readily solved with those approximation methods that are very popular in economics (see Judd (1998) and Marimon and Scott (1998) for a long list). However, the GEE introduces some subtleties that generate severe computational problems when approached with standard methods.

To discuss some of these problems we now switch our attention to the simplest problem that embodies a generalized Euler equation of the type that we have described. This problem is that of a single household with time-inconsistent preferences of the type known as quasi-geometric discounting: each currently alive agent uses a sequence of discount factors equalling $(1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots)$, where it is presumed that β and δ are both less than 1. This problem has been studied in the literature by Strotz (1956), Phelps and Pollak (1968) and Pollack (1971), and more recently by Laibson (1997), Christopher and Laibson (2000), and Krusell, Kuruşçu, and Smith (2000). Numerical studies of this problem appear in Krusell, Kuruşçu, and Smith (2000) and Angeletos, Laibson, Repetto, Tobacman, and Weinberg (2001).

The problem of this type of household is best modeled as a game between the present self and his future selves. This problem does not have a sequential implementation like those studied in this paper (due to the fact that the preferences of the different selves are not the same). Yet its equilibria can be characterized as the solutions of a functional equation not dissimilar to that of the benevolent government that populates this paper. Such a GEE is

$$u'[f(k) - h(k)] = \beta\ \delta\ u'\{f[h(k)] - h[h(k)]\} \quad \{f'[h(k)] - h'[h(k)](1 - 1/\beta)\} \tag{26}$$

where the function that we are looking for is h and where the functional equation has to hold for all k (at least in the interval between zero and the maximum sustainable capital stock). Note also, that like our GEE, this functional equation has as arguments not only the value of the function, but also its derivatives. In particular, the last term on the right-hand side is the only new component compared to a standard first-order conditino for savings. The new term captures the extra return from savings that appears due to strategic savings. The next self decreases consumption by h' when given one more unit of capital, accounting for the minus sign. In return, that added saving increases utility by its effect on consumption two periods from now and later. In a standard model, that effect is worth exactly h', too, in consumption units, due to the envelope theorem. Here, however, it is worth precisely a factor $1/\beta$ more for the current self than for the next-period self. Since $\beta < 1$, so that $1/\beta - 1 > 0$, this is therefore an additional positive return for the current self from saving more. For a detailed discussion of how to derive and interpret this equation, see (Krusell, Kuruşçu, and Smith 2000).

A natural way of trying to solve this functional equation perhaps is through a global approximation method by, say, an orthogonal family of polynomials. However, this has not been useful in practice for solving the functional equation, and the other iterative procedures tried have not converged either. Krusell and Smith (2000)) suggest that this finding is related to a fundamental property of the GEE: it has infinitely many Markov solutions. In particular, Krusell and Smith (2000)) prove that the multiple-selves game has a continuum of Markov equilibria: it has a continuum of steady states, and even a continuum of decision rules supporting each steady state. These solutions are all found by construction. They consist of *discontinuous* solutions for g: the different gs are all step functions. Thus, trigger-like solutions can be supported with discontinuities in the decision rules. The discontinuities build up incentives precisely by using the disagreements among the different selves; in the case there is agreement, no discontinuous equilibrium exists.

It thus seems possible that one can find many (discontinuous) Markov solutions in the present context as well. We are not interested in these, however; they are a feature of the infinite time horizon, and they have the reputation feature that we set out to avoid. However, for computation, the indeterminacy is important, because it is likely, as Krusell and Smith (2000)) argue, the reason why flexible global nonlinear methods fail. Specifically, even a steady state is hard to pin down, since there are a continuum of steady states.

