BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2007-3

Time-Consistent Control in Non-Linear Models

by Steve Ambler and Florian Pelgrin

Bank of Canada Working Paper 2007-3

February 2007

Time-Consistent Control in Non-Linear Models

by

Steve Ambler[1] and Florian Pelgrin[2]

[1]CIRPEE
Université du Québec à Montréal
Montréal, Quebec, Canada H3C 3P8
ambler.steven@uqam.ca

[2]Département d'économétrie et économie politique
HEC, Université de Lausanne
Lausanne, Switzerland
and CIRANO, EUREQUA, Université de Paris
florian.pelgrin@unil.ch

ISSN 1701-9397

Acknowledgements

We thank Stéphane Adjémian, Antoine d'Autume, Jésus Fernández-Villaverde, Kenneth Judd, Michel Juillard, Paul Klein, Fernando Martin, Pierre-Olivier Weil, and seminar participants at the Society for Computational Economics and the Econometric Society European meetings for helpful comments and discussions.

Abstract

We show how to use optimal control theory to derive optimal time-consistent Markov-perfect government policies in nonlinear dynamic general equilibrium models, extending the result of Cohen and Michel (1988) for models with quadratic objective functions and linear dynamics. We replace private agents' costates by flexible functions of current states in the government's maximization problem. The functions are verified in equilibrium to an arbitrarily close degree of approximation. They can be found numerically by perturbation or projection methods. We use a stochastic model of optimal public spending to illustrate the technique.

JEL classification: E61, E62, C63
Bank classification: Fiscal policy; Monetary policy framework

Résumé

Les auteurs montrent comment la théorie du contrôle optimal permet d'élaborer des politiques optimales temporellement cohérentes en équilibre markovien parfait à l'aide de modèles d'équilibre général dynamiques non linéaires, dans la lignée des résultats obtenus par Cohen et Michel (1988) à partir de modèles dynamiques linéaires où la fonction-objectif est de forme quadratique. Les multiplicateurs de Lagrange du problème de maximisation des agents du secteur privé sont remplacés par des fonctions flexibles des variables d'état de la période en cours dans le problème de maximisation du bien-être collectif. À l'équilibre, ces fonctions se vérifient jusqu'à un degré quelconque d'approximation. Elles peuvent être résolues numériquement à l'aide de méthodes de perturbation ou de projection. Les auteurs illustrent l'emploi de leur technique au moyen d'un modèle stochastique formalisant le niveau optimal des dépenses publiques.

Classification JEL : E61, E62, C63
Classification de la Banque : Politique budgétaire; Cadre de la politique monétaire

1 Introduction

An appealing feature of solving Ramsey (1927) problems to derive optimal second-best government policies in dynamic general equilibrium models is their relative analytical tractability. It is often possible to use the so-called *primal approach*, in which private agents' first order conditions and budget constraints are combined to derive an *implementability constraint*,[1] allowing prices and policy variables to be substituted out of the problem. The choice variables of the optimal policy problem are the allocations themselves. Prices and policies that support the optimal allocations can be derived once the allocations themselves are known. Using the primal approach leads to equations in which expected future allocations have an influence on agents' current behavior. Therefore, optimal policies derived in this manner are generally *time-inconsistent*.[2] The government must be able to commit credibly to its announced policies. Otherwise, it will optimally revise them as time goes by, in which case its announced policies will not be believed by private agents.

It is often interesting to compare the optimal allocations under credible precommitment by the government to optimal allocations where this precommitment is not possible, possibly for institutional or political reasons. In the latter case, dynamic programming can be used to compute optimal Markov-perfect strategies for the government. In special cases, the envelope theorem can be used to eliminate the government's value function from the system of equations.[3] Alternatively, it is possible to linearize the laws of motion of the economy and use quadratic approximations to agents' preferences, so that the value function takes a known form.[4] This approach may be less than satisfactory in the presence of important nonlinearities. In addition, using linear-quadratic approximations may lead

[1] See Chari and Kehoe (1999) for a detailed discussion.

[2] Solutions to the Ramsey problem can be made time-consistent in special cases. The most well-known case is Lucas and Stokey (1983).

[3] See Judd (1998, section16.9), Azzimonti-Renzo, Sarte and Soares (2003), Klein, Krusell and Ríos-Rull (2004) and Ortigueira (2004) for examples.

[4] See Ambler and Paquet (1996, 1997) and Ambler and Cardia (1997)

to misleading welfare comparisons if the deterministic steady state is not a Pareto optimum.[5] An approximation to the government's value function can be found by discretizing the model's state space, but this approach suffers from a curse of dimensionality: it is computationally burdensome with more than a small number of state variables. It would be useful to have an alternative general methodology for analyzing optimal time-consistent Markovian policies.

In this paper, we show how to use optimal control theory to derive time-consistent Markovian government policies in nonlinear dynamic general equilibrium models, extending the insight of Cohen and Michel (1988). They showed that in linear-quadratic environments time-consistent policies compatible with Markov-perfect equilibria can be found using optimal control theory by imposing a linear relationship between predetermined state variables and the costate variables from private agents' maximization problems.[6] We show that by restricting private agents' costates to be a nonlinear function of current predetermined state variables, the optimal control problem of the government becomes recursive (in a sense to be defined below), whereas in the Ramsey problem it typically is not. The nonlinear function is verified in equilibrium to an arbitrarily close degree of approximation. Projection methods or perturbation methods can be used to approximate the function.[7] The equilibria found using this approach are Markov-perfect since the government's policy function is time invariant and depends only on the current state of the economy.[8] The technique can be used to find Markov-perfect equilibria in stochastic models. Many previous treatments

[5]See Kim and Kim (2003). See Woodford (2003, chapter 6) for conditions under which the linear-quadratic approach is justified.

[6]Their methodology was utilized to analyze optimal government policies by a number of researchers. See Currie and Levine (1993), Oudiz and Sachs (1985) and Miller and Salmon (1985) for examples.

[7]See Judd (1998), McGrattan (1999), and Aruoba, Fernndez-Villaverde and Rubio-Ramírez (2004). Projection methods can have better global properties than perturbation methods around a particular equilibrium point.

[8]See Bernheim and Ray (1989) and Maskin and Tirole (1993) for rigorous treatments of the concept of Markov-perfect equilibrium. We exclude more complex strategies that are history-dependent. For examples of the latter, see Benhabib and Rustichini (1997), Benhabib, Rustichini and Velasco (1996), Benhabib and Velasco (1996) and Chari and Kehoe (1990).

of optimal time-consistent government policies have been limited to deterministic models.[9]

The paper is structured as follows. In the following section, we develop an abstract model of the interaction between a representative private agent and a government. In section three, we review how the time consistency problem arises by analyzing a Ramsey problem applied to the abstract model. In the fourth section, we show how to extend Cohen and Michel's (1988) approach to nonlinear models. In section five, we formally demonstrate the recursivity of the government's problem. In the sixth section, we discuss how to calculate a numerical solution to the optimal control problem. In the seventh section, we present a simple model of public spending in order to illustrate the technique. Conclusions are in section eight.

2 The Model

The economy consists of a representative household,[10] a representative competitive firm, and a government.[11] The household has an infinite planning horizon and maximizes its utility taking as given all relative prices and the government's policy rule. The government chooses its policies to maximize social welfare, which in this framework leads it to maximize the utility of the representative household, subject to the first order conditions of the household.

[9] Papers include Klein, Krusell and Ríos-Rull (2004) and Ortigueira (2004).

