Optimal Fiscal Policy When Agents are Learning*

Francesco Caprioli†

15 February 2010

Abstract

In this paper I investigate the optimal fiscal policy when private agents are boundedly rational. A benevolent government has to choose taxes on labor income and one-period state-contingent bonds to finance public spending, taking into account the expectation formation mechanism of private agents. The main result I find is that the government should use fiscal variables to manipulate agents' expectations. I 'rationalize' the popular view that in periods of pessimism the government should reduce taxes and increase public spending, and vice versa in periods of optimism. Moreover, the expectation-dependent fiscal plan prescribes taxes which are less smooth than in a fully rational expectations framework. Finally, I re-examine the validity of some tests for market completeness and debt sustainability in light of my results.

JEL Classification: E62, H63, D83.

Keywords: Optimal Taxation, Learning, Market Completeness.

*I am heavily indebted to Albert Marcet for his excellent supervision and encouragement. I would like to thank Klaus Adam, Richard Anton Braun, Vasco M. Carvalho, Omar Licandro, Eva-Carceles Poveda, Jordi Galí, Mark Giannoni, Bruce Preston, Michael Reiter and Thijs Van Rens for their very useful comments and suggestions. I am also grateful to Sofia Bauducco, Filippo Ferroni, Stefano Gnocchi, Antonio Mele and Andrea Tesei for detailed feedback, which greatly improved this work. This paper was awarded the Best Paper mention in the XIII Workshop on Dynamic Macroeconomics, Vigo, Spain.

†E-mail: francesco.caprioli@upf.edu

1 Introduction

An important issue in public finance theory is how to collect revenues to pay for government expenditures. When lump-sum transfers are not available, fiscal authorities must resort to taxes which distort people's decisions and move the economy away from the first-best. The optimal taxation literature [1] focuses on identifying the tax profile which minimizes the associated distortionary costs. The key insight of this literature is that taxes should be smooth across time.

This conclusion has been derived under the assumption of rational expectations (RE). Under this assumption households have a complete knowledge of the systematic aspects of the economy, including the way in which the policymaker sets its policy. Although very useful, the RE assumption is quite strong, as in many real-world situations the households' knowledge may not be so deep. Structural changes, like the stage 3 of the EMU, or the recent crisis, are examples in which this assumption is at least questionable. Moreover, Sargent (2005) provides historical examples of ambiguity in American monetary and fiscal policy. Then a natural question arises: What is the optimal fiscal policy when households do not know exactly how the government sets distortionary taxes?

This paper offers an answer to this question. I consider a closed production economy with no capital and infinitely lived agents. I start assuming that public spending is an exogenous shock, as it is the usual reference point in the public finance literature.[2] To finance the public spending, the government levies a proportional tax on labor income and has access to a complete set of one-period state-contingent bonds. The only difference between this framework and the Lucas and Stokey (1983) one is that households do not know the mapping between aggregate shock and tax rate, but instead they learn it. As their expectations about the tax rate are not model-consistent, their expectations about next-period's marginal utility of consumption are neither. Households act like econometricians and to forecast next period's contingent marginal utility of consumption they use a weighted average of past values of it. Given the realization of the shock, each period they update their belief about the marginal utility of consumption contingent on that specific realization.

The government is benevolent and chooses distortionary taxes on labor income and state-contingent debt to maximize households' expected utility, subject to the feasibility constraint, households' optimality conditions and the way in which they update their beliefs.

I find that the government should set fiscal variables to manipulate private agents' expectations. To give an intuition, assume that the public expenditure is constant and that the govern-

[1] The optimal taxation literature is immense and offering a comprehensive survey goes beyond the scope of this paper. See Barro (1979, 1989, 1995, 1997), Bohn (1990), Kydland and Prescott (1980), Lucas and Stokey (1983), Chari et al. (1994), Chari and Kehoe (1999), Aiyagari et al. (2002), Zhu (1992) among many other.

[2] Later on I extend the analysis to the case in which the government decides the amount of public consumption.

ment has zero initial wealth. Under rational expectations, the optimal fiscal rule prescribes a balanced budget: the government sets the tax rate to collect enough revenues to finance expenditure. When agents are pessimistic, (i.e. they expect the one-step-ahead tax rate to be higher than they would expect it to be if they were fully rational) government optimality conditions require public expenditure to be financed mainly through debt. In this way the low current tax rate induces agents to revise downwards their expectations about the next period's tax rate.[3] In the long-run the tax rate is higher than in a rational expectations framework because the government has to finance the interest paid on a positive amount of debt.[4]

In this sense the agents' initial beliefs have an effect on the long-run mean value of the tax rate and debt: the more pessimistic (optimistic) the agents are, the higher is the government debt (wealth) in the long-run. One implication of this result is that the model can help explain the wide dispersion across countries in the level of government debt and tax rate.

As expectations are not model-consistent, taxes are less smooth than under rational expectations. The reason is that the government has to minimize the welfare costs associated with distortionary taxes on one hand, and with distorted expectations on the other. When expectations are rational only the first distortion is present, and to minimize the associated losses taxes have to be smooth. But when both distortions are present, this is no longer optimal. The case-study of a perfectly anticipated war is a clear example of the tension between the two conflicting goals the government wants to achieve, tax smoothing on one hand and manipulation of beliefs on the other. Under rational expectations it is optimal for the government to accumulate assets before the war and sell them during the war-time. In this way the tax rate is constant in all periods before and after the war. By contrast, in a learning framework pessimistic agents do not trust the promises made by the government of higher-than-expected future consumption. The government sets low tax rates to manipulate agents' expectations, accumulating less assets (than in a rational expectations framework) before the war. As a consequence, the war is financed issuing more debt than in a rational expectations framework. The tax rate after the big shock is much higher than before.[5]

Since tax rates and debt have a unit-root behavior, bounded rationality affects the power of some widely used tests to check for market completeness and debt sustainability. In line with Marcet and Scott (2008) I find that looking at the behaviour of debt is a much more reliable way to test the bond market structure than looking at the behavior of tax rates. Similarly, the standard unit-root test in the debt/GDP ratio used to discriminate between responsible and

[3]One implication of this result is that restricting how much a government can become indebted can delay the learning process.

[4]The analysis is symmetric for the case of optimistic agents.

[5]Manipulating expectations can explain why a benevolent government should run a deficit during peacetime periods, an implication that the Lucas and Stokey (1983) does not have and for which has been criticized.

non-responsible governments can be misleading, since it may cause a fiscal policy plan to be declared unsustainable when instead it is sustainable by construction. Augmenting this test to include the primary surplus in the regressors is a sharper way to distinguish the optimal and sustainable fiscal policy from an unsustainable policy.

Finally, I extend the model to the case in which the government chooses public consumption. I find that, when agents are pessimistic, the fiscal authority increases public spending above the rational expectations level, financing it mainly through debt. This conclusion is in line with some proposals to deal with the recent distress.

Many authors have studied the impact of learning on monetary policy design, either when the central bank follows some ad hoc policy rules (see *inter alia* Orphanides and Williams (2006), Preston (2005a,b, 2006), Preston and Eusepi (2007b,a)) or when it implements the optimal monetary policy (see *inter alia* Evans and Honkapohja (2003, 2006), Molnar and Santoro (n.d.)). Perhaps surprisingly, fiscal policy has received much less attention. Evans et al. (2007) study the interest rate dynamic learning path in a non-stochastic economy in which the fiscal authority credibly announces a future change in government purchases. Karantounias et al. (2007) and Svec (2008) study the optimal fiscal policy when agents do not trust the transition probabilities of the public expenditure suggested by their approximating model. Up to my knowledge, this is the first paper studying the influence of learning on fiscal policy design.

The paper proceeds as follows. Section 2 solves for the optimal fiscal policy under learning. In Section 3 I characterize the fiscal plan restricting the government expenditure shock to a specific form. Section 4 gives some policy implications. In section 5 I extend the basic model to the case of endogenous government expenditure. Section 6 deals with the problem of discriminating between a complete markets model with learning and an incomplete markets model with rational expectations. Section 7 focuses on debt sustainability and debt limits. Section 8 concludes.

2 The Model

I consider an infinite horizon economy where the only source of aggregate uncertainty is represented by a government expenditure shock.[6] Time is discrete and indexed by $t = 0, 1, 2....$ In each period $t \geq 0$ there is a realisation of a stochastic event $g_t \in G$. The history of events up and until time t is denoted by $g^t = [g_t, g_{t-1}, g_{t-2}, ..., g_0]$. The conditional probability of g^r given g^t is denoted by $\pi(g^r|g^t)$. For notational convenience, I let $\{x\} = \{x(g^t)\}_{g^t \in G}$ represent the entire state-contingent sequence for any variable x throughout the paper.

[6] In section 5 I extend the analysis to the case in which the fiscal authority chooses the amount of public consumption.

In subsection 2.1 I briefly review the Lucas and Stokey (1983) model. As it is the workhorse in the optimal taxation literature, it is the benchmark model against which I will assess my results. The economy is populated by a representative household and a government. To finance an exogenous stream of public consumption, the government levies a proportional tax on labor income and has access to a complete set of one-period state-contingent bonds. Both the household and the government have rational expectations. The solution to this model states policy rules for labor tax rate and bond-holdings which maximize households' welfare subject to the restriction that taxes are distortionary. In subsection 2.2 I study the problem of a government which internalizes the fact that agents are boundedly rational.

2.1 The Lucas and Stokey (1983) model

Consider a production economy where the technology is linear in labor. The household is endowed with 1 unit of time that can be used for leisure and labor. Output can be used either for private consumption or public consumption. The resource constraint is

$$c_t + g_t = 1 - l_t \tag{1}$$

where c_t, l_t and g_t denote respectively private consumption, leisure, and public consumption. The problem of the household is to maximise his lifetime discounted expected utility

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t) \tag{2}$$

subject to the period-by-period budget constraint

$$b_{t-1}(g_t) + (1 - \tau_t)(1 - l_t) = c_t + \sum_{g^{t+1}|g^t} b_t(g_{t+1}) p_t^b(g_{t+1}) \tag{3}$$

where β is the discount factor, τ_t is the state-contingent labour tax rate and $b_t(g_{t+1})$ denotes the amount of bonds issued at time t contingent on period $t+1$ government shock at the price $p_t^b(g_{t+1})$. $vb_t \equiv -\sum_{g^{t+1}|g^t} b_t(g_{t+1}) p_t^b(g_{t+1})$ is defined as the value of government debt.

The household's optimality condition are

$$1 - \tau_t = \frac{u_{l,t}}{u_{c,t}} \tag{4}$$

$$p_t^b(g_{t+1}) = \beta \frac{u_{c,t+1}(g_{t+1})}{u_{c,t}} \pi(g^{t+1}|g^t) \tag{5}$$

together with the budget constraint 3.

The government pursues an optimal taxation approach: given an initial amount of inherited debt, b_{-1}^g, she chooses the sequence of tax rates and state-contingent bonds to maximise consumer's welfare. The solution to this dynamic optimal taxation problem is called a Ramsey plan. Lucas and Stokey (1983) show that under complete markets and rational expectations the Ramsey plan has to satisfy the following restriction

$$E_0 \sum_{t=0}^{\infty} \beta^t (u_{c,t} c_t - u_{l,t}(1 - l_t)) = u_{c,0} b_{-1} \tag{6}$$

which can be thought of as the intertemporal consumer budget constraint with both prices and taxes replaced by the households' optimality conditions, (4) and (5). Constraint (6) is the implementability condition. The Ramsey plan satisfies

$$\tau_t = T(g_t, b^g_{-1}) \forall t > 0 \tag{7}$$

$$b^g_t(g_{t+1} = \bar{g}) = D(\bar{g}, b^g_{-1}) \forall t > 0 \tag{8}$$

$$vb^g_t = V(g_t, b^g_{-1}) \forall t > 0 \tag{9}$$

The allocation is a time invariant function of the only state variable in this model, g_t. The initial holding of government bonds matters for the allocation because it determines the value of the Lagrange multiplier attached to the implementability condition. The state-contingent bond holding is a time invariant function and does not depend on the current state of the economy, and the market value of debt is influenced by the current shock only through variations in the state-contingent interest rates.