Smoothness of the g function, then, becomes key in finding equilibria. For the discontinuous equilibrium decision rules do *not* satisfy the GEE—they are not differentiable.[5] Heuristically, it turns out that the key to achieve convergence is through an accurate computation of the steady state. But this is not trivial matter, unlike in standard models where the computation of a steady state just amounts to solving one equation (the functional equation) and one unknown (the point where the unknown function has a fixed point). To see that this is not a trivial undertaking note that at a point where $k = h(k)$, the value of the functional equation (26) depends on the derivative of h at that point. In other words, we have one equation but two unknowns, so it cannot be solved. The only way to proceed is to make assumptions about the derivative of h at the steady state. The obvious first guess might be to set $h'(k)$ to zero at the steady state. This yields as a solution that of the standard consumer ($\beta = 1$) which is obviously not quite correct and thus delivers the wrong steady state. But note that the functional equation (26) holds for all k. This means that so does its derivative, and it can be readily computed by differentiation of both sides of the GEE. Now its derivative

[5] Locally, however, they "almost" satisfy it, which is why numerical problems occur.

has terms in h, h' and unfortunately also in h'' (the function and its first and second derivatives). Now we have two equations, but three unknowns, again an impossible task, but we can make an assumption about the second derivative at the steady state and use the two equations to compute the value of the function and its derivative. And we could go on and on, taking an additional derivative which gives us an additional equation and getting an additional unknown about which me make an assumption. This delivers an infinite sequence, unless a functional-form happenstance presents us with a natural short-cut—we shall see an example of that in the next section. However, what we are really interested in is in whether its first two terms (the steady state and the derivative of the decision rule evaluated in the steady state) converge, as we let the number of included derivatives increase toward infinity. In practice, though, these derivatives can always be computed numerically (with finite-difference methods) and they involve solving a nonlinear equation system.

In terms of an interpretation of why this method works—and it does in the applications encountered so far: fast, monotonic convergence is achieved for all the examples investigated—note that it critically exploits the differentiability of g. Thus, the large set of step function equilibria, which are not differentiable and do not satisfy the GEE that is being differentiated here, will not be numerical attractors. Appendix B describes, in algorithmic form, how to find the steady state of the multiple-selves model.

The iterative procedure we use is not dissimilar to what is called a perturbation method—see the extensive discussion in Judd (1998)). The use of the perturbation method is special here, however, since Judd's examples are all nicely behaved: they do not contain derivatives of decision rules, and can be solved with many other methods.

Having discussed the multiple-selves setup, let us point to the obvious similarity with the present problem—that of a sequence of governments who disagree. They do not disagree, however, from having different preferences, but from perceiving the private-sector responses to tax policy as being different (future governments think the responses to future tax policy is weaker than what the current government thinks). Appendix C shows how to apply the procedure outlined above to the model in this paper.

Now it is time to look at some actual model economies.

7 Optimal policy for actual economies

Partly to illustrate the power of our methods and partly due to their intrinsic interest, we proceed next to look at the actual 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 answers to the optimal policy under the first best (lump-sum taxation) and those implied by a benevolent government that has access to commitment but not to a technology to save resources, this is, the Ramsey government is subject to a period by period balanced budget constraint.[6]

We specify a general per period utility function of the CES class as

$$u(c,\ell,g) = \frac{\left[(1-\alpha_p)\left(\alpha_c c^\rho + (1-\alpha_c)\ell^\rho\right)^{\Psi/\rho} + \alpha_p g^\Psi\right]^{\frac{1-\sigma}{\Psi}} - 1}{1-\sigma}. \tag{27}$$

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

$$u(c,\ell,g) = (1-\alpha_p)\alpha_c \ln c + (1-\alpha_p)(1-\alpha_c)\ln \ell + \alpha_p \ln g \tag{28}$$

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

Our parameterization of the baseline economy is very basic and does not need a lot of discussion. We want 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 standard in the macroeconomic literature.

[6]Stockman (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.

For clarity we choose the baseline economy to have log utility which makes preference separable (making cross derivatives zero).[7] We report the values of the parameters that implement our choices in Table 1.