[10] The approach here could be extended to models of heterogeneous agents, but the notation would be cumbersome. See Ríos-Rull (1995) for a good introduction to heterogeneous agent models.

[11] Although the analysis is framed in terms of optimal government policy, it is clear that it could be used to derive time-consistent feedback rules in any dynamic game with a Stackelberg leader.

2.1 The Household

The utility function of the household can be written as[12]

$$U = E_t \sum_{i=0}^{\infty} \beta^i r \left(z_{t+i}, g_{t+i}, S_{t+i}, s_{t+i}, D_{t+i}, d_{t+i} \right), \tag{1}$$

where z_t is a vector of exogenous state variables of dimension η_zx1, g_t is a η_gx1 vector of government policy variables, s_t is a η_sx1 vector of endogenous state variables under the control of the individual household, S_t is a η_sx1 vector of endogenous aggregate state variables, which are the aggregate counterparts of s_t, d_t is a η_dx1 vector of the household's control variables, D_t is a η_dx1 vector of the aggregate counterparts of d_t, and E_t denotes mathematical expectations conditional on information available at time t. The household chooses $\{d_{t+i}\}_{i=0}^{\infty}$ in order to maximize its utility, subject to the following set of constraints: the law of motion of the household's state variables,

$$s_{t+1} = b \left(z_t, g_t, S_t, s_t, D_t, d_t \right); \tag{2}$$

the law of motion of the aggregate state variables,

$$S_{t+1} = B \left(z_t, g_t, S_t, D_t, \right); \tag{3}$$

the feedback rule for the aggregate control variables,

$$D_t = D \left(z_t, g_t, S_t \right); \tag{4}$$

and the feedback rule for the government's policy variables,

$$g_t = g \left(z_t, S_t \right). \tag{5}$$

The assumption that the law of motion for the household's state variables is an explicit function for s_{t+1} is not innocuous. If there were an implicit relationship between s_{t+1} and current states and controls, the household's first order condition

[12]The notation is patterned after Hansen and Prescott (1995).

for the choice of d_t would depend on the future state of the economy. Solving for the private sector's control variables as an explicit function of current state variables and costate variables as in (10) below would no longer be possible. The solution to this problem leads to a feedback rule of the form

$$d_t = d\left(z_t, g_t, S_t, s_t\right). \tag{6}$$

As equilibrium conditions, we will impose *aggregate consistency conditions*. The laws of motion for S_t and s_t must satisfy

$$b\left(z_t, g_t, S_t, S_t, D_t, D_t\right) = B\left(z_t, g_t, S_t, D_t, \right), \tag{7}$$

and the feedback rules for D_t and d_t must be consistent:

$$d\left(z_t, g_t, S_t, S_t\right) = D\left(z_t, g_t, S_t\right). \tag{8}$$

The Lagrangian of the household's problem can be written as

$$\begin{aligned}\mathcal{L}_t = E_t \sum_{i=0}^{\infty} \beta^i \Big\{ & r\left(z_{t+i}, g_{t+i}, S_{t+i}, s_{t+i}, D_{t+i}, d_{t+i}\right) \\ & +\lambda_{t+i}\Big[s_{t+i+1} - b\left(z_{t+i}, g_{t+i}, S_{t+i}, s_{t+i}, D_{t+i}, d_{t+i}\right)\Big]\Big\}.\end{aligned} \tag{9}$$

The household chooses $\{d_{t+i}, s_{t+i+1}, \lambda_{t+i}\}_{i=0}^{\infty}$. The first order conditions with respect to variables chosen at time t can be written as follows:

$$\begin{aligned} d_t: &\quad \frac{\partial r(\cdot)}{\partial d_t} - \lambda_t \frac{\partial b(\cdot)}{\partial d_t} = 0, \\ s_{t+1}: &\quad E_t\left(\lambda_t + \beta\frac{\partial r(\cdot)}{\partial s_{t+1}} - \beta\lambda_{t+1}\frac{\partial b(\cdot)}{\partial s_{t+1}}\right) = 0, \\ \lambda_t: &\quad s_{t+1} = b\left(z_t, g_t, S_t, s_t, D_t, d_t\right). \end{aligned}$$

When we impose the aggregate consistency constraints, the first order condition with respect to d_t gives a set of η_dx1 static equations. We assume that it is possible to solve the equations explicitly for D_t as a function of states and costates:

$$D_t = \tilde{D}\left(z_t, g_t, S_t, \lambda_t\right). \tag{10}$$

3 A Ramsey Problem

Models like the one in the previous section are often used to set up Ramsey (1927) problems, in which the government maximizes social welfare subject to the first order conditions of private agents. In the present context, this leads to the following Lagrangian for the government's problem:[13]

$$\mathcal{L}_t^g = E_t \sum_{i=0}^{\infty} \beta^i \Bigg\{ r^g \left(z_{t+i}, g_{t+i}, S_{t+i}, S_{t+i}, D_{t+i}, D_{t+i} \right)$$

$$+\pi^1_{t+i} \Big[S_{t+i+1} - b \left(z_{t+i}, g_{t+i}, S_{t+i}, S_{t+i}, D_{t+i}, D_{t+i} \right) \Big]$$

$$+\pi^2_{t+i} \left[\frac{\partial r(\cdot)}{\partial d_{t+i}} - \lambda_{t+i} \frac{\partial b(\cdot)}{\partial d_{t+i}} \right]'$$

$$+\pi^3_{t+i} \left[\lambda_{t+i} + \beta \frac{\partial r(\cdot)}{\partial s_{t+i+1}} - \beta \lambda_{t+i+1} \frac{\partial b(\cdot)}{\partial s_{t+i+1}} \right]' \Bigg\}, \tag{11}$$

where the household's control variables and state variables are replaced by their aggregate per capita counterparts. The government maximizes the representative agent's utility. This assumption is not necessary but it simplifies the analysis. The government maximizes the Lagrangian with respect to its control variables $\{g_{t+i}\}_{i=0}^{\infty}$, the aggregate equivalents of the private sector's control variables, and the Lagrange multipliers.

The force of the time inconsistency argument is made clear if we consider the first order conditions for the optimal choice of g_{t+1}. We have:

$$\frac{\partial \mathcal{L}_t^g}{\partial g_{t+1}} = 0 = E_t \Bigg\{ \beta \frac{\partial r^g(\cdot)}{\partial g_{t+1}} - \beta \pi^1_{t+1} \frac{\partial b(\cdot)}{\partial g_{t+1}}$$

$$+\beta \pi^2_{t+1} \left[\frac{\partial^2 r(\cdot)'}{\partial g_{t+1} \partial d_{t+1}} - \frac{\partial}{\partial g_{t+1}} \left(\frac{\partial b(\cdot)}{\partial d_{t+1}}' \lambda'_{t+1} \right) \right]$$

$$+\beta \pi^3_t \left[\frac{\partial^2 r(\cdot)'}{\partial g_{t+1} \partial s_{t+1}} - \frac{\partial}{\partial g_{t+1}} \left(\frac{\partial b(\cdot)}{\partial s_{t+1}}' \lambda'_{t+1} \right) \right] \Bigg\} \tag{12}$$

[13] As noted in the introduction, it is often possible to simplify the government's Lagrangian using the primal approach. This approach is not applicable to the highly abstract model presented here.