2.2 Optimal Fiscal Policy When Agents Are Learning

In this subsection we describe how households form expectations about variables which are relevant for their decision problem. The households' optimality conditions, which we repeat for convenience, are

$$\frac{u_{l,t}}{u_{c,t}} = 1 - \tau_t \tag{10}$$

$$p^b_t(g_{t+1}) = \beta \frac{\tilde{u}_{c,t+1}(g_{t+1})}{u_{c,t}} \pi(g^{t+1}|g^t) \tag{11}$$

The implementation of equation (11) requires agents to forecast their own state-contingent consumption one-period-ahead. This approach of modeling boundedly rational behavior may seem strange at first glance, but it is commonly used in the learning literature (see Evans et al. (2003), Carceles-Poveda and Giannitsarou (2007), Milani (2007) among many others). It is a very useful short-cut to model households' lack of knowledge about market determined variables, which are outside of agents' control although they are relevant to their decision problem. In the current setup, non-rational expectations about future consumption can be interpreted as non rational expectations about the tax policy rule followed by the government. In fact, forwarding equation (10) one period, agents understand that consumption at $t+1$ depends on the tax rate the government will set at $t+1$; as far as expectations about tax rate are not-model consistent, expectations about consumption are neither.

In order to simplify the analysis I assume that the government expenditure shock can take only two realizations, g_H and g_L, with $g_H > g_L$.

Let $\gamma_t^i \equiv \tilde{u}_{c,t+1}(g_{t+1} = g^i)$ for $i = H, L$. Agents update their beliefs about next-period state-contingent marginal utility of consumption according to the following scheme

$$\gamma_t^i = \begin{cases} \gamma_{t-1}^i + \alpha_t(u_{c,t}(g_t = g_i) - \gamma_{t-1}^i), & \text{if } g_t = g_i \\ \gamma_{t-1}^i, & \text{if } g_t = g_j \end{cases} \tag{12}$$

with $i = H, L$ and where α_t follows an exogenous law of motion.[7]

Definition 1. *A competitive equilibrium with boundedly rational agents is an allocation $\{c_t, l_t, g_t\}_{t=0}^{\infty}$, state-contingent beliefs about one-step-ahead marginal utility of consumption $\{\gamma_t^i\}_{t=0}^{\infty}$ for $i = H, L$, a price system $\{p_t^b\}_{t=0}^{\infty}$ and a government policy $\{g_t, \tau_t, b_t\}_{t=0}^{\infty}$ such that (a) given the price system, the beliefs and the government policy the households' optimality conditions are satisfied; (b) given the allocation and the price system the government policy satisfies the sequence of government budget constraint* (3)*; and (c) the goods and the bond markets clear.*

Let

$$x_t = [\gamma_t^H I(g_{t+1} = g_H) + \gamma_t^L I(g_{t+1} = g_L)] \tag{13}$$

where I is the indicator function and define

$$A_t \equiv \prod_{k=0}^{t} \frac{x_{k-1}}{u_{c,k}} \tag{14}$$

Taking logs to both sides we get

$$logA_t = \sum_{k=0}^{t} (log(x_{k-1}) - log(u_{c,k})) \tag{15}$$

[7]In Appendix A.6 we discuss a measure of the 'quality' of this learning scheme.

The log of A_t is the sum of the log-differences between expected and actual marginal utility of consumption from period 0 to period t. This variable has a very natural interpretation as the sum of all past forecast errors agents have made up to period t in predicting next-period log consumption. Under rational expectations, A_t is constant and equal to 1, while under learning it is not, unless the initial beliefs coincide with the rational expectations ones.

As shown in subsection 2.1, using households' optimality conditions to substitute out prices and taxes from the government budget constraint, Lucas and Stokey (1983) show that under complete markets and rational expectations the competitive equilibrium imposes one single intertemporal constraint on allocations. Using a similar argument, I show that under complete markets and bounded rationality the following result holds.

Proposition 1. *Assume that for any competitive equilibrium $\beta^t A_t u_{c,t} \to 0$ a.s.*[8] *Given b_{-1}, γ^H_{-1}, γ^L_{-1}, a feasible allocation $\{c_t, l_t, g_t\}_{t=0}^{\infty}$ is a competitive equilibrium if and only if the following constraint is satisfied*

$$E_0 \sum_{t=0}^{\infty} \beta^t A_t (u_{c,t} c_t - u_{l,t}(1 - l_t)) = A_0 u_{c,0} b_{-1} \tag{16}$$

with initial condition $A_{-1} = 1$

Proof. We relegate the proof to the appendix. □

Equation (16) is the bounded rationality version of the implementability condition derived in equation 6. The difference between equations (16) and (6) arises through the effect that out-of-equilibria expectations exert on state-contingent prices. As expectations are not model-consistent, the primary surplus at time t, expressed in terms of marginal utility of consumption, is weighted by the product of ratios of expected to actual marginal utility from period 0 till period t.

2.3 The government problem

Using the *primal approach* to taxation we recast the problem of choosing taxes and state-contingent bonds as a problem of choosing allocations maximizing households' welfare over competitive equilibria. At this point a clarification is needed. When the households and the benevolent government share the same information, they maximize the same objective function. But when the way in which they form their expectations differ, as in this setup, their objective functions differ as well. Therefore it is no longer obvious which objective function the benevolent government should maximize. In what follows I assume that it maximizes the representative consumer's welfare *as if* he were rational. Two reasons justify this assumption. First, as agents

[8] Using the results of Proposition 3 we show that this is actually the case.

form model-consistent expectations in the long-run, in the long-run agents are going to be fully rational. Second, the government understands how agents behave and form their beliefs, and it understands that these beliefs are distorted. Consequently, it uses this information to give the allocation which is best for them from an objective point of view. This is consistent with a paternalistic vision of the government.[9]

Definition 2. *The government problem under learning is*

$$\max_{\{c_t, l_t, \gamma_t^H, \gamma_t^L, A_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t)$$

subject to

$$E_0 \sum_{t=0}^{\infty} \beta^t A_t (u_{c,t} c_t - u_{l,t}(1 - l_t)) = A_0 u_{c,0} b_{-1} \tag{17}$$

$$A_t = A_{t-1} \frac{\gamma_{t-1}^H I(g_t = g^H) + \gamma_{t-1}^L I(g_t = g^L)}{u_{c,t}} \tag{18}$$

$$\gamma_t^i = \begin{cases} \gamma_{t-1}^i + \alpha_t (u_{c,t}(g_t = g_i) - \gamma_{t-1}^i), & \text{if } g_t = g_i \\ \gamma_{t-1}^i, & \text{if } g_t = g_j \end{cases} \tag{19}$$

$$c_t + g_t = 1 - l_t \tag{20}$$

Equation (17) constraints the allocation to be chosen among competitive equilibria. Equation (18) is the recursive formulation for A_t, obtained directly from its definition, given in equation (14). Equation (19) gives the law of motion of beliefs. Equation (20) is the resource constraint. Since A_t and γ_t^i for $i = L, H$ have a recursive structure, the problem becomes recursive adding A_{t-1} and γ_{t-1}^i for $i = L, H$ as state variables.

Leaving the details about the derivation in appendix A.2, first order necessary conditions[10] with respect to consumption and leisure impose that

$$\begin{aligned} & u_{c,t} + \Delta A_t (u_{cc,t} c_t + u_{c,t}) - \lambda_{1,t} \alpha_t u_{cc,t} I(g_t = g^H) \\ & - \lambda_{2,t} \alpha_t u_{cc,t} I(g_t = g^L) - \Delta \frac{u_{cc,t}}{u_{c,t}} E_t \sum_{j=0}^{\infty} \beta^j A_{t+j} (u_{c,t+j} c_{t+j} - u_{l,t+j}(1 - l_{t+j})) = \lambda_{3,t} \end{aligned} \tag{21}$$

$$u_{l,t} + \Delta A_t (u_{l,t} - u_{ll,t}(1 - l_t)) = \lambda_{3,t} \tag{22}$$

[9]The same assumption is made in Karantounias et al. (2007).

[10]As standard in the optimal fiscal policy literature, it is not easy to establish that the feasible set of the Ramsey problem is convex. To overcome this problem in our numerical calculations we check that the solution to the first-order necessary conditions of the Lagrangian is unique.

The first term on the left side of equation (21) represents the benefit for the government from increasing consumption by one unit. The second one measures the impact of the implementability constraint on the allocation, weighted by the distortion A_t represented by not fully rational expectations. The third and fourth terms reflect the fact that the government takes into account how agents update their expectations on the basis of the current consumption. The last term on the left represents the derivative of all future expected discounted primary surpluses with respect to current consumption. This is because from equation (16) each primary surplus (in terms of marginal utility) at $t+j, \forall j \geq 0$ is pre-multiplied by the product of past ratios of expected to actual marginal utility. In choosing optimal consumption today, the government is implicitly choosing the factor at which all future primary surpluses are discounted through its effect on A_t. The term on the right is the shadow value of output. A similar interpretation holds for the optimality condition with respect to leisure, equation (22).

Several comments are necessary. First, the optimal allocation is history-dependent through the presence of the state variable A_{t-1}: differently from Lucas and Stokey (1983), the allocation is not any more a time-invariant function of the current realization of the government shock only, but depends on what happened in the past. Second, in appendix A.3 I show that the optimality conditions in a complete markets and rational expectations framework are a special case of equation (21) and (22). Third, using the recursive formulation of A_t, the intertemporal budget constraint at t, which we repeat here for convenience

$$b_{t-1}A_t u_{c,t} = E_t \sum_{j=0}^{\infty} \beta^j A_{t+j}(u_{c,t+j}c_{t+j} - u_{l,t+j}(1-l_{t+j}))$$

and combining equations (21) and (22) the optimal allocation satisfies the following equation:

$$\begin{aligned} u_{c,t} + \Delta A_t(u_{cc,t}(c_t - b_{t-1}) + u_{c,t}) - \lambda_{1,t}\alpha_t u_{cc,t} I(g_t = g^H) - \lambda_{2,t}\alpha_t u_{cc,t} I(g_t = g^L) = \\ u_{l,t} + \Delta A_t(u_{l,t} - u_{ll,t}(1-l_t)) \end{aligned} \tag{23}$$

Equation (23) looks very similar to the first-order condition with respect to consumption in the incomplete markets model of Aiyagari et al. (2002). In both frameworks the excess burden of taxation is not constant, although for different reasons. In the absence of a full set of state-contingent bonds, as in Aiyagari et al. (2002), the excess burden of taxation is time-varying because of the incomplete insurance offered by the financial markets: since the interest payment on inherited debt is fixed across realizations of the current government shock, the government in each period has to adjust the stream of all future taxes to ensure solvency.[11] In a complete markets model with learning, what makes the excess burden of taxation time-varying is the

[11]In a complete market framework with rational expectations the excess burden of taxation is constant because the variable which adjusts to ensure solvency is the pay-off of the portfolio of contingent bonds.

cost of issuing state-contingent debt. Although market completeness implies that in each period the government can fully insure against expenditure shocks, the state contingent interest rates change as time goes by because agents' expectations change. Denote by T the time at which agents stop updating their beliefs because the forecast error is zero. Then $A_{T+j} = A_T \forall j \geq 1$, and the excess burden of taxation from T onwards becomes constant again.