Parameter Values		
$\theta = .36$	$\alpha_c = .30$	$\alpha_p = .13$
$\beta = .96$	$\delta = .08$	$\rho = .0$
$\Psi = .0$	$\sigma = 1.0$	

Table 1: Parameterization of the Baseline Model Economy.

7.1 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, meaning, more precisely: 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 we look at the Markov equilibrium). Table 2 reports the steady-state allocations of these three economies.

Let us comment on some of the properties of these allocations. 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

[7] When, 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.

Labor taxes, Endogenous g			
Steady State	Type of Government		
Statistic	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; Separable utility in logs.

allocation (dah) while the Ramsey economy has a distortionary tax that discriminates against produced goods and in favor of leisure. As a result, leisure is a lot larger 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) is a lot smaller 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.[8]

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; 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 (25) the optimal policy of the

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

Markov case amounted to strike 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 in the leisure–private consumption margin. This balance does not imply setting the margin between the public and the private good to zero. Note that indeed, the term $u_g - u_c$ is positive in the Markov case, making the second term of equation (25) positive while the first is negative. The difference with the Ramsey case can be, perhaps, 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 $t-1$. A Markov policy-maker treats the latter as a bygone.

7.2 Capital income taxation

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

Capital Income taxes, Endogenous g			
Steady State	Type of Government		
Statistic	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.

This tax is in general very distortionary. The Ramsey government understands this very clearly well 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 influnce the future choices is of course smaller than that of the Ramsey government, and as a result its capital tax rate is larger and capital and output is 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 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 exactly $(1-\alpha_c)$, 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.

7.3 Taxes on total income

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

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

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.

increase yesterday's incentives to work, and in addition it tries to affect the future tax rate choices, all of which induces a lower than optimal government sector. This result is perhaps surprising because are consistent with an impression that the Markov government view its taxes as less distortionary than does the Ramsey government.

Besides 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.

Note that some of the different characteristics of the three tax alternatives is that 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.[9] 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

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

income tax rate for the same tax rate, or less distortionary for the same revenue. In addition, the Markov government understands that its choice affects the savings of the private sector and, indirectly, next period's tax rate, which in the case of total income or capital income is what determines the rate of return on current savings and hence current savings itself.

With respect to tax outcome, first, as should have been expected, the larger is the role of capital income taxes (which implies an ordering with capital income first, followed by total income and last labor income), the lower is the stock of capital, and hence the lower is output. The differences are large. The second point to note is that hours worked are actually varying very little across environments. Third, and, 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, while considering taxes to be non-distortionary, would do 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.

8 Conclusion

In this paper we have characterized the set of functional equations that are required for characterizing Markov equilibria of 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 in a specific way. The so posed description of the responses of the private sector is unknown in its initial formulation and has to be solved simultaneously with the behavior of the government. This leads to a natural characterization of government behavior in terms of the first-order conditions of such a problem that we term the generalized Euler equation.

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 believe that the methods that we have developed for solving for our Markov equilibrium are very general and can be applied to a vast set of environments beyond 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.

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.

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.

Angeletos, G., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001): "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation," Manuscript.

Calvo, G. (1978): "On the Time Consistency of Optimal Policy in a Monetary Economy," *Econometrica*, 46, 1411–1428.

Chari, V. V., and P. J. Kehoe (1990): "Sustainable Plans," *Journal of Political Economy*, 98(4), 784–802.

Christopher, H., and D. Laibson (2000): "Dynamic Choices and Hyperbolic Consumers," Forthcoming in *Econometrica*.

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, Cambirdge.

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 A. A. Smith (2000): "Tax Policy with Quasi-Geometric Discounting," Manuscript.

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," *qje*, pp. 443–477.

Marimon, R., and A. Scott (eds.) (1998): *Computational Methods for the Study of Dynamic Economics*. Oxford University Press.

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*.

Pollack, R. (1971): "Additive Utility Functions and Linear Engel Curves," *Review of Economic Studies*, 38, 401–413.