The term in π_t^3 gives the influence of future policy on the *current* behavior of households, via its effect on the forward-looking costate variables λ_t. If we allow the government to reoptimize at time $t+1$, the first order conditions for the choice of g_{t+1} become:

$$\frac{\partial \mathcal{L}_{t+1}^g}{\partial g_{t+1}} = 0 = \beta \frac{\partial r^g(\cdot)}{\partial g_{t+1}} - \beta \pi_{t+1}^1 \frac{\partial b(\cdot)}{\partial g_{t+1}}$$
$$+\beta \pi_{t+1}^2 \left[\frac{\partial^2 r(\cdot)'}{\partial g_{t+1} \partial d_{t+1}} - \frac{\partial}{\partial g_{t+1}} \left(\frac{\partial b(\cdot)}{\partial d_{t+1}}' \lambda_{t+1}' \right) \right] \tag{13}$$

Bygones are bygones. The effect of the government's controls at time $t+1$ on the household's actions at time t no longer appears. Even in the absence of unanticipated shocks, the government will in general revise its optimal plans.

Since the values of the private sector's costate variables λ_t are not pinned down by initial conditions, one of optimality conditions for the government's problem has to be

$$\pi_t^3 = 0.$$

The private sector's costates give the marginal value of the state variables to the representative agent's utility. Since the government's welfare function is just the utility function of the representative agent, a necessary condition to maximize welfare is that the contribution of a marginal change in these costates to welfare be zero. The future values of π_{t+i}^3, $i > 0$ are determined by the endogenous dynamics of the economy. After time t, they will only be zero by coincidence. However, if the government is allowed to reoptimize at time $t+i$, with $i > 0$, it will once again want to set

$$\pi_{t+i}^3 = 0.$$

In so doing, its optimal strategy changes. Time inconsistency arises here because the government's problem is not *recursive* in the sense of Sargent (1987, p.19). An agent's problem is recursive if its control variables dated t influence states dated $t+1$ and later and influence returns dated t and later. The household's current actions depend partly on its expectations of future government actions. In the Ramsey problem, the government's announced or future policies influence private agents' current behavior and therefore the current return via the function $r(\cdot)$.

4 Time-Consistent Control

Cohen and Michel (1988) studied a linear model with a single state variable and a quadratic objective function. They imposed a linear relationship between the predetermined state variable and the representative private agent's costate variable. They showed that the relationship is verified in equilibrium. Imposing it ties the government's hands. It is not allowed to choose its policy in order to set the initial marginal values of the costates equal to zero. If allowed to reoptimize, it is not tempted to change its policy in order to bring the marginal values of the costates back to zero.

This insight can be extended to nonlinear models as follows. In the present context, we can replace the last two terms associated with the π_t^3 constraint in the government's problem as

$$\text{Let} \quad \beta E_t \left(\frac{\partial r(\cdot)}{\partial s_{t+1}} - \lambda_{t+1} \frac{\partial b(\cdot)}{\partial s_{t+1}} \right)' = \phi\left(z_t, S_t\right). \tag{14}$$

Note that these terms include the value in time $t+1$ of the private agent's costate variables. As shown below, replacing these terms with a function of the model's state variables can be used with projection methods or perturbation methods to solve the model numerically, once we do the same thing for analogous terms from the government's first order conditions.[14] A byproduct of this is that λ_t, the vector of costate variables, becomes just a function of the current exogenous and endogenous state variables:

$$\lambda_t = -\phi\left(z_t, S_t\right)'. \tag{15}$$

This is just the nonlinear equivalent of the constraint imposed by Cohen and Michel (1988). It makes the government's problem recursive, as shown in the next section.

[14]See Judd (1998). The substitution is identical to that proposed by Den Haan and Marcet (1990). Their parameterized expectations solution technique is an example from the class of projection methods.

5 The Recursivity of the Government's Problem

We now assume that the government maximizes the utility of the representative private agent, as in the Ramsey problem described above, subject to the additional constraint given by (15). We can then show that the government's problem becomes recursive. We can write it as a dynamic programming problem in which the period-t return function does not depend on the future values of the government's controls.

We need one further assumption to demonstrate recursivity. We assume that the set of η_g equations associated with the π_t^2 constraint, once λ_t is replaced by $-\phi\left(z_t, S_t\right)$, can be solved out to find an explicit set of feedback rules for D_t which in equilibrium is just equation (4). Then, we can show the following:

Proposition: Subject to the constraint (15), the government's maximization problem is recursive.

Proof: Substituting in the constraint, we have the following expression for the difference between the government's Lagrangian at time t and the discounted value of its Lagrangian at time $t+1$, which gives the government's one-period return function:

$$\begin{aligned}\mathcal{L}_t^g - \beta E_t \mathcal{L}_{t+1}^g &= r^g\left(z_t, g_t, S_t, S_t, D_t, D_t\right) \\ &+ \pi_t^1\left(S_{t+1} - b\left(z_t, g_t, S_t, S_t, D_t, D_t\right)\right) \\ &+ \pi_t^2\left(\frac{\partial r(\cdot)}{\partial d_t} + \phi(z_t, S_t)' \frac{\partial b(\cdot)}{\partial d_t}\right). \end{aligned} \tag{16}$$

The one-period return function of the government does not depend directly or indirectly on g_{t+1}, since we suppose that D_t can be written as a function of only current state variables and g_t. The government's value function can be written as

$$\begin{aligned}V_t^g\left(z_t, S_t\right) = \max_{g_t} \Big\{ & r^g\left(z_t, g_t, S_t, S_t, D\left(z_t, g_t, S_t\right), D\left(z_t, g_t, S_t\right)\right) \\ & + \beta E_t\left(V_{t+1}^g\left(z_{t+1}, S_{t+1}\right)\right) \Big\}, \end{aligned} \tag{17}$$

q.e.d.

The maximization is subject to the law of motion of the aggregate state variables S_t, and to the first order conditions for the household's choice of its controls d_t, with the household's Lagrange multipliers λ_t substituted out using the constraint given in (15). Note that the government's problem becomes recursive partly because one of the underlying assumptions of this approach is that there is a time-invariant feedback rule for D_t which depends only on the current state of the economy. This leads to a feedback rule for the government compatible with (5) that depends only on the current state of the economy. It is as if the current government derives its optimal policy using dynamic programming techniques, under the assumption that all future governments will derive their optimal policies in the same way.[15]

In the context of the Ramsey problem described earlier, we have instead

$$\begin{gathered}\mathcal{L}_t^g - \beta E_t \mathcal{L}_{t+1}^g = r^g\left(z_t, g_t, S_t, S_t, D_t, D_t\right) \\ +\pi_t^1\left(S_{t+1} - b\left(z_t, g_t, S_t, S_t, D_t, D_t\right)\right) \\ +\pi_t^2\left(\frac{\partial r(\cdot)}{\partial d_t} - \lambda_t \frac{\partial b(\cdot)}{\partial d_t}\right)' \\ +\pi_t^3\left(\lambda_t + \beta E_t \frac{\partial r(\cdot)}{\partial s_{t+1}} - \beta E_t \lambda_{t+1} \frac{\partial b(\cdot)}{\partial s_{t+1}}\right)'.\end{gathered}$$

Because of the presence of the future value of the household's constraint λ_{t+1}, the government's problem fails to be recursive.

6 Numerical Solution

Using the Lagrangian in (11) above, after substituting out λ_t and eliminating the third constraint, the first order conditions for the government's problem can be

[15] One interpretation of optimal time-consistent policy is that the current government is playing a game against the private sector *and* future governments, taking the feedback rules of the future governments as given.

written as follows:

$$\frac{\partial \mathcal{L}_t^g}{\partial g_t} = 0 = \frac{\partial r^g(\cdot)}{\partial g_t} - \pi_t^1 \frac{\partial b(\cdot)}{\partial g_t}$$

$$+\pi_t^2 \Big[\frac{\partial^2 r(\cdot)'}{\partial g_t \partial d_t} + \frac{\partial}{\partial g_t} \left(\frac{\partial b(\cdot)'}{\partial d_t} \phi \left(z_t, S_t \right) \right) \Big], \tag{18}$$

$$\frac{\partial \mathcal{L}_t^g}{\partial S_{t+1}} = 0 = E_t \Bigg\{ \pi_t^1 + \beta \frac{\partial r(\cdot)}{\partial S_{t+1}} - \beta \pi_{t+1}^1 \frac{\partial B(\cdot)}{\partial S_{t+1}}$$

$$+\beta \pi_{t+1}^2 \frac{\partial^2 r(\cdot)}{\partial S_{t+1} \partial d_{t+1}} - \beta \pi_{t+1}^2 \frac{\partial}{\partial S_{t+1}} \left(\frac{\partial b(\cdot)}{\partial d_{t+1}}' \phi \left(z_{t+1}, S_{t+1} \right) \right) \Bigg\} \tag{19}$$

$$\frac{\partial \mathcal{L}_t^g}{\partial D_t} = 0 = \frac{\partial r^g(\cdot)}{\partial D_t} - \pi_t^1 \frac{\partial B(\cdot)}{\partial D_t}$$

$$+\pi_t^2 \Big[\frac{\partial^2 r(\cdot)'}{\partial D_t \partial d_t} + \frac{\partial}{\partial D_t} \left(\frac{\partial b(\cdot)}{\partial d_t}' \phi \left(z_t, S_t \right) \right) \Big]. \tag{20}$$

Several remarks are in order. First, the government's first order equations constitute a time-autonomous set of nonlinear difference equations.[16] Second, if the system is saddle-point stable, the initial conditions of the government's costate variables are those that place the system on the multi-dimensional convergent manifold of the system. The initial conditions of the costates therefore depend on the current state of the economy, given by the values of z_t and S_t. We can suppose that

$$E_t \Bigg\{ \beta \frac{\partial r(\cdot)}{\partial S_{t+1}} - \beta \pi_{t+1}^1 \frac{\partial B(\cdot)}{\partial S_{t+1}}$$

$$+\beta \pi_{t+1}^2 \frac{\partial^2 r(\cdot)}{\partial S_{t+1} \partial d_{t+1}} - \beta \pi_{t+1}^2 \frac{\partial}{\partial S_{t+1}} \left(\frac{\partial b(\cdot)}{\partial d_{t+1}}' \phi \left(z_{t+1}, S_{t+1} \right) \right) \Bigg\}'$$

$$= \psi \left(z_t, S_t \right). \tag{21}$$

[16] In the Ramsey problem, the optimality condition that the Lagrange multipliers related to private agent's costates be equal to zero at the moment the government optimizes, independently of the state of the economy, means that the resulting dynamical equation system is not time-autonomous. Its optimal policy is not a time-invariant function of the state of the economy, but rather depends on when it optimized. This is another way of interpreting the time inconsistency of optimal policy in the Ramsey problem.

We have

$$\pi_t^1 = -\psi\left(z_t, S_t\right)'. \tag{22}$$

Third, if the government's costates are functions of the current state of the economy, its policy variables are also implicit functions of the current state of the economy. The government's behavior is memoryless, and a solution to the dynamic equations that also satisfies the government's and private agents' optimality conditions is a Markov-perfect equilibrium. Fourth, the government's and private agents' Euler conditions will hold if equations (14) and (21) are satisfied. Fifth, for given parameterizations of the $\phi(\cdot)$ and $\psi(\cdot)$ functions, the dynamical system becomes recursive and can easily be solved numerically by iterating forward. Sixth, we do not have a general proof of existence or uniqueness of our Markov-perfect equilibrium. Krusell and Smith (2003) showed in a similar context that there is an indeterminacy of Markov-perfect equilibria. However, the multiple equilibria are associated with discontinuous decision rules. If we parameterize the $\psi(\cdot)$ function as a continuous function of the economy's state variables, then implicitly the government's policy rules are also continuous functions of the economy's state variables: Klein, Krusell and Ríos-Rull (2004) argue that imposing a differentiability requirement on the government's policy function is an important "refinement tool" for reducing the number of Markov-perfect equilibria. In some applications, such as the application presented below, proof of unicity will be straightforward.

The use of projection or perturbation methods in conjunction with control theory allows for an arbitrarily close approximation to the exact solution to the underlying problem. To illustrate how to solve the model numerically, we use a variation of the methodology described by Den Haan and Marcet (1990) and Marcet and Lorenzoni (1999). This method falls into the class of projection methods, and uses Monte Carlo methods to solve for the unknown parameters of the functional equations. The model can be simulated using the following steps:

- Parameterize the $\phi(\cdot)$ and $\psi(\cdot)$ functions using flexible functional forms such as polynomials or orthogonalized polynomials.

- Initialize the parameter values of these functions.
- For given parameter values of the $\phi(\cdot)$ and $\psi(\cdot)$ functions, simulate the model for a large number of time periods. Aside from the laws of motion for S_t and z_t, all of the equations that need to be solved are static. The laws of motion themselves are recursive.
- Estimate the parameters in the $\phi(\cdot)$ and $\psi(\cdot)$ functions by nonlinear regression, with the dependent variables being the series generated by numerical simulation, in order to minimize the sum of squared forecasting errors.
- Repeat the simulation and estimation steps until the change in the parameters of the expectations functions between iterations is sufficiently small.

7 Application

We apply the techniques developed in above to a simple model of optimal public spending. The utility function of the representative private agent depends on both private consumption spending and on government purchases. The government chooses public spending in order to maximize social welfare, which is just the expected utility of the representative private agent, It finances this spending via a proportional tax on total income.

The representative private agent maximizes expected utility, which is given by

$$U = E_t \sum_{i=o}^{\infty} \beta^i \left\{ \ln(c_{t+i}) + \mu \ln(g_{t+i}) \right\}, \tag{23}$$

where c_t is private consumption and g_t is public spending. The private agent holds the capital stock and rents it to firms. Its period budget constraint is given by

$$(1 - \tau_t)\left(w_t + (r_t - \delta)k_t\right) + k_t = c_t + k_{t+1}, \tag{24}$$

where w_t is the competitive real wage, r_t is the competitive real rental rate of capital, k_t is capital held by the individual, and τ_t is the rate of taxation on total income. The time endowment of the individual is normalized to equal one, so that before-tax labor income is just given by w_t.

The aggregate production function is given by

$$y_t = a_t {k_t}^{\alpha}, \tag{25}$$

where y_t is GDP. The law of motion for a_t is given by

$$\ln(a_t) = \rho \ln(a_{t-1}) + \varepsilon_t, \tag{26}$$

where ε_t is a white noise shock with variance σ_{ε}^2.

The government finances public investment via a proportional tax on total income. We rule out lump sum taxation in order to make the policy problem one of finding the second-best outcome, which leads to a distinction between time-consistent policies and time inconsistent policies. The government's budget is balanced in each period, so that

$$\tau_t \left(w_t + (r_t - \delta)k_t\right) = g_t. \tag{27}$$

The individual's first order conditions for utility maximization imply:

$$\frac{1}{c_t} = \lambda_t, \tag{28}$$

$$\lambda_t = \beta E_t \left(\lambda_{t+1} \left[1 + (1 - \tau_{t+1})(r_{t+1} - \delta)\right]\right), \tag{29}$$

$$k_{t+1} = (1 - \tau_t) y_t + (1 - \delta) k_t - c_t \tag{30}$$

The government's budget constraint can be used to substitute out public spending, so that τ_t is the only policy instrument. Under commitment, the government's maximization problem can be expressed as follows, after substituting out the representative agent's costate variable using $\lambda_t = 1/c_t$:

$$\mathcal{L} = E_t \sum_{i=0}^{\infty} \beta^i \Bigg\{ \ln(c_{t+i}) + \mu \ln(\tau_{t+i}) + \mu \ln(y_{t+i} - \delta k_{t+i})$$

$$+\pi_{t+i}^1 \left((1 - \delta) k_{t+i} + y_{t+i} - \tau_{t+i}(y_{t+i} - \delta k_{t+i}) - c_{t+i} - k_{t+i+1}\right)$$

$$+\pi^2_{t+i}\left(\frac{1}{c_{t+i}}-\frac{\beta}{c_{t+i+1}}\left(\left(1-\tau_{t+i+1}\right)\alpha\frac{y_{t+i+1}}{k_{t+i+1}}+1-\delta(1-\tau_{t+i+1})\right)\right)\Bigg\} \quad (31)$$

The first-order conditions imply:

$$\tau_{t+i}: \quad \frac{\mu}{\tau_{t+i}}-\pi^1_{t+i}(y_{t+i}-\delta k_{t+i})+\frac{\pi^2_{t+i-1}}{c_{t+i}}\left(\alpha\frac{y_{t+i}}{k_{t+i}}-\delta\right)=0, \quad i>0,$$

$$\tau_t: \quad \frac{\mu}{\tau_t}-\pi^1_t(y_t-\delta k_t)=0,$$

$$c_{t+i}: \quad \frac{1}{c_{t+i}}-\pi^1_{t+i}-\frac{\pi^2_{t+i}}{(c_{t+i})^2}$$

$$+\frac{\pi^2_{t+i-1}}{(c_{t+i})^2}\left(\left(1-\tau_{t+i}\right)\alpha\frac{y_{t+i}}{k_{t+i}}+1-\delta(1-\tau_{t+i})\right)=0, \quad i>0,$$

$$c_t: \quad \frac{1}{c_t}-\pi^1_t-\frac{\pi^2_t}{(c_t)^2}=0,$$

$$\pi^1_{t+i}: \quad k_{t+i+1}=(1-\tau_{t+i})y_{t+i}+(1-\delta(1-\tau_{t+i}))k_{t+i}-c_{t+i}, \quad i\geq 0,$$

$$\pi^2_{t+i}: \quad \frac{1}{c_{t+i}}$$

$$-\beta E_t\left(\frac{1}{c_{t+i+1}}\left(\left(1-\tau_{t+i+1}\right)\alpha\frac{y_{t+i+1}}{k_{t+i+1}}+1-\delta(1-\tau_{t+i+1})\right)\right)=0, \quad i\geq 0,$$

$$k_{t+i+1}: \quad \pi^1_{t+i}=-\beta E_{t+i}\left(\pi^2_{t+i}\frac{\alpha(\alpha-1)}{c_{t+i+1}}\frac{(1-\tau_{t+i+1})y_{t+i+1}}{k^2_{t+i+1}}\right)$$

$$+\beta E_t\left(\pi^1_{t+i+1}\left(1+\left(1-\tau_{t+i+1}\right)\left(\alpha\frac{y_{t+i+1}}{k_{t+i+1}}-\delta\right)\right)+\mu\frac{\alpha\frac{y_{t+i+1}}{k_{t+i+1}}-\delta}{y_{t+i+1}-\delta k_{t+i+1}}\right),$$

$$i\geq 0,$$

where y_t and a_t are defined respectively in (25) and (26).

It is straightforward to verify numerically whether this system of equations is saddlepoint stable, using a first-order approximation of the equilibrium conditions around its deterministic steady state. If local stability is satisfied, which is the case for our base-case calibration of the model and for a wide range of parameter values used for sensitivity analysis, then a solution that converges to the deterministic steady state is the unique solution that satisfies the transversality conditions for the optimization problems by private agents and the government.

7.1 Discretion

To find the time-consistent optimal policy under discretion (the absence of an ability to precommit), we replace the private agent's Euler equation by

$$\frac{1}{c_t} - \phi(a_t, k_t) = 0,$$

where $\phi(a_t, k_t)$ is a function to be approximated. After using the government's budget constraint to substitute out g_t, the Lagrangian can be written as:

$$\begin{aligned}\mathcal{L} \;=\; & E_t \sum_{i=0}^{\infty} \beta^i \Bigg\{ \ln(c_{t+i}) + \mu \ln(\tau_{t+i}) + \mu \ln(y_{t+i} - \delta k_{t+i}) \\ & + \pi^1_{t+i} \left((1-\delta)k_{t+i} + y_{t+i} - \tau_{t+i}(y_{t+i} - \delta k_{t+i}) - c_{t+i} - k_{t+i+1}\right) \\ & + \pi^2_{t+i} \left(\frac{1}{c_{t+i}} - \phi(a_t, k_t)\right) \Bigg\}.\end{aligned}$$

The first-order conditions imply:

$$\tau_t : \quad \frac{\mu}{\tau_t} - \pi^1_t (y_t - \delta k_t) = 0,$$

$$c_t : \quad \frac{1}{c_t} - \pi^1_t - \frac{\pi^2_t}{(c_t)^2} = 0,$$

$$\pi^1_t : \quad k_{t+1} = (1-\tau_t) y_t + (1 - \delta(1-\tau_t)) k_t - c_t,$$

$$\pi^2_t : \quad \frac{1}{c_t} - \phi\left(a_t, k_t\right) = 0,$$

$$\begin{aligned} k_{t+1} : \quad \pi^1_t = \beta E_t \Bigg\{ & \pi^1_{t+1} \left(1 + (1-\tau_{t+1}) \left(\alpha \frac{y_{t+1}}{k_{t+1}} - \delta\right)\right) - \\ & \pi^2_{t+1} \frac{\partial \phi}{\partial k_{t+1}} (a_{t+1}, k_{t+1}) + \mu \frac{\alpha \frac{y_{t+1}}{k_{t+1}} - \delta}{y_{t+1} - \delta k_{t+1}} \Bigg\}.\end{aligned}$$

7.1.1 Markov-perfect equilibrium

We define equilibrium in our model with discretion as follows. The following equations are satisfied in equilibrium:

$$\ln(a_t) = \rho \ln(a_{t-1}) + \varepsilon_t, \tag{32}$$

$$\pi_t^1 = \psi\left(a_t, k_t\right), \tag{33}$$

$$y_t = a_t {k_t}^{\alpha}, \tag{34}$$

$$c_t = \frac{1}{\phi\left(a_t, k_t\right)}, \tag{35}$$

$$\tau_t = \frac{\mu}{\pi_t^1 (y_t - \delta k_t)}, \tag{36}$$

$$\pi_t^2 = c_t - \pi_t^1 (c_t)^2, \tag{37}$$

$$k_{t+1} = (1-\tau_t) y_t + (1-\delta(1-\tau_t)) k_t - c_t, \tag{38}$$

and where, for a suitably-defined sample of artificial data based on simulating this system of equations and suitable parameterizations of the $\phi(\cdot)$ and $\psi(\cdot)$ functions, we have:

$$\phi(a_t, k_t) = \beta \widehat{E}_t \left\{ \left(\frac{1}{c_{t+1}} \right) \left[1 + (1-\tau_{t+1})(r_{t+1} - \delta) \right] \right\}, \tag{39}$$

$$\begin{aligned} \psi(a_t, k_t) = \beta \widehat{E}_t \Bigg\{ & \pi_{t+1}^1 \left(1 + (1-\tau_{t+1}) \left(\alpha \frac{y_{t+1}}{k_{t+1}} - \delta \right) \right) - \\ & \pi_{t+1}^2 \frac{\partial \phi}{\partial k_{t+1}} (a_{t+1}, k_{t+1}) + \mu \frac{\alpha \frac{y_{t+1}}{k_{t+1}} - \delta}{y_{t+1} - \delta k_{t+1}} \Bigg\}, \end{aligned} \tag{40}$$

where the $\widehat{E}_t$ operator refers to sample means. Equation (39) ensures that the representative private agent's Euler equation is satisfied to an arbitrary degree of accuracy. Equation (40) ensures that the government's Euler equation for the optimal choice of the capital stock is satisfied to an arbitrary degree of accuracy.

For given initial conditions (values of a_{t-1} and k_t) and given parameterizations of the $\phi(\cdot)$ and $\psi(\cdot)$ functions, the system of equations to solve for the economy's equilibrium for any period t is recursive. As long as the $\phi(\cdot)$ function is parameterized to avoid zero values, a solution exists. It can be shown that the solutions for output, aggregate consumption and the tax rate are unique.

It is clear that private agents' consumption (and hence savings) policies and the government's taxation policy are dependent only on the current values of the technology shock a_t and the capital stock k_t. In this sense, the equilibrium is a Markov-perfect equilibrium.

The substitution of the private agent's and government's Euler equations using the $\phi(\cdot)$ and $\psi(\cdot)$ functions renders the dynamical equation system completely recursive for given values of the parameters used to approximate the functions. Local stability of this system for suitable approximations of $\phi(\cdot)$ and $\psi(\cdot)$ means that agents' first order conditions hold and that transversality conditions are satisfied.[17]

7.2 Results

The parameter values used to simulate the model are summarized in Table 1. They are standard values used in the real business cycle literature. To solve the time-consistent and Ramsey problems, we use a version of the parameterized expectations algorithm (PEA) of Den Haan and Marcet (1990) and Marcet and Lorenzoni (1999). In both cases, we need to find two interpolating functions (one for each Euler equation). We describe in detail the methodology for the time-consistent problem. There are two state variables, k_t and a_t, so that the two interpolating functions, ϕ and ψ, should be functions of both k_t and $a_t,$ and verify

$$\frac{1}{c_t} - \phi\left(a_t, k_t\right) = 0 \quad \text{and} \quad \pi_t^1 - \psi\left(a_t, k_t\right) = 0.$$

We derive results for both first-order and second-order polynomial approximations. The first order approximations are given by:

$$\phi(\cdot) = \exp(\beta_1 + \beta_2 \ln(K_t) + \beta_3 \ln(a_t));$$

$$\psi(\cdot) = \exp(\theta_1 + \theta_2 \ln(K_t) + \theta_3 \ln(a_t)).$$

The second-order approximations are given by:

$$\phi(\cdot) = \exp(\beta_1 + \beta_2 \ln(K_t) + \beta_3 \ln(a_t) + \beta_4 (\ln(K_t))^2 + \beta_5 (\ln(a_t))^2 + \beta_6 \ln(K_t) \ln(a_t));$$

[17] A proof of the unicity of the Markov-perfect equilibrium for this application is available on request.

$$\psi(\cdot) = \exp(\theta_1 + \theta_2 \ln(K_t) + \theta_3 \ln(a_t) + \theta_4 (\ln(K_t))^2 + \theta_5 (\ln(a_t))^2 + \theta_6 \ln(K_t) \ln(a_t)).$$

Below, we report results for DM-tests (see Den Haan and Marcet, 1994) of the adequacy of the approximations. We use nonlinear least squares to estimate the parameters of the polynomials, and use a dampening coefficient $0 < \lambda \leq 1$ to foster convergence. We iterate on:

$$\psi_{s+1} = (1 - \lambda)\psi_s + \lambda\hat{\psi},$$

where $\hat{\psi}$ gives the current vector of estimates of the parameters of the approximations. As starting values, we use the steady-state (and the decision rules) of the commitment solution. The estimated parameters of the PEA functions are given in Table 2 below.

Figure 1 displays impulse-response functions (for the first-order PEA functions) in response to a one-standard-deviation shock to technology. The horizontal axis measures the number of quarters after the shock and the vertical axis measures the deviations in logs from the steady state. Following a positive shock to technology, output increases as does consumption. The tax rate initially falls, so that the initial impact on public spending is proportionately lower than the impact on private consumption spending. The gradual accumulation of capital in response to the shock imparts a hump-shaped response of both private consumption spending and public spending. These two variables peak just before the capital stock reaches its maximum level, and then the capital stock, public spending, private consumption, and output all converge to their steady state levels from above. The response of the tax rate is non-monotonic. After an initial drop, the tax rate surpasses its steady-state value shortly after the capital stock peaks, and then converges to the steady state from above.

As explained before, the form of the PEA is critical. In this context, we proceed in two steps. First, we compare deterministic simulations using the first-order and second-order PEA. Second, we develop the adequacy test, as suggested by Den Haan and Marcet (1994) in a multivariate setting. Figure 2 reports the deterministic paths of variables of interest under the optimal (discretionary) policy, assuming an initial capital stock that is below its steady-state level. The dotted lines represent the results from the second-order approximation. A visual inspection shows that the difference in polynomial orders matters only slightly in this application.

The PEA functions are approximated in order to satisfy orthogonality conditions. If the approximations are adequate, the realized future values of the nonlinear functions of future variables to which the conditional expectations operator is applied should be unpredictable given information available at time t. This is the rationale of the DM test proposed by Den Haan and Marcet (1994). To implement the test, we regressed the realized time $t+1$ forecast errors of the PEA functions on a constant, the capital stock, and the technology shock, both measured at time t. We also tried regressions in which we added four lags of the capital stock and four lags of the technology shock. The results were robust to the specification of the regressions.

To reduce the probability of a type I error (rejection of the null when it is true: see explanations in Heer and Maussner, page 502), we follow the procedure of Den Haan and Marcet (1994). For a given sequence of shocks, we first compute the two approximate solutions ϕ and ψ for a large T ($T = 5000$). Second we use this solution and draw a new sequence of shocks such that the corresponding sample size, T_1 is smaller than T. After computing the time path of the variables of interest, the DM statistic can be calculated for these observations. Then, this procedure is repeated very often and we can compute the percentage of the DM-statistics that are below the lower or above the upper 2.5 percent critical values of the $\chi^2(m)$ distribution.

In the benchmark case and for first-degree approximations of the PEA functions, the DM-statistic was 8.6746, which leads us to accept (at a five-percent level) the the null hypothesis that the forecast errors are unpredictable based on the time t information set. The number of DM-statistics (out of 500) below (above) the 2.5^{th} and 97.5^{th} percentile of a χ^2 distribution with 6 degrees of freedom was close to the theoretical five percent.

Given that the first-order PEA functions are adequate, we compared the deterministic convergence to the steady-state under commitment and discretion. The solid lines in Figure 3 reproduce the time paths from Figure 2 and illustrate the convergence to the steady state of the model under optimal discretionary (time consistent) policy. The dotted lines illustrate the convergence of the model under optimal policy with commitment. The steady states are different under the two types of optimal policy, and the behavior of consumption, public expenditures and the tax rate is substantially different when the government first optimizes. Overall, taxes are higher under commitment but this leads to higher public expenditures, which increases the conditional welfare of the representative agent.

Finally, to assess the differences between the commitment and discretion solutions, we calculated conditional welfare. This is measured at time t (when the government optimizes) for the same initial conditions (same technology level, same level of the capital stock, which is below its deterministic steady-state level under either discretion or commitment). Table 3 gives values for conditional welfare.

8 Conclusions

The methodology proposed in this paper is quite general. It leads to systems of dynamical equations which can easily be simulated with available computer technology and relatively parsimonious numerical solutions using projection or perturbation techniques. Deriving time-consistent government policies using these methods is conceptually as straightforward as solving Ramsey problems. The technique should allow researchers to do normative analysis, comparing the levels of welfare attainable with and without precommitment by the government. It should also be useful for positive analysis, for example comparing the predictions of a given model for comovements between government policy variables and other macroeconomic aggregates with and without precommitment. As suggested by Judd (1998), with current advances in computer technology it should become more and more common to use numerical methods to advance our understanding of economic theory.

References

Ambler, Steve and Emanuela Cardia (1997), "Optimal Government Spending in a Business Cycle Model," *Business Cycles and Macroeconomic Stability: Should We Rebuild Built-In Stabilizers?* Jean-Olivier Hairault, Pierre-Yves Hénin and Franck Portier, editors, Kluwer Academic Press

Ambler, Steve and Alain Paquet (1996), "Fiscal Spending Shocks, Endogenous Government Spending and Real Business Cycles," *Journal of Economic Dynamics and Control* 20, 237-256

Ambler, Steve and Alain Paquet (1997), "Recursive Methods for Computing Equilibria of General Equilibrium Dynamic Stackelberg Games," *Economic Modelling* 14, 155-173

Aruoba, S. Boragan, Jess Fernndez-Villaverde and Juan Francisco Rubio-Ramírez (2003), "Comparing Solution Methods for Dynamic Equilibrium Economies," working paper 2003-27, Federal Reserve Bank of Atlanta

Azzimonti-Renzo, Marina, Pierre-Daniel Sarte and Jorge Soares (2003), "Optimal Public Investment with and without Government Commitment," working paper 03-10, Federal Reserve Bank of Richmond

Benhabib, Jess and Aldo Rustichini (1997), "Optimal Taxes without Commitment," *Journal of Economic Theory* 77, 231-259

Benhabib, Jess, Aldo Rustichini and Andrés Velasco (1996), "Public Capital and Optimal Taxes without Commitment," C.V. Starr Center working paper 96-19, New York University

Benhabib, Jess and Andrés Velasco (1996), "On the Optimal and Best Sustainable Taxes in an Open Economy," *European Economic Review* 40, 135-154

Bernheim, B. Douglas and Debraj Ray (1989), "Markov Perfect Equilibria in Altruistic Growth Economies with Production Uncertainty," *Journal of Economic Theory* 47, 195-202

Chari, V.V. and Patrick J. Kehoe (1990), "Sustainable Plans," *Journal of Political Economy* 98, 783-802

Chari, V.V. and Patrick J. Kehoe (1999), "Optimal Fiscal and Monetary Policy," *Handbook of Macroeconomics*. John Taylor and Michael Woodford, editors, Amsterdam, North Holland

Cohen, Daniel and Philippe Michel (1988), "How Should Control Theory Be Used to Calculate a Time-Consistent Government Policy?," *Review of Economic Studies* 55, 263-274

Currie, David and Paul Levine (1993), *Rules, Reputation and Macroeconomic Policy Coordination*. Cambridge, Cambridge University Press

Den Haan, W.J. and A. Marcet (1990), "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business and Economic Statistics* 8, 31-34

Den Haan, W.J. and A. Marcet (1994), "Accuracy in Simulations," *Review of Economic Studies* 61, 3-18

Hansen, Gary D. and Edward C. Prescott (1995), "Recursive Methods for Computing Equilibria of Business Cycle Models," in Thomas F. Cooley (ed.), *Frontiers of Business Cycle Research*. Princeton, Princeton University Press

Heer, Burkhard and Alfred Maussner (2005), *Dynamic General Equilibrium Modelling: Computational Methods and Applications*. Springer Verlag

Judd, Kenneth L. (1998), *Numerical Methods in Economics*. Cambridge, MA, MIT Press

Kim, Jinill and Sunghyung Kim (2003), "Spurious Welfare Reversals in International Business Cycle Models," *Journal of International Economics* 60, 471-500

Klein, Paul, Per Krusell and José-Victor Ríos-Rull (2004), "Time-Consistent Public Expenditures," working paper 4582, Centre for Economic Policy Research

Ljungqvist Lars, and Thomas Sargent (2000), *Recursive Macroeconomic Theory*. Cambridge, MA, MIT Press

Lucas, Robert E. and Nancy Stokey (1983), "Optimal Fiscal and Monetary Policy in an Economy without Capital," *Journal of Monetary Eocnomics* 12, 55-93

Marcet, Albert and Giovanni Lorenzoni (1999), "The Parameterized Expectations Approach: Some Practical Issues," in R. Marimon and A. Scott, eds., *Computational Methods for the Study of Dynamic Economies* Oxford, Oxford University Press

Maskin, Eric and Jean Tirole (1993), "Markov Perfect Equilibrium," mimeo, Harvard University

McGrattan, Ellen (1999), "Application of Weighted Residual Methods to Dynamic Economic Models," *Computational Methods for the Study of Dynamic Economies*. Ramn Marimon and Andrew Scott, editors, London, Oxford University Press

Miller, Marcus and Mark Salmon (1985), "Policy Coordination and the Time Inconsistency of Optimal Policies in Open Economies", *Economic Journal* 95, 124-137

Ortigueira, Salvador (2004), "Markovian Optimal Taxation: The Macroeconomic Implications of Two Notions of Time Consistency," draft, Department of Economics, European University Institute

Oudiz, G. and J. Sachs (1985), "International Policy Coordination in Dynamic Macroeconomic Models," *International Economic Policy Coordination*. Willem Buiter and Richard Marston, editors, Cambridge, Cambridge University Press

Ramsey, Frank (1927), "A Contribution to the Optimal Theory of Taxation," *Economic Journal* 37, 47-61

R'ıos-Rull, José-Victor (1995), "Models with Heterogeneous Agents,"*Frontiers of Business Cycle Research*. Thomas Cooley, editor, Princeton, Princeton University Press

Sargent, Thomas J. (1987), *Dynamic Macroeconomic Theory*. Cambridge, MA, Harvard University Press

Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, Princeton University Press

Table 1: Model Parameter Values

Parameter	Value
β	0.987
α	0.300
μ	0.300
δ	0.025
ρ	0.950
σ_ε	0.030

Table 2: Model Parameter Estimates

Parameter	β_1	β_2	β_3	β_4	β_5	β_6
$\phi(\cdot)$	0.8877	-0.4796	-0.4699	-	-	-
$\phi(\cdot)$	0.8903	-0.4827	-0.5348	0.011	-0.2480	0.0374
Parameter	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6
$\psi(\cdot)$	1.0804	-0.5472	-0.4812	-	-	-
$\psi(\cdot)$	0.8988	-0.3740	-0.8939	-0.0403	-0.2375	0.1886

Table 3: Conditional Welfare

commitment	-17.4792
discretion, 1st-order approximation	-17.7404
discretion, 2nd-order approximation	-17.7280

Figure 1: Impulse-response functions after a technology shock

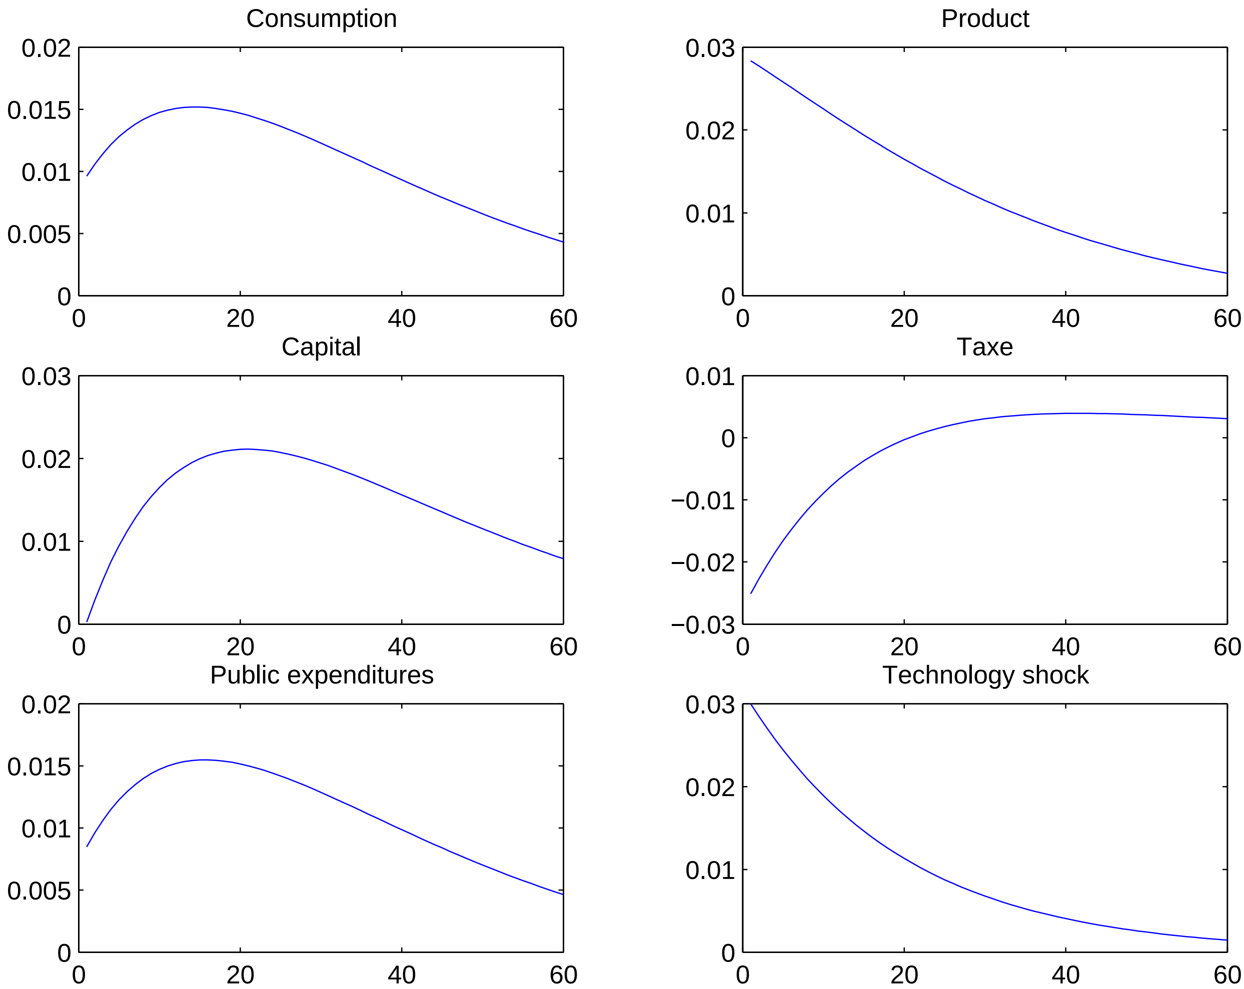

Figure 2: First-order versus Second-order PEA

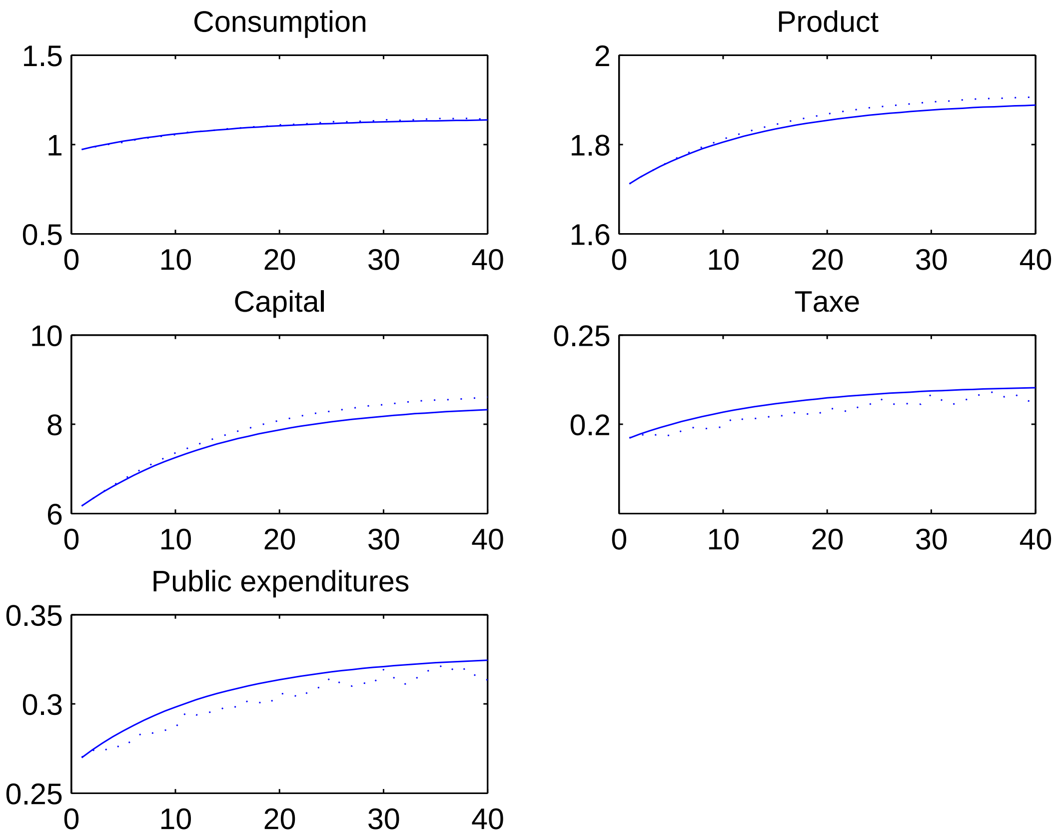

Figure 3: First-order PEA versus Commitment