Equation (21) expresses the actual marginal utility of consumption as a function of agents' beliefs about it. Figure 1 shows this relation for a log-log utility function and a given value of A_{t-1}.[12] The left panel displays the actual marginal utility of consumption contingent on the government expenditure shock being low (average with respect to the expected marginal utility of consumption contingent on the government expenditure shock being high), as a function of the previous period belief, γ_{t-1}^L. The right panel displays the same relation for the government expenditure shock being high. Figures 2 and 3 show the tax rate and the state-contingent bond policy functions which guarantee that the convergence between actual and expected marginal utility holds. The tax rate is a decreasing function of the previous period expected marginal utility; symmetrically, the state-contingent bond is an increasing function of it.

3 Some examples

In order to characterize the optimal fiscal policy in the framework we are studying, in what follows I consider some examples restricting the government expenditure shock to a specific form.

3.1 Constant government expenditure

Consider the case in which the government expenditure is known to be constant and the initial amount of bond holdings is zero. The Lagrangian associated to the government problem is

$$\begin{aligned}\mathcal{L} = \sum_{t=0}^{\infty} \beta^t [u(c_t, l_t) + \Delta A_t (u_{c,t} c_t - u_{l,t}(1 - l_t)) \\ + \lambda_{1,t}(\gamma_t - (1 - \alpha_t)\gamma_{t-1} - \alpha_t u_{c,t}) + \lambda_{2,t}(1 - l_t - c_t - g)] - \Delta A_0 u_{c,0} b_{-1}\end{aligned} \quad (24)$$

where the notation is the same as before and $x_t = \gamma_t$.

The optimality conditions $\forall t \geq 0$ are:

[12] The shape of the mapping is robust to different values of this variable.

$$u_{c,t} + \Delta A_t(u_{cc,t}c_t + u_{c,t}) - \lambda_{1,t}\alpha_t u_{cc,t} -$$
$$\Delta \frac{u_{cc,t}}{u_{c,t}} \sum_{j=0}^{\infty} \beta^j A_{t+j}(u_{c,t+j}c_{t+j} - u_{l,t+j}(1 - l_{t+j})) = u_{l,t} + \Delta A_t(u_{l,t} - u_{ll,t}(1 - l_t)) \tag{25}$$

$$\lambda_{1,t} - \beta(1 - \alpha_t)\lambda_{1,t+1} + \beta \Delta b_t A_t = 0$$

Equation (25) gives the mapping T between agents' beliefs about (marginal utility of) consumption and actual (marginal utility of) consumption. In the next proposition I characterize the properties of this mapping.

Proposition 2. *Assume the utility function*

$$u(c_t, l_t) = log c_t + l_t \tag{26}$$

and that the gain α_t is small.

Then, in the set $\gamma_{t-1} > 0$ the mapping T: $R_+ \rightarrow R_+$ has the following properties:

- *T is increasing and concave.*
- *T has one fixed point.*
- *The least squares learning converges to it.*

Proof. We relegate the proof to the appendix. □

Proposition 2 shows that the expected marginal utility converges to actual one, so that in the long-run agents' expectations are model-consistent and the forecast error is zero. The next proposition characterizes the value towards which expectations converge.

Proposition 3. *Given an initial value for the government bond holding b_{-1} the allocation under learning does not converge to the allocation under rational expectations implied by the same initial bond holding. However, for any initial belief held by agents, there exists a b_{-1} such that*

$$lim_{t\rightarrow\infty} c_t^L(\gamma_{-1}) = c_t^{RE}(b_{-1}) \tag{27}$$

Proof. We relegate the proof to the appendix. □

Proposition 3 shows for any initial belief about the marginal utility of consumption, there is always an initial level of government wealth such that the allocation under learning converges to the one under rational expectations starting with that initial government wealth. Figure 4 shows this relation assuming that in equation (24) $b_{-1} = 0$. Given the parameters values used, the solution of the Ramsey problem under rational expectations implies that the marginal utility of consumption is constant and equal to 2.5. For all values of initial belief higher (lower) than this reference value, the learning allocation coincides with the solution of a Ramsey problem in an economy populated by rational agents and endowed with a positive (negative) initial government debt.

3.1.1 Policy implications

The example of constant government consumption highlights the impact of expectations on the optimal fiscal plan. Under rational expectations, the only distortion is the one associated to taxes. In order to smooth this distortion over time, taxes are set to balance the government budget every period. In this way agents can enjoy a perfectly constant allocation. By contrast under learning, there are two distortions in the economy, one associated with taxes and the other one associated with agents' expectations. Therefore, although the government could follow a balanced-budget rule, it decides not to do it because in this way it would not minimize the *overall* distortions. To influence out-of-equilibria expectations the government animates initially pessimistic agents setting a low tax rate at the beginning and financing the public consumption with debt. As time goes by, the tax rate has to increase in order to ensure government solvency.[13]. This stabilization policy is resistant to a selection of robustness checks. For example, it holds if 1) we suppose that agents use lagged value of marginal utility of consumption to update their current beliefs, 2) the government has access to consumption taxes instead of labor ones.

Figures 5, 6 and 7 offer a graphical interpretation of the result. The solid lines show the optimal fiscal plan under rational expectations. Whereas the dashed lines show the optimal fiscal plan when agents adopt a constant gain algorithm to update their belief while supposing that in the initial period the expected consumption is lower than the actual consumption prevailing at $t = 0$.[14]

3.2 Perfectly Foreseen War

Consider the case in which expenditure is constant in all periods apart from T, when $g_T > g_t$. Both the government and households know the entire path of the expenditure, so that the shock in T is perfectly anticipated. Under rational expectations the government runs a positive

[13]The analysis is symmetric for the case of initially optimistic agents

[14]Since for optimistic agents the evolution of the system is symmetric, we do not report it.

primary surplus from period 0 to $T-1$, using it to buy bonds. At T the government finances the high public consumption level by selling the accumulated assets and possibly by levying a tax rate on labor income. From period $T+1$ onwards the tax rate is just sufficient to cover the expenditure and to service the interest on the bonds issued at T. By contrast, in an economy populated by pessimistic agents, the government can accumulate less assets because it has to stimulate the economy to manipulate expectations. The big shock at T is financed by increasing debt much more than under rational expectations. Figures 8-10 illustrate the optimal plan under rational expectations (the solid lines) and under learning (the dashed lines), assuming $T=10$, $g_t = 0.1$ and $g_T = 0.2$.

3.2.1 Policy implications

The example of a perfectly anticipated war is useful for two reasons. First, it clarifies how the tax smoothing result is altered by the presence of boundedly rational agents. Under rational expectations the government spreads over time the cost of financing the war in T thorough distortionary taxes. As a result, the tax rate is perfectly constant in all periods before and after the war: taxes are smooth in the sense that they have a smaller variance than a balanced-budget rule would imply. By contrast, when agents are learning, they do not trust the promises made by the fiscal authority in terms of future consumption. The government uses taxes and debt to correct agents' distorted expectations, in a way that the tax rate is more volatile than under rational expectations.

The example is relevant also because it reconciles the complete markets framework with the empirical evidence that during peacetime periods countries run a primary deficit. The Lucas and Stokey (1983) model is unable to fit this evidence, as the government runs a primary surplus to accumulate assets before the war.

3.3 Perfectly foresen, cyclical wars

Suppose that

$$\begin{aligned} g_{t'} &= g_H \quad \forall t' = j \times H \leq \overline{T} \\ g_t &= g_L \quad otherwise \end{aligned}$$

with $j = 1, 2, ..., \frac{\overline{T}}{H}$. H is the length of time over two subsequent bad shocks and $\overline{T}$ is the last period in which a bad shock can occur. The rational expectations policy recipe is the same as before: the tax rate is constant in all periods when $g_t = g_L$ and increases very little when the bad shock hits the economy, due to the assets the government accumulates during the good shock periods. Under learning with pessimistic agents, before the first realization of the bad shock the tax rate is lower than under RE and increases between any two subsequent bad shocks, generating resources devoted to reducing debt, which increases whenever the bad shock occurs. After the last bad realization of the shock, the tax rate falls and then gradually increases over time to ensure intertemporal solvency.

3.4 War of unknown duration

Suppose that the shock can take two realizations, g_L and g_H, with the following transition probabilities matrix

$$P = \begin{pmatrix} 1 & 0 \\ \pi_{H,L} & \pi_{H,H} \end{pmatrix}$$

where $\pi_{i,j}$ is the probability that tomorrow the shock is in state j, being today in state i, with $g_t = g_H$ at $t = 0$. This example corresponds to an absorbing Markov chain, where the low realization of the shock is the absorbing state and the high one is the transient state. Under rational expectations, the government finances the bad shocks partly through taxes and partly by issuing debt. Numerical results, not reported here, confirm the role of fiscal policy as stabiliser of expectations: the accumulation of public debt is higher and longer under learning than under rational expectations, the difference being due to the opportunity of inducing the agents to revise their expectations downwards.

3.5 Serially correlated shock

Suppose that the shock can take two realizations, g_L and g_H, with the following transition probabilities matrix

$$P = \begin{pmatrix} \pi_{L,L} & \pi_{L,H} \\ \pi_{H,L} & \pi_{H,H} \end{pmatrix}$$

We set $g_L = 0.05$,$g_H = 0.1$ and $\pi_{H,H} = \pi_{L,L} = 0.8$.[15] As in the previous examples, we assume a discount factor equal to 0.95 and a gain parameter equal to 0.02.[16]

Table 1 summarizes some statistics for the allocation and the fiscal variables under rational expectations. Table 2 summarizes the same statistics under learning after convergence of beliefs for initially pessimistic and optimistic agents. Reported values are average across 1000

[15] For the case of i.i.d shock case the results are very similar to those with a serially correlated shock, and therefore they are not reported.

[16] The choice of the updating parameter is not easy because it requires a trade-off between filtering noises and tracking structural changes. Milani (2007) estimates a New-Keynesian model and finds that the best fitting specification has a gain coefficient in a range between 0.015 and 0.03. Orphanides and Williams (2004) find that a value for k in the range $0.01 - 0.04$ fits the expectations data from the Survey of Professional Forecasters better than using higher or lower values. Evans et al. (2007) also use the same value for k.

simulations. Comparing the two tables we can observe that in the long run initially pessimistic (optimistic) agents consume less (more) than if they had been rational. This result is in line with the examples in sections 3.1 and 3.2. Since at the beginning consumers were pessimistic, the accumulation of debt necessary to induce them to revise upwards their expectations about consumption requires higher taxes in the long run than under rational expectations. Because of this, consumption is lower and leisure is higher (than under rational expectations). The average primary surplus is higher as well. In this sense we can say that beliefs are self-fulfilling in the long run: the lower is the initial expected consumption, the lower the actual consumption after convergence. Exactly the opposite is true with initially optimistic agents.

Table 3 shows the same statistics for the system during the first 30 periods of transition, under rational expectations and initially pessimistic agents. Although all the endogenous variables are more volatile before convergence than after convergence, the market value of government debt and the labor tax rate are the most volatile ones. For example, the tax rate volatility before convergence is double that after convergence. This is due to the fact that the government implements an expectation-dependent fiscal plan. When beliefs are distorted, fiscal variables react to correct this distortion. As time passes and agents' expectations become model-consistent, the government stops using fiscal variables to influence distorted beliefs.

4 Responsible and irresponsible governments

The analysis in section 3 has characterized the optimal fiscal plan that a benevolent government should implement when it realizes that agents are boundedly rational. Stabilizing out-of-equilibria expectations has a cost and a benefit. The cost is represented by the fact that taxes are less smooth, but the benefit is that under the expectation-dependent fiscal plan agents learn the tax rate policy rule much faster than under the rational expectations optimal fiscal plan. Figure 11 illustrates this point graphically. The solid line shows the next-period marginal utility of consumption forecast error made by initially pessimistic agents when the government follows the rational expectations recipe. The dashed line shows the same series when the government instead implements the optimal policy under learning. The lower (than RE) tax rate set at the beginning following the optimal fiscal policy plan induces agents to correct their pessimism much faster than if the fiscal policy suggested by Lucas and Stokey (1983), which is the optimal one under rational expectations, were followed.

The way in which the government should use fiscal variables to manipulate agents' distorted expectations in some sense resembles a standard Keynesian-inspired stabilization policy. However, it is important to stress that the government should not stimulate economic activity indiscriminately. Actually it is very important to implement the right policy at the right moment. In what follows I show that an expansionary fiscal policy, implemented when agents' expectations require a restrictive one, generates a sub-optimal volatility in the system.

Suppose for simplicity that the public consumption shock is constant and that the government wants to animate the economy *even if* agents are optimistic. To this aim, it implements the following tax-rate rule

$$\tau_t = \begin{cases} \tau_t^{pess}, & \forall t \leq T \\ \xi_t \tau_t^{pess} + (1-\xi_t)\tau_t^{bb}. & \forall t > T \end{cases} \tag{28}$$

According to equation (28) the fiscal authority stimulates the economy till period T setting the tax rate at the (low) optimal level when agents are pessimistic, τ_t^{pess}, and that from T onwards sets the tax rate as a weighted average between this value and the one which raises enough revenues to pay-back both the interests on the inherited debt and the current government expenditure shock, τ_t^{bb}. The weight is given by $\xi_t = k^{t-T}$, with $0 < k < 1$. In order to ensure that the transversality condition is not violated, it is necessary to impose the restriction that the weight ξ_t goes to zero quickly. Otherwise the fiscal revenues raised through distortionary taxes would not be enough to finance the interests on the debt that the government has accumulated. The parameter k is set small enough to rule-out any Ponzi-scheme.

For simplicity I assume the utility function:

$$log(c_t) + log(l_t) \tag{29}$$

and set $\beta = 0.95$, $g = 0.1$, $T = 24$ and $k = 0.7$.

The dashed lines in figure 14 show the optimal tax rate and bond holdings when agents are optimistic and the government implements the fiscal policy taking into account that they are optimistic, while the solid lines show the same variables when agents are optimistic and the government implements the rule given by equation 28. The dashed lines in figure 15 show the consumption, leisure and forecast error when agents are optimistic and the government implements the optimal fiscal plan conditioning on this, while the solid lines show the same variables under the wrong stimulus.

Looking at the figures it can be noticed that the allocation is more volatile under the wrong stimulus, and consequently households' welfare is lower. The welfare losses in terms of consumption-equivalence units of animating the economy when the optimum requires depressing it is equal to 0.2 percent. Increasing T and/or k increases the welfare losses. Using the same parameters as before but with $T = 25$, the welfare losses become equal to 0.25; similarly with $k = 0.78$ the welfare losses is equal to 0.27. The intuition is that the higher is T and/or the higher is k, the more the fiscal policy is expansionary instead of being restrictive, as it should be since agents are optimistic.

The sub-optimally higher volatility generated by an unjustifiable government's desire to stimulate the economy translates into a longer time for agents to learn. While under the optimal fiscal plan the forecast error is zero after 50 periods, under the wrong stimulus it fluctuates much more and it is still not zero after 300 periods. In conclusion, setting a low tax rate to stimulate the economy when the opposite is required inefficiently induces instability into the system. While a government can accumulate debt even when it is not necessary, a responsible government will only accumulate it when necessary.

5 Endogenous Government Spending

Up to now I have considered public consumption as a completely exogenous shock. This assumption seems quite restrictive, as governments can decide how much to spend. In this section I consider the same model as in subsection 3.1 but we allow the government to choose the amount of public spending.

I assume the utility function

$$u(c_t, l_t, g_t) = log(c_t) + log(l_t) + \alpha log(g_t)$$

with $\alpha < 1$. The Lagrangian associated to the government problem is

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t [u(c_t, l_t, g_t) + \Delta A_t (u_{c,t} c_t - u_{l,t}(1 - l_t)) \\ + \lambda_{1,t}(\gamma_t - (1 - \alpha_t)\gamma_{t-1} - \alpha_t u_{c,t}) + \lambda_{2,t}(1 - l_t - c_t - g_t)] - \Delta A_0 u_{c,0} b_{-1}$$

The dashed lines in figures 12 and 13 show the rational expectations equilibrium, whereas the solid lines show the learning optimal allocation when agents are pessimistic. In line with the previous results, at the beginning the government chooses an expansionary fiscal policy, setting higher public spending than under rational expectations and financing it mainly through debt.

6 Testing Complete Versus Incomplete Markets

In section 2 I have shown that the first order condition with respect to consumption in a complete markets model with learning looks very similar to the one in an incomplete markets model with rational expectations. In both cases the excess burden of taxation changes over time, although for different reasons.

Assessing whether markets are complete or incomplete is not an obvious issue, since there are theoretical justifications in both directions: while transaction costs and limited commitment push in favour of market incompleteness, the possibility of replicating the complete markets equilibrium through a portfolio of bonds with different maturities favours market completeness. The tests proposed in the literature to discriminate between complete and incomplete markets (see inter alia Scott (2007), Marcet and Scott (2008), and Faraglia et al. (2006)) are based on two discriminating features between the two regimes:

1. Under complete markets fiscal variables (tax rate and market value of debt) inherit the serial properties of the underlying shocks hitting the economy, while under incomplete markets they have a unit-root component.

2. Under complete markets the market value of debt and the primary deficit co-move negatively, while under incomplete markets they co-move positively.

Tests based on the first feature are persistence tests, and those based on the second are impact tests. The aim of this section is to show that the tests belonging to the first category are not able to discriminate between an incomplete markets model and a complete markets model when agents are boundedly rational. In particular I argue that these tests would be prone to accept the wrong hypothesis that markets are incomplete if, in fact, agents learn and markets are complete. The reason is simply that learning creates persistence in the system.[17]

We replicate persistence tests proposed in Scott (1999, 2007) and Faraglia et al. (2007) to check market completeness. The first test is based on the presence of unit root in the labor tax rate. Assume that the government expenditure shock is stationary; under complete markets the labor tax rate is stationary, while under incomplete markets it contains a unit root. We simulate the model described in Section 2 and we apply the Augmented Dickey Fuller test to the tax rate using the first 50 periods of data. Out of 1000 simulations, the probability of accepting the unit-root test, and therefore of concluding (erroneously) that markets are incomplete, is equal to 0.999.

The second test is to estimate whether the excess burden of taxation has a unit root: under rational expectations and complete markets the excess burden of taxation is constant, while under rational expectations and incomplete markets it has a unit root, as shown in Aiyagari et al. (2002). Using the same sample period as before, the probability of accepting the unit-root test, and therefore of concluding (erroneously) that markets are incomplete, is equal to 0.922.

The third test in based on the fact that under complete markets the market value of debt and the primary deficit have the same persistence, while under incomplete markets the first is more

[17]One way to compare persistence under learning and under rational expectations is to look at autoregressions of tax rates in the two frameworks when the shock is *i.i.d.*. Table 4 shows that while under rational expectations the coefficient on the lagged tax rate is close to zero and not statistically significant, under learning it is high and statistically significant.

persistent than the second. Once again this result does not hold in a boundedly rationality framework. The top panel in Fig. 16 displays the persistence of the debt/GDP ratio, the primary surplus/GDP ratio and the government expenditure shock when agents have rational expectations and markets are complete. The bottom panel displays the same statistics when agents are boundedly rational and markets are complete.[18]

Results in figure 16 would induce the researcher to accept, once again erroneously, the hypothesis of market incompleteness.

However a model with complete markets and learning is *not* observationally equivalent to a model with incomplete markets and rational expectations. Actually impact tests, already used in the literature, are able to capture the differences between the two frameworks. Consider for example the co-movement between primary deficit and government debt. Independently of the way in which agents form their expectations, under complete markets this co-movement is negative, while under incomplete markets it is positive, and so is in data.

To conclude, in line with Marcet and Scott (2008) looking at the behavior of debt is a much more reliable way to test the bond market structure than looking at the behavior of tax rate.

7 Debt Sustainability and Debt Limits

The literature has recently emphasized the opportunity of imposing limits on the amount of debt a government can accumulate.[19] In a context of boundedly rational agents the long run market value of debt depends on the initial beliefs held by agents: the higher the initial pessimism in the economy, the higher is the long-run level of debt. Since this debt accumulation is "good", in the sense that it allows for convergence between actual and expected marginal utility, there is not necessarily a correspondence between keeping the debt/GDP ratio low and optimal fiscal policy considerations. Moreover, debt limits may fail to discriminate between "good" and "bad" governments. Consider two countries, hit by the same realization of the government expenditure shock which differ only as to the vector of initial beliefs. Figure 17 shows the probability that the debt limit (set equal to 60 per cent of steady state GDP) is binding for the two countries conditioning on the fact that each of the two governments implements the optimal fiscal plan taking as given the initial degree of pessimism. Since the first country is populated by less pessimistic agents than the second, the long-run debt level is lower in the first than in the

[18] To measure the persistence of a variable, say y, we use the k-variance ratio, defined as

$$P_y^k = \frac{Var(y_t - y_{t-k})}{kVar(y_t - y_{t-1})}$$

[19] In Chari and Kehoe (2004), debt constraints are beneficial if the monetary authority cannot commit to solve the time inconsistency problem of deflating the nominal debt issued by the fiscal authorities of the member states.

second. But this does not mean that the government in the first country is more responsible than the one in the second just because it accumulated less debt. The only reason for the difference in the long-run debt level is that in the first country initial beliefs were less distorted than in the second, and the government had to intervene less to correct them.

The main advantage of debt constraints is that they are helpful in ensuring sustainability of fiscal policy.[20] Assessments of debt sustainability performed by international institutions are usually based on medium-term simulations (generally 5-10 years) of the debt/GDP ratio. A declining trend in debt/GDP ratio is interpreted as a signal that the government follows a sustainable fiscal policy, whereas an increasing one raises doubts about intertemporal solvency. In a model with boundedly rational agents assessing sustainability is particularly cumbersome, exactly because at the beginning government debt displays a trend.

Suppose that an agency wants to test for the presence of unit root in the debt/GDP ratio, in which case the fiscal policy plan is declared unsustainable. We show that actually this test can perform very poorly if the government follows the optimal fiscal policy plan when agents are learning.

Suppose that agents in the economy are initially pessimistic and that the government implements the fiscal policy accordingly. The agency is asked to evaluate the government solvency and to do that it applies an Augmented Dickey Fuller test on the market value of debt using the first 50 periods of observations using this regression

$$\frac{debt_t}{GDP_t} = \alpha + \beta_T^{OLS} \frac{debt_{t-1}}{GDP_{t-1}} + \gamma_T^{OLS} \left(\frac{debt_{t-1}}{GDP_{t-1}} - \frac{debt_{t-2}}{GDP_{t-2}}\right) + \epsilon_t \tag{30}$$

Over 1000 simulations, the probability that the agency would declare the fiscal plan to be unsustainable when instead it is sustainable by construction is equal to 0.697.[21] The reason for this is that debt is used to manipulate agents' expectations: since these are persistent, they impart persistence to debt as well. In other words, bounded rationality increases the lack of power of unit root tests. Suppose now that the agency applies the Augmented Dickey Fuller test to the following equation

$$\begin{aligned} \left(\frac{debt}{GDP}\right)_t = \alpha^{OLS} + \beta_T^{OLS} \left(\frac{debt}{GDP}\right)_{t-1} + \gamma_T^{OLS} \left(\frac{s}{GDP}\right)_t + \delta_T^{OLS} \Delta\left(\frac{debt}{GDP}\right)_{t-1} + \\ + \nu_T^{OLS} \Delta\left(\frac{s}{GDP}\right)_{t-1} + \mu_T^{OLS} \Delta\left(\frac{s}{GDP}\right)_{t-2} + \epsilon_t \end{aligned} \tag{31}$$

where we added two lagged difference terms of the primary surplus/GDP ratio to obtain white noise residuals. In this case the probability of getting the wrong answer of debt unsustainability

[20] A set of tests has been proposed by the literature to check sustainability, among others by Hamilton and Flavin (1986), Trehan and Walsh (1991) and Bohn (1998).

[21] The result is robust to a number of alternative specifications: adding lagged difference terms of the dependent variable the probability would be equal to 0.77; including a time trend would lower the probability to 0.652.

is be equal to 0.163, much lower than before. We conclude that one way to disentangle between persistence and sustainability of debt is to consider the evolution of primary surpluses.

8 Conclusions

In this paper I characterize the optimal fiscal policy when agents are boundedly rational and taxes are distortionary. There are two main results. The first one is that the policymaker should manipulate agents' beliefs by setting low (high) taxes in a context of pessimism (optimism). This conclusion is in line with some suggestions to handle the recent distress. The second one is that the complete market solution under learning is history-dependent. This fact makes assessing market completeness more challenging, since unit-root test in the tax rate can mix evidence of bounded rationality with evidence of market incompleteness. In line with Marcet and Scott (2008) I find that looking at the behavior of debt is a much more reliable way to test the bond market structure. Also gauging debt sustainability is more complicated because of the persistence induced by agents' expectations.

Several important issues are still open question. First, for simplicity I restrict the analysis to one-period state-contingent bonds. In reality governments issue uncontingent debt at different maturities. Including these two features in the analysis is important both from a normative and a positive point of view. Second, I assume that while agents do not know how aggregate variables are determined, the government has full information about the structure of the economy. Other authors (*inter-alia*, Primiceri (2005) and Cogley and Sargent (2005)) followed the opposite approach. A natural framework to analyze is one in which neither the households know the government policy rules nor the government knows the households' response to these rules. Third, as in several papers on optimal taxation I abstract from monetary issues. On the other hand, the literature studying the impact of learning on the monetary policy design abstracts from fiscal policy considerations, such as distortionary taxes. A natural step would be to unify these two strands and to understand how the interaction of fiscal and monetary policy can help agents to form their expectations. We leave these issues to a future exercise.

References

Aiyagari, R., Marcet, A., Sargent, T. and Seppala, J.: 2002, Optimal taxation without state-contingent debt, *Journal of Political Economy* **110**, 1220–1254.

Barro, R.: 1979, On the determination of public debt, *Journal of Political Economy* **87(5)**, 940–71.

Barro, R.: 1989, The neoclassical approach to fiscal policy, *Modern Business Cycle Theory Cambridge, M.A: Harvard University Press* pp. 1–15.

Barro, R.: 1995, Optimal debt management, *NBER WP 5327* .

Barro, R.: 1997, Optimal management of indexed and nominal debt, *NBER WP 6197* .

Bohn, H.: 1990, Tax smoothing with financial instruments, *American Economic Review* **80**, 1217–1230.

Bohn, H.: 1998, The behavior of u.s. public debt and deficit, *The Quarterly Journal of Economics* **113**, 949–963.

Carceles-Poveda, E. and Giannitsarou, C.: 2007, Adaptive learning in practice, *Journal of Economic Dynamics and Control* .

Chari, V., Christiano, L. and Kehoe, P.: 1994, Optimal fiscal policy in a business cycle model, *The Journal of Political Economy* .

Chari, V. and Kehoe, P.: 1999, Optimal fiscal and monetary policy, *Handbook of Macroeconomics* .

Chari, V. and Kehoe, P. J.: 2004, On the deiderability of fiscal constraints in a monetary union, *NBER Working Paper No. 10232* .

Cogley, T. and Sargent, T.: 2005, The conquest of u.s. inflation: Learning and robustness to model uncertainty, *ECB W.P. 478* .

Evans, G. and Honkapohja, S.: 2003, Expectations and the stability problem for optimal monetary policies, *Review of Economic Sudies* **70**, 807–824.

Evans, G. and Honkapohja, S.: 2006, Monetary policy, expectations and commitment, *Scandinavian Journal of Economics* **108**, 15–38.

Evans, G., Honkapohja, S. and K., M.: 2003, Notes on agents' behavioral rules under adaptive learning and recent studies on monetary policy.

Evans, G., Honkapohja, S. and K., M.: 2007, Anticipated fiscal policy and adaptive learning.

Faraglia, E., Marcet, A. and Scott, A.: 2006, In search of a theory of debt management.

Faraglia, E., Marcet, A. and Scott, A.: 2007, Fiscal insurance and debt management in oecd economies.

Hamilton, J. and Flavin, M.: 1986, On the limitations of government borrowing: A framework for empirical testing, *American Economic Review* **76**, 808–819.

Karantounias, A., Hansen, L. and Sargent, T.: 2007, Ramsey taxation and the fear of misspecification.

Kydland, F. and Prescott, E.: 1980, Dynamic optimal taxation, rational expectations and optimal control, *Journal of Economic Dynamics and Control* .

Lucas, R. and Stokey, N.: 1983, Optimal fiscal and monetary policy in an economy without capital, *Journal of Monetary Economics* .

Marcet, A. and Nicolini, J. P.: 2003, Recurrent hyperinflations and learning, *The American Economic Review* **93, No. 5**, 1476–1498.

Marcet, A. and Scott, A.: 2008, Debt and deficit fluctuations and the structure of bonds markets, *Journal of Economic Theory, forthcoming* .

Milani, F.: 2007, Expectations, learning and macroeconomic persistence, *Journal of Monetary Economics* .

Molnar, C. and Santoro, S.: n.d., Optimal monetary policy when agents are learning.

Orphanides, A. and Williams, G.: 2006, Monetary policy with imperfect knowledge, *Journal of the European Economic Association* **4**, 366–375.

Orphanides, A. and Williams, J. C.: 2004, Robust monetary policy with imperfect knowledge.

Preston, B.: 2005a, Adaptive learning in infinite horizon decision problems, p. Manuscript.

Preston, B.: 2005b, Learning about monetary policy rules when long-horizon expectations matter.

Preston, B. Eusepi, S.: 2006, Expectations driven fluctuations and stabilization policy.

Preston, B. and Eusepi, S.: 2007a, Central bank communication and expectations stabilization.

Preston, B. and Eusepi, S.: 2007b, Stabilizing expectations under monetary and fiscal coordination.

Primiceri, G.: 2005, Why inflation rose and fell: Policymakers's beliefs and us postwar stabilization policy, *NBER WP, No. 11147* .

Sargent, T.: 2005, Ambiguity in american monetary and fiscal policy.

Scott, A.: 1999, Does tax smoothing imply smooth taxes, *CEPR Discussion Paper No. 2172* .

Scott, A.: 2007, Optimal taxation and oecd labor taxes, *Journal of Monetary Economics* **54**, 925–944.

Svec, J.: 2008, Optimal fiscal policy with robust control.

Trehan, B. and Walsh, C. E.: 1991, Testing intertemporal budget constraints: Theory and application to u.s. federal budget and current account deficit, *Journal of Money, Credit and Banking* **23**, 206–223.

Zhu, X.: 1992, Optimal fiscal policy in a stochastic growth model, *JET* **58**, 250–289.

A Appendix

A.1 Proof of Proposition 1

First I show that constraints (3), (10), (11) and (12) imply (16).

Consider the period-by-period budget constraint after substituting for the household optimality conditions:

$$b_{t-1}(g_t) = \frac{u_{c,t}(g_t)s_t(g_t)}{u_{c,t}(g_t)} + \beta\frac{\gamma_t^H b_t(H)}{u_{c,t}(g_t)}\pi(H) + \beta\frac{\gamma_t^L b_t(L)}{u_{c,t}(g_t)}\pi(L) \tag{32}$$

where $s_t \equiv c_t - \frac{u_{l,t}}{u_{c,t}}(1-l_t)$, $b_t(i)$ for $i = H, L$ is the amount of state-contingent bond holdings, and γ_t^i for $i = H, L$ are the (state contingent) marginal utilities that agents expect in the next period and $\pi(i) = \pi(g_{t+1} = g^i|g_t)$ for $i = H, L$. Since the fiscal authority has full information about the economy, in t it will issue state-contingent bonds such that the budget constraint in the next period is satisfied for any realization of the government shock. Forwarding equation (32) one period

$$b_t(H) = \frac{u_{c,t+1}(H)s_{t+1}(H)}{u_{c,t+1}(H)} + \beta\frac{\gamma_{t+1}^H b_{t+1}(H)}{u_{c,t+1}(H)}\pi(H) + \beta\frac{\gamma_{t+1}^L b_{t+1}(L)}{u_{c,t+1}(H)}\pi(L) \tag{33}$$

$$b_t(L) = \frac{u_{c,t+1}(L)s_{t+1}(L)}{u_{c,t+1}(L)} + \beta\frac{\gamma_{t+1}^H b_{t+1}(H)}{u_{c,t+1}(L)}\pi(H) + \beta\frac{\gamma_{t+1}^L b_{t+1}(L)}{u_{c,t+1}(L)}\pi(L) \tag{34}$$

Substituting equations (33) and (34) into equation (32), and multiplying both sides by $x_{t-1} \equiv [\gamma_{t-1}^H I(g_t = g^H) + \gamma_{t-1}^L I(g_t = g^L)]$

$$\begin{aligned} b_{t-1}(g_t)x_{t-1} = u_{c,t}s_t\frac{x_{t-1}}{u_{c,t}} + \frac{x_{t-1}}{u_{c,t}}\beta\{&\frac{\gamma_t^H}{u_{c,t+1}(H)}[u_{c,t+1}(H)s_{t+1}(H)+ \\ &\beta\gamma_{t+1}^H b_{t+1}(H)\pi(H) + \beta\gamma_{t+1}^L b_{t+1}(L)\pi(L)]\pi(g_{t+1} = g^H|g_t) + \frac{\gamma_t^L}{u_{c,t+1}(L)}[u_{c,t+1}(L)s_{t+1}(L)+ \\ &\beta\gamma_{t+1}^H b_{t+1}(H)\pi(H) + \beta\gamma_{t+1}^L b_{t+1}(L)\pi(L)]\pi(g_{t+1} = g^L|g_t)\} \end{aligned} \tag{35}$$

Define $W_t = \frac{x_{t-1}}{u_{c,t}}$. Keeping substituting forward we get

$$b_{t-1}(g_t)x_{t-1} = E_t\sum_{j=0}^{\infty}\beta^j \tilde{A}_{t+j}u_{c,t+j}s_{t+j} \tag{36}$$

where

$$\tilde{A}_{t+j} \equiv \prod_{k=t}^{t+j} W_k \tag{37}$$

Multiplying each side of equation (36) by $A_{t-1} \equiv \prod_{k=0}^{t-1} W_k$ we get

$$b_{t-1}(g_t)x_{t-1}H_{t-1} = E_t \sum_{j=0}^{\infty} \beta^j A_{t+j} u_{c,t+j} s_{t+j} \tag{38}$$

where $A_{t+j} = H_{t-1} \prod_{k=t}^{t+j} W_k$.

Notice that the A_{t+j} has a recursive formulation given by:

$$\begin{aligned} A_{t+j} &= \prod_{k=0}^{t-1} W_k \times \prod_{k=t}^{t+j} W_k = \prod_{k=0}^{t+j} W_k = \prod_{k=0}^{t+j-1} W_k \times \frac{x_{t+j-1}}{u_{c,t+j}} = \\ &= A_{t+j-1} \frac{x_{t+j-1}}{u_{c,t+j}} \end{aligned} \tag{39}$$

To prove the reverse implication, take any feasible allocation $\{c_{t+j}, l_{t+j}\}_{j=0}^{\infty}$ that satisfies equation (16). Then it is always possible to back out the state-contingent bond holding such that the period-by-period budget constraint is satisfied.

From equation (38), define

$$b_{t-1}(g_t) = E_t \sum_{j=0}^{\infty} \beta^j A_{t+j} u_{c,t+j} s_{t+j} \frac{1}{x_{t-1} A_{t-1}} \tag{40}$$

It follows that

$$b_t(g_{t+1}) = E_{t+1} \sum_{j=0}^{\infty} \beta^j A_{t+1+j} u_{c,t+1+j} s_{t+1+j} \frac{1}{x_t A_t} \tag{41}$$

$$\begin{aligned} b_{t-1}(g_t) &= \frac{A_t u_{c,t} s_t}{x_{t-1} A_{t-1}} + E_t \sum_{j=1}^{\infty} \beta^j A_{t+j} u_{c,t+j} s_{t+j} \frac{1}{x_{t-1} A_{t-1}} = \\ &= \frac{A_t u_{c,t} s_t}{x_{t-1} A_{t-1}} + \beta E_t \sum_{j=0}^{\infty} \beta^j A_{t+1+j} u_{c,t+1+j} s_{t+1+j} \frac{1}{x_{t-1} A_{t-1}} = \\ &= \frac{A_t u_{c,t} s_t}{x_{t-1} A_{t-1}} + \frac{\beta}{x_{t-1} A_{t-1}} E_t \{x_t A_t [\frac{\sum_{j=0}^{\infty} \beta^j A_{t+1+j} u_{c,t+1+j} s_{t+1+j}}{x_t A_t}]\} = \\ &= \frac{A_t u_{c,t} s_t}{x_{t-1} A_{t-1}} + \frac{\beta}{x_{t-1} A_{t-1}} E_t \{x_t A_t [E_{t+1} \frac{\sum_{j=0}^{\infty} \beta^j A_{t+1+j} u_{c,t+1+j} s_{t+1+j}}{x_t A_t}]\} = \\ &= \frac{A_t u_{c,t} s_t}{x_{t-1} A_{t-1}} + \frac{\beta}{x_{t-1} A_{t-1}} E_t \{x_t A_t B_t\} = \end{aligned} \tag{42}$$

Using equation (**??**) when $j = 0$ we get

$$\begin{aligned} b_{t-1}(g_t) &= s_t + \frac{\beta}{u_{c,t}} E_t([\gamma_t^H I(g_{t+1} = g^H) + \gamma_t^L I(g_{t+1} = g^L)] b_t(g_{t+1})) = \\ &= s_t + \frac{\beta}{u_{c,t}} [\gamma_t^H b_t(g_{t+1} = g^H) \pi(g_{t+1} = g^H | g_t) + \gamma_t^L b_t(g_{t+1} = g^L) \pi(g_{t+1} = g^L | g_t)] \end{aligned} \tag{43}$$

A.2

Attach the multipliers Δ, $\beta^t\pi_t(g^t)\lambda_{1,t}(g^t)$, $\beta^t\pi_t(g^t)\lambda_{2,t}(g^t)$, $\beta^t\pi_t(g^t)\lambda_{3,t}(g^t)$ and $\beta^t\pi_t(g^t)\lambda_{4,t}(g^t)$ to constraints (17), (19) for $i = H, L$, (20) and to (18).

The Lagrangian is

$$\begin{aligned}\mathcal{L} =& E_0\sum_{t=0}^{\infty}\beta^t\{u(c_t,l_t) + \Delta(A_t(u_{c,t}c_t - u_{l,t}(1-l_t)))\\ &+ \lambda_{1,t}((\gamma_t^H - \gamma_{t-1}^H)I(g_t = g_L) + (\gamma_t^H - (1-\alpha_t)\gamma_{t-1}^H - \alpha_t u_{c,t})I(g_t = g_H))\\ &+ \lambda_{2,t}((\gamma_t^L - \gamma_{t-1}^L)I(g_t = g_H) + (\gamma_t^L - (1-\alpha_t)\gamma_{t-1}^L - \alpha_t u_{c,t})I(g_t = g_L))\\ &+ \lambda_{3,t}(1 - l_t - c_t - g_t)\} + \lambda_{4,t}(A_t - A_{t-1}\frac{\gamma_{t-1}^H I(g_t = g^H) + \gamma_{t-1}^L I(g_t = g^L)}{u_{c,t}}) - \Delta A_0 u_{c,0}b_{-1}\end{aligned}$$

Assuming $b_{-1} = 0$, the first-order necessary conditions $\forall t \geq 0$ are:

- c_t:

$$\begin{aligned}&u_{c,t} + \Delta A_t(u_{cc,t}c_t + u_{c,t}) - \lambda_{1,t}\alpha_t u_{cc,t}I(g_t = g^H)\\ &- \lambda_{2,t}\alpha_t u_{cc,t}I(g_t = g^L) + u_{cc,t}\lambda_{4,t}A_{t-1}\frac{\gamma_{t-1}^H I(g_t = g^H) + \gamma_{t-1}^L I(g_t = g^L)}{u_{c,t}^2} = \lambda_{3,t}\end{aligned} \tag{44}$$

- l_t:

$$u_{l,t} + \Delta A_t(u_{l,t} - u_{ll,t}(1-l_t)) = \lambda_{3,t} \tag{45}$$

- γ_t^H:

$$\begin{aligned}&\lambda_{1,t} - \beta E_t\{\lambda_{1,t+1}I(g_{t+1} = g^L) + (1-\alpha_{t+1})\lambda_{1,t+1}I(g_{t+1} = g^H) +\\ &+ \frac{\lambda_{4,t+1}A_t}{u_{c,t+1}}I(g_{t+1} = g^H)\} = 0\end{aligned} \tag{46}$$

- γ_t^L:

$$\begin{aligned}&\lambda_{2,t} - \beta E_t\{\lambda_{2,t+1}I(g_{t+1} = g^H) + (1-\alpha_{t+1})\lambda_{2,t+1}I(g_{t+1} = g^L) +\\ &+ \frac{\lambda_{4,t+1}A_t}{u_{c,t+1}}I(g_{t+1} = g^L)\} = 0\end{aligned} \tag{47}$$

- A_t:

$$\Delta(u_{c,t}c_t - u_{l,t}(1-l_t)) + \lambda_{4,t} - \beta E_t \lambda_{4,t+1} \frac{\gamma_t^H I(g_{t+1} = g^H) + \gamma_t^L I(g_{t+1} = g^L)}{u_{c,t+1}} \tag{48}$$

From equation (48)

$$\lambda_{4,t} = -\Delta(u_{c,t}c_t - u_{l,t}(1-l_t)) + \beta E_t \lambda_{4,t+1} \frac{\gamma_t^H I(g_{t+1} = g^H) + \gamma_t^L I(g_{t+1} = g^L)}{u_{c,t+1}} \tag{49}$$

Multiplying both sides by A_t we get

$$\begin{aligned} \lambda_{4,t} A_t = -\Delta A_t(u_{c,t}c_t - u_{l,t}(1-l_t)) + \beta E_t \lambda_{4,t+1} \frac{A_t \gamma_t^H I(g_{t+1} = g^H) + \gamma_t^L I(g_{t+1} = g^L)}{u_{c,t+1}} = \\ = -\Delta A_t(u_{c,t}c_t - u_{l,t}(1-l_t)) + \beta E_t \lambda_{4,t+1} A_{t+1} \end{aligned} \tag{50}$$

where the last equality follows from equation (18).

Iterating forward we obtain

$$\lambda_{4,t} A_t = -\Delta E_t \sum_{j=0}^{\infty} \beta^j A_{t+j}(u_{c,t+j}c_{t+j} - u_{l,t+j}(1-l_{t+j})) \tag{51}$$

Inserting (51) into (44) we get

$$\begin{aligned} & u_{c,t} + \Delta A_t(u_{cc,t}c_t + u_{c,t}) - \lambda_{1,t}\alpha_t u_{cc,t} I(g_t = g^H) \\ & - \lambda_{2,t}\alpha_t u_{cc,t} I(g_t = g^L) - \Delta \frac{u_{cc,t}}{u_{c,t}} E_t \sum_{j=0}^{\infty} \beta^j A_{t+j}(u_{c,t+j}c_{t+j} - u_{l,t+j}(1-l_{t+j})) = \lambda_{3,t} \end{aligned} \tag{52}$$

A.3

Under rational expectations the following equalities hold

$$\begin{aligned} \gamma_{t-1}^H &= u_{c,t}(g_t = g^H) \forall t \\ \gamma_{t-1}^L &= u_{c,t}(g_t = g^L) \forall t \end{aligned}$$

which implies that $A_t = A_{t-1} = 1 \forall t$. The Lagrangian collapses to

$$\begin{aligned} \mathcal{L} = & E_0 \sum_{t=0}^{\infty} \beta^t [u(c_t, l_t) + \Delta(u_{c,t}c_t - u_{l,t}(1-l_t)) + \lambda_{1,t}(\gamma_t^H - \gamma_{t-1}^H) + \lambda_{2,t}(\gamma_t^L - \gamma_{t-1}^L) \\ & + \lambda_{3,t}(1 - l_t - c_t - g_t)] - \Delta u_{c,0} b_{-1} \end{aligned}$$

The first-order conditions with respect to γ_t^H and γ_t^L are

- γ_t^H :

$$\lambda_{1,t} = \beta E_t \lambda_{1,t+1} \tag{53}$$

- γ_t^L :

$$\lambda_{2,t} = \beta E_t \lambda_{2,t+1} \tag{54}$$

which imply that the only solution is $\lambda_{1,t} = \lambda_{2,t} = 0$. The first-order condition with respect to consumption and leisure are

$$u_{c,t} + \Delta(u_{cc,t}c_t + u_{c,t}) = \lambda_{3,t} \tag{55}$$

$$u_{l,t} + \Delta(u_{l,t} - u_{ll,t}(1 - l_t)) = \lambda_{3,t} \tag{56}$$

which are exactly the optimality conditions found in a rational expectations framework (see Lucas and Stokey (1983)) in which expectations do not depend on the current consumption level and in which there is no distortion into agents' beliefs that the government has to manipulate optimally.

A.4 Proof of Proposition 2

The period-by-period budget constraint implies that the following equality

$$b_{t-1}A_{t-1}\gamma_{t-1} = \sum_{j=0}^{\infty} \beta^j A_{t+j}(u_{c,t+j}c_{t+j} - u_{l,t+j}(1 - l_{t+j})) \tag{57}$$

Equation (25) can be written as

$$\begin{aligned} &u_{c,t} + \Delta A_t(u_{cc,t}c_t + u_{c,t}) - \lambda_{1,t}\alpha_t u_{cc,t} - \\ &\Delta \frac{u_{cc,t}}{u_{c,t}} b_{t-1}A_{t-1}\gamma_{t-1} = u_{l,t} + \Delta A_t(u_{l,t} - u_{ll,t}(1 - l_t)) \end{aligned} \tag{58}$$

Assuming that agents' initial belief over marginal utility of consumption is close to the actual marginal utility under rational expectations, $A_{t-1} \approx 1$. Moreover, beliefs close to rational expectations implies that also the optimal fiscal policy under learning is close to that under rational expectations, therefore $b_{t-1} \approx 0$. Under these assumptions and the one that α_t is small enough, equation (25) can be rewritten as

$$u_{c,t} + \Delta A_t(u_{cc,t}c_t + u_{c,t}) \approx u_{l,t} + \Delta A_t(u_{l,t} - u_{ll,t}(1 - l_t)) \tag{59}$$

For the utility function given in equation (26) the previous equation can be expressed as

$$u_{c,t} = \frac{2\Delta\gamma_{t-1}A_{t-1}}{-1 + \sqrt{1 + 4\Delta A_{t-1}\gamma_{t-1}}} \tag{60}$$

1. Taking the first derivative of (60) we get

$$\frac{\partial u_{c,t}}{\partial \gamma_{t-1}} = \frac{-1+\sqrt{1+4\Delta\gamma_{t-1}} - 2\Delta\gamma_{t-1}\frac{1}{\sqrt{1+4\Delta\gamma_{t-1}}}}{(-1+\sqrt{1+4\Delta\gamma_{t-1}})^2} \tag{61}$$

which is positive being $\Delta > 0$ and $\gamma_{t-1} > 0$.

The second derivative with respect to γ_{t-1} is equal to

$$\frac{\partial^2 u_{c,t}}{\partial \gamma_{t-1}^2} = \frac{\frac{4\Delta^2\gamma_{t-1}}{\sqrt{x}x}(-1+\sqrt{x})^2 - \frac{4\Delta}{\sqrt{x}}(-1+\sqrt{x})(-1+\sqrt{x}-\frac{2\Delta\gamma_{t-1}}{\sqrt{x}})}{(-1+\sqrt{x})^4} \tag{62}$$

where $x = \sqrt{1+4\Delta\gamma_{t-1}}$. After some algebra, it can be shown that equation (62) is negative, being $\Delta > 0$ and $\gamma_{t-1} > 0$

2. Imposing $T(\gamma^*, A_{t-1}) = \gamma^*$ we get that the fixed point is given by

$$\gamma^* = 1 + \Delta A_{t-1} > 0 \tag{63}$$

3. Imposing $\gamma_{t-1} < \frac{3}{4\Delta}$ implies that $-1+\sqrt{1+4\Delta\gamma_{t-1}} < 1$. It follows that the learnability condition

$$\frac{\partial u_{c,t}}{\partial \gamma_{t-1}} \mid_{\gamma^*} < 1 \tag{64}$$

is satisfied because

$$-\frac{2\Delta\gamma^*}{\sqrt{1+4\Delta A_{t-1}}} < 0 \tag{65}$$

A.5 Proof of Proposition 3

Suppose that if in period t $u_{c,t} - \gamma_{t-1} < 0$, then $u_{c,t+1} - \gamma_t < 0$ as well. This condition is satisfied for small enough values of α_t.

Using the argument in the text, the optimal consumption function under learning is a time invariant function of the previous period belief and of a function of past forecast errors and of the initial bond holdings. So

$$c_t^L = c(\gamma_{t-1}, A_{t-1}, b_{-1}) \tag{66}$$

Consumption under rational expectations is a special case of the previous equation, i.e. when the agents' belief about today's marginal utility coincides with the actual marginal utility and the product of the past ratios between expected and actual marginal utility is 1.

$$c^{RE} = c^L(u_{c,t}, 1, b_{-1}) \tag{67}$$

When the government expenditure is constant, the law of motion for A_t is given by

$$A_t = A_{t-1}\frac{\gamma_{t-1}}{u_{c,t}} \tag{68}$$

Substituting backwards in the definition of (68) and taking log we get that

$$lim_{t\to\infty} log A_t = lim_{t\to\infty} \sum_{j=0}^{t} log \frac{\gamma_{j-1}}{u_{c,j}} \tag{69}$$

Proposition 2 shows that the expected marginal utility converges to the actual one. Define N the time when this happens. Then

$$lim_{t\to\infty} log A_t = \sum_{j=0}^{N} log \frac{\gamma_{j-1}}{u_{c,j}} \tag{70}$$

Being the finite sum of finite numbers, $logA_t$ converges to a strictly positive value for initial pessimistic belief and to a strictly negative value for initial optimistic belief. Since the arguments of equation (??) do not converge to the ones in equation (??), the allocation does not either.

To show the second part of the proposition, assume the same utility function as in equation (26). Using equations (55), (56) and (1), consumption under rational expectations is given by

$$c_t^{RE} = \frac{1}{1+\Delta} \forall t \geq 1 \tag{71}$$

At $t = 0$, for any given b_{-1}, consumption under RE is equal to

$$c_0 = \frac{1 + \sqrt{(}1 + 4\Delta b_{-1}(1+\Delta))}{2(1+\Delta)} \tag{72}$$

where the equilibrium Δ is the value that guarantees that the implementability condition under rational expectations

$$\sum_{t=0}^{\infty} \beta^t (u_{c,t} c_t - u_{l,t}(1-l_t)) = u_{c,0} b_{-1} \tag{73}$$

evaluated at the optimal allocation, is satisfied. In particular,

$$\begin{aligned} l_0 + \sum_{t=1}^{\infty} \beta^t l_t = 1 - c_0 - g + \frac{\beta}{1-\beta}(1 - \frac{1}{1+\Delta} - g) = \\ \frac{1-g}{1-\beta} - \frac{\beta}{1-\beta}\frac{1}{1+\Delta} = b_{-1} \frac{4(1+\Delta)^2 + (1+\sqrt{(}1+4\Delta b_{-1}(1+\Delta))^2}{2(1+\sqrt{(}1+4\Delta b_{-1}(1+\Delta)))(1+\Delta)} \end{aligned} \tag{74}$$

Denote the positive root of equation (74) as $\Delta^* = \Delta(b_{-1})$. Inserting Δ^* into equation (71) we get

$$c_t^{RE} = \frac{1}{1+\Delta^*} = \frac{1}{1+\Delta(b_{-1})} \tag{75}$$

Let $c^L = lim_{t\to\infty} c_t^L(\gamma_{-1})$. The initial holding of bonds such that the consumption under learning converges to the one under rational expectations starting with that amount of bond is defined by the equation

$$c^L = \frac{1}{1+\Delta(b_{-1})} \tag{76}$$

A.6

To discuss the quality of the learning equations used by agents to predict one-step-ahead state-contingent marginal utility of consumption, I use the Epsilon-Delta Rationality criterion (EDR), as formalized in Marcet and Nicolini (2003). Define

$$\pi^{\epsilon,T} \equiv P(\frac{1}{T}\sum_{t=0}^{T}[u_{c,t}-\gamma_{t-1}]^2 < \frac{1}{T}\sum_{t=0}^{T}[u_{c,t}-E_{t-1}u_{c,t}]^2+\epsilon)$$

which is a function of ϵ and T. $E_{t-1}u_{c,t}$ denote the expectations of an agent who knows the whole economic structure of the model. For small ϵ, $\pi^{\epsilon,T}$ is the probability that, after T periods, the sample prediction error made by boundedly rational agents is almost as small as the sample prediction error made by fully rational agents.

The learning mechanism (12) with $\alpha_t = \alpha$ satisfies EDR for (ϵ,δ,T) if $\pi^{\epsilon,T} \geq 1-\delta$. Table 5 shows $\pi^{\epsilon,T}$ for different values of ϵ (across columns) and T (across rows). Reported values are computed out of 1000 simulations. After 10 periods, there is an 80 percent probability that the prediction error made by boundedly rational agents is at most 3 percent higher than the prediction error made by fully rational agents. This result suggests that ι agents' initial beliefs about marginal utility of consumption are quite close to their rational expectation values and that $\iota\iota$ agents follow a learning scheme that generates good forecasts even along the transition towards the full rationality.

Table 1: Statistics of the allocation under rational expectations

	Mean	St.Dev.	Autocorr
consumption	0.423	0.01	0.6
leisure	0.5	0.01	0.6
labor tax rate	0.16	0.003	0.6
market value of debt	0.04	0.03	0.6
primary surplus	0.003	0.02	0.6

Table 2: Statistics under learning after convergence of beliefs

	Initially pessimistic agents			Initially optimistic agents		
	Mean	St.Dev.	Autocorr	Mean	St.Dev.	Autocorr
consumption	.416	.015	.6	.45	1e-4	.6
leisure	.51	.01	.6	.47	1e-4	.6
labor tax rate	.18	.01	.6	.04	3e-6	.6
market value of debt	.38	.04	.6	-0.35	3e-4	.6
primary surplus	.015	.02	.6	-.056	.02	.6

Table 3: Statistics under RE and learning for the first 30 periods

	Rational expectations			Initially pessimistic agents		
	Mean	St.Dev.	Autocorr	Mean	St.Dev.	Autocorr
consumption	.42	.012	.53	.43	.015	.62
leisure	.5	.012	.53	.49	.014	.73
labor tax rate	.16	.003	.526	.14	.03	.95
market value of debt	.04	.03	.526	.28	.08	.86
primary surplus	.001	.02	.526	-.006	.02	.73

Table 4: OLS estimates and t-statistics (in parenthesis) with i.i.d. government shock

	α	β	R^2
$\tau_t^{RE} = \alpha + \beta\tau_{t-1}^{RE} + \varepsilon_t$	0.1562 (7.0467)	-0.0171 (-0.1207)	0.9996
$\tau_t^{L} = \alpha + \beta\tau_{t-1}^{L} + \varepsilon_t$	0.0288 (2.3923)	0.8213 (10.7910)	0.9960

Table 5: $\pi^{\epsilon,T}$

$T\backslash\epsilon$	0.04	0.03	0.02
5	1	.4	0
10	1	.8	0
15	1	1	.06
20	1	1	.4

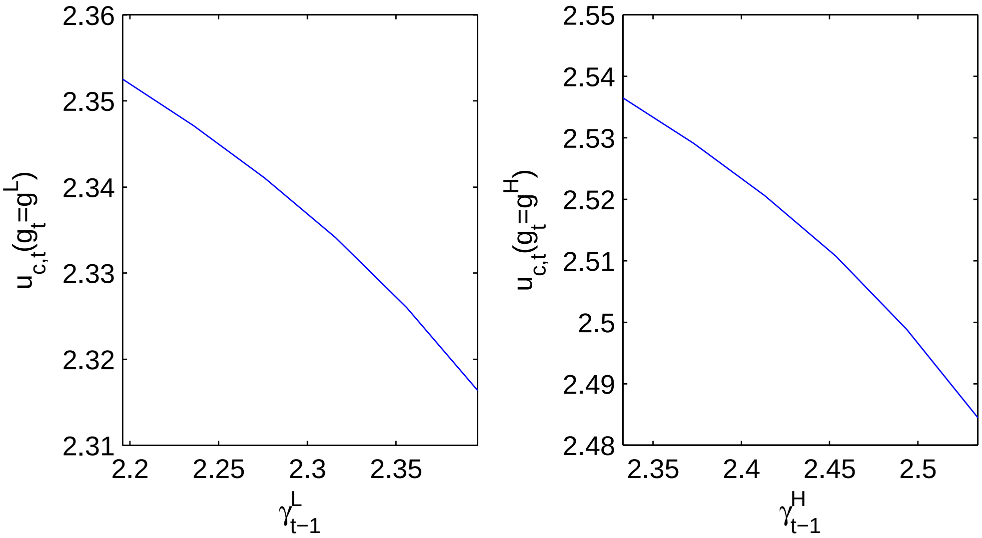

Figure 1: T-mapping

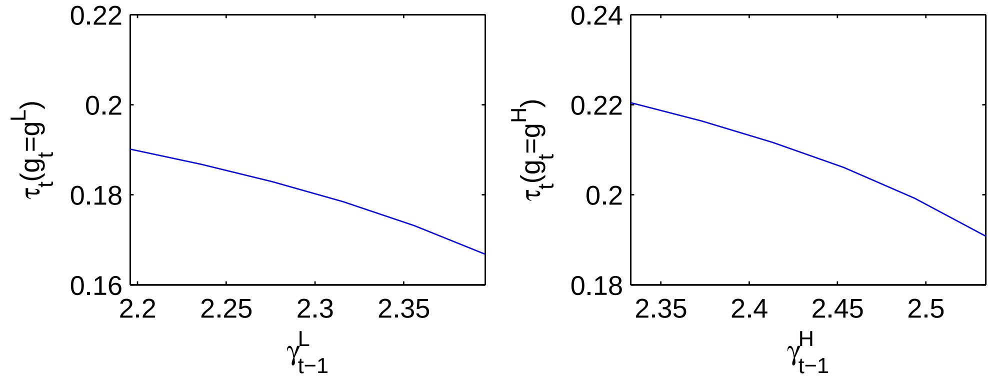

Figure 2: Tax policy function

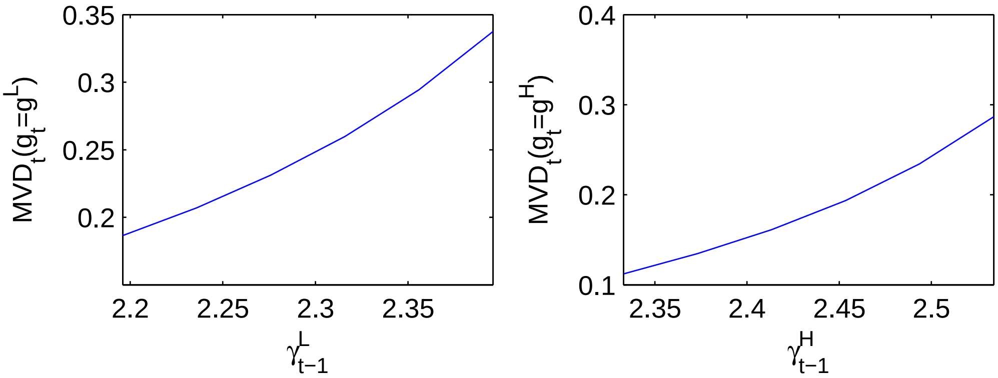

Figure 3: Bond holdings policy function

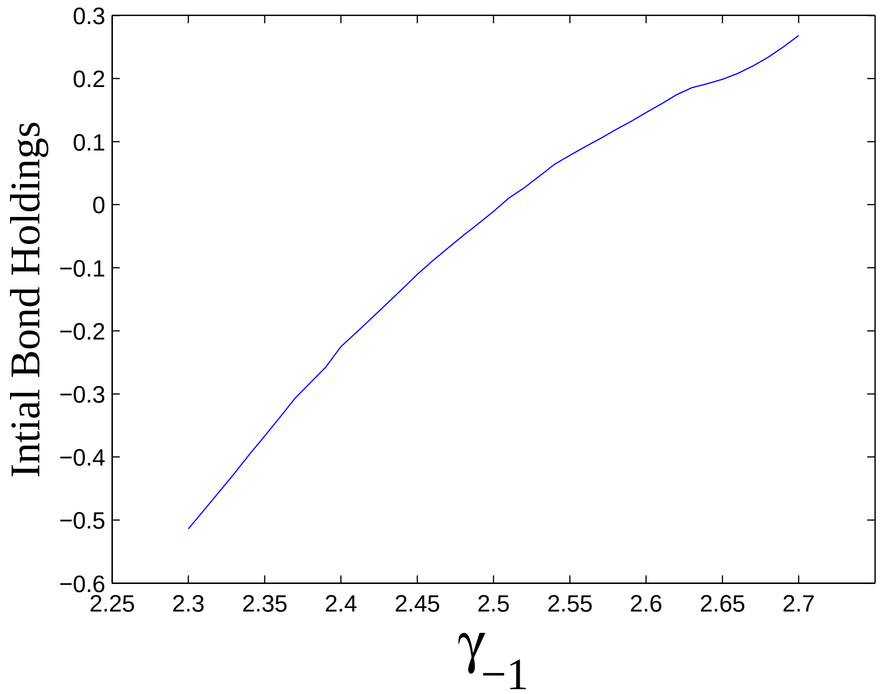

Figure 4: Beliefs and initial debt

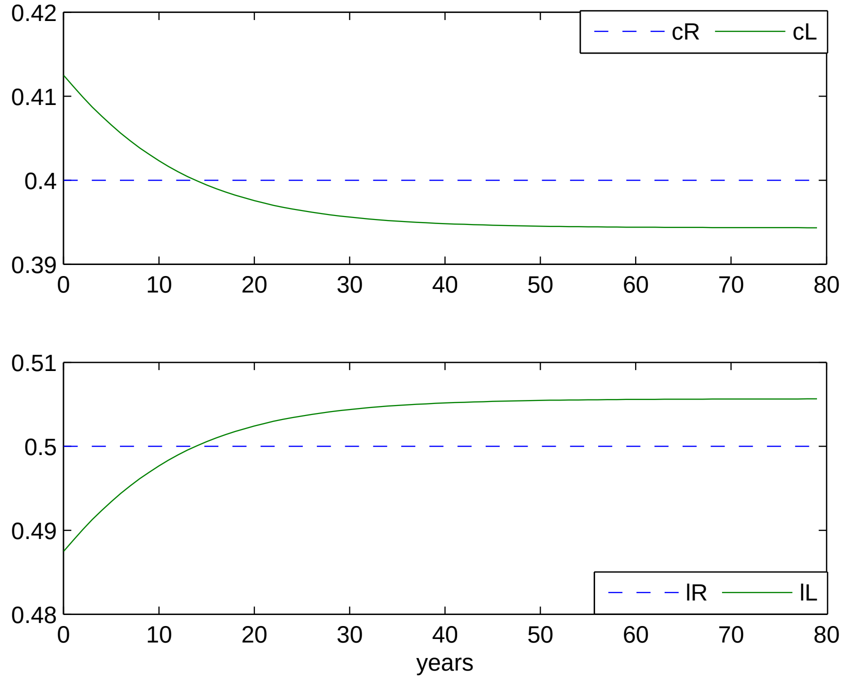

Figure 5: Consumption and leisure under RE and under learning dynamics

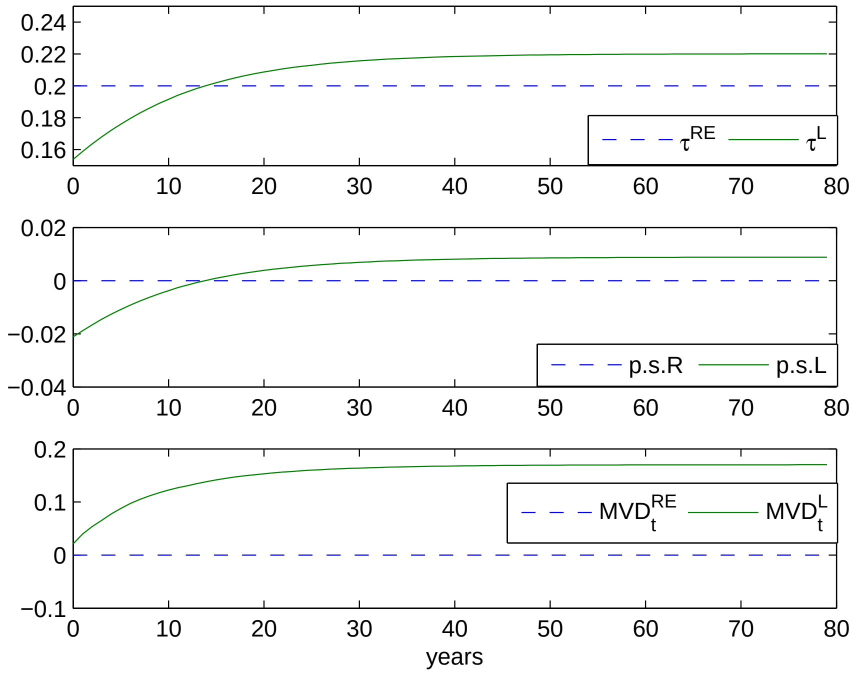

Figure 6: Taxes, primary surplus and debt under RE and under learning dynamics

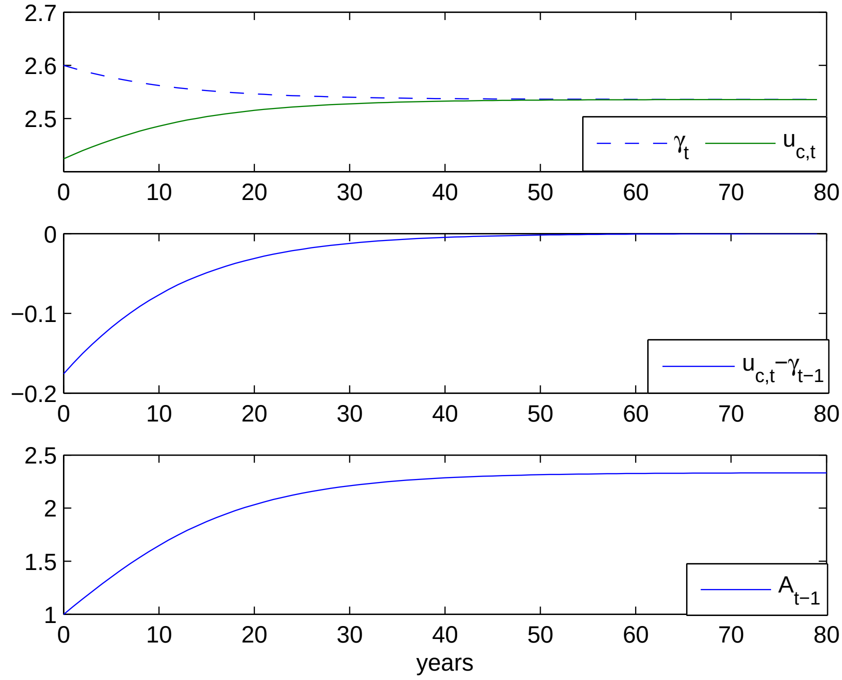

Figure 7: Forecast errors, history and non convergence to the RE values

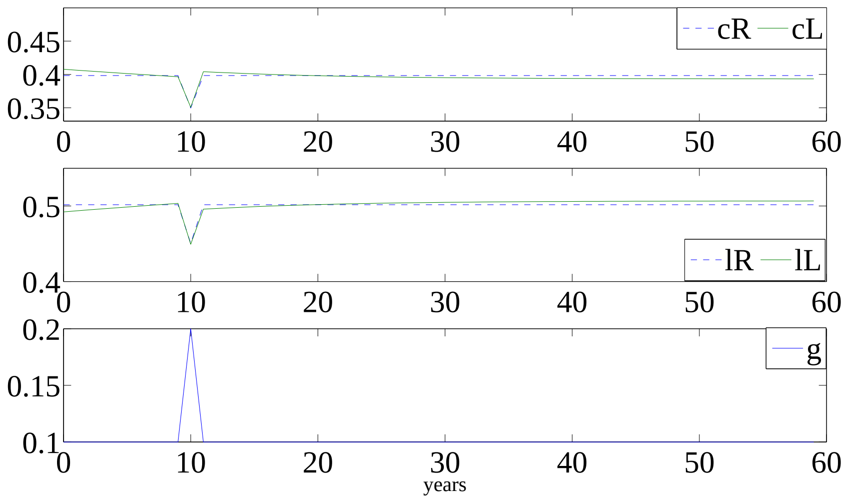

Figure 8: Consumption and leisure under RE and under learning dynamics