Stockman, D. R. (1998): "Balanced-Budget Rules: Welfare Loss and Optimal Policies," Mimeo, Universiy of Delawere.

Strotz, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *res*, 23, 165–80.

Appendix

A Definition of the recursive competitive equilibrium

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

Definition 2 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,

$$\begin{aligned} \Omega(k, K, \tau) &= \max_{c,l,k'} \; u(c, l, g) + \beta \; \Omega(k', K', \tau') && (29) \\ \{\ell(k, K, \tau), h(k, K, \tau)\} &\in \operatorname{argmax}_{c,l,k'} \; u(c, l, g) + \beta \; \Omega(k', K', \tau') && (30) \end{aligned}$$

   subject to

$$\begin{aligned} c + k' &= k + [w(K, \tau) \; l + r(K, \tau) \; k] \; (1 - \tau) && (31) \\ g &= \mathcal{G}(K, \tau) && (32) \\ K' &= \mathcal{H}(K, \tau) && (33) \\ \tau' &= \Psi(K) && (34) \end{aligned}$$

2. The agent is representative

$$\begin{aligned} \ell(K, K, \tau) &= \mathcal{L}(K, \tau) && (35) \\ h(K, K, \tau) &= \mathcal{H}(K, \tau) && (36) \end{aligned}$$

3. Factor prices are marginal productivities.

$$\begin{aligned} r(K, \tau) &= f_K[K, \mathcal{L}(K, \tau)] - \delta && (37) \\ w(K, \tau) &= f_L[K, \mathcal{L}(K, \tau)] && (38) \end{aligned}$$

4. The government satisfies its budget constraint

$$\mathcal{G}(K, \tau) = \tau \; \{f[K, \mathcal{L}(K, \tau)] - \delta K\}. \quad (39)$$

B Algorithm to find the steady state of the problem of the quasi–geometric consumer

- To find a steady state: 1 equation and two unknowns.
- Computational solution: Perturbation methods.
  - Assume $g'(k) = 0$. Solve for Steady–state $k_0^* = g(k_0^*)$.
  - Assume $g''(k) = 0$. Solve for Steady–state $k_1^* = g(k_1^*)$ and $g'(k_1^*)$.
  - Keep going
  - Assume $g^n(k) = 0$. Solve for Steady–state $k_n^* = g(k_n^*)$, and $g'(k_n^*)$, up to $g^n(k_n^*)$.
  - Hope k_n^* converges (so far it has).

C Outline of an algorithm for the Markov taxing problem

- Recall that the fundamental difference between this problem and the standard problems of finding a recursive competitive equilibrium or solving a Ramsey problem is The steady state is **not** the solution to a finite-dimensional system of nonlinear equations.
- To see this, notice that the GEE features unknown *derivatives* of the private decision rules. Thus the unknowns are $\bar{L}$, $\bar{K}$, $\bar{g}$, $\bar{\mathcal{H}}_\tau$, $\bar{\mathcal{H}}_K$, $\bar{\mathcal{L}}_\tau$, and $\bar{\mathcal{L}}_K$. Thus there are seven unknowns but just three equations.
- So, differentiate the private first-order conditions with respect to τ and K and the GEE with respect to K. This gives us five new equations. It also gives us several more unknowns, including $\bar{\Psi}_K$ as well as *second* derivatives of $\mathcal{H}$ and $\mathcal{L}$ evaluated at the steady–state.
- This never stops. There are always more unknowns than equations.
- Then: set all derivatives of degree n and greater equal to zero. Then increase n until further increases make a negligible difference.
- The simplest case: $n = 2$. Then we have eight unknowns and eight equations. With $n = 3$ we have 15 equations and 15 unknowns. And so on.

This procedure provides both a robust mechanism to find the steady state, and a local approximation, as well as an anchor to solve for a global approximation using global methods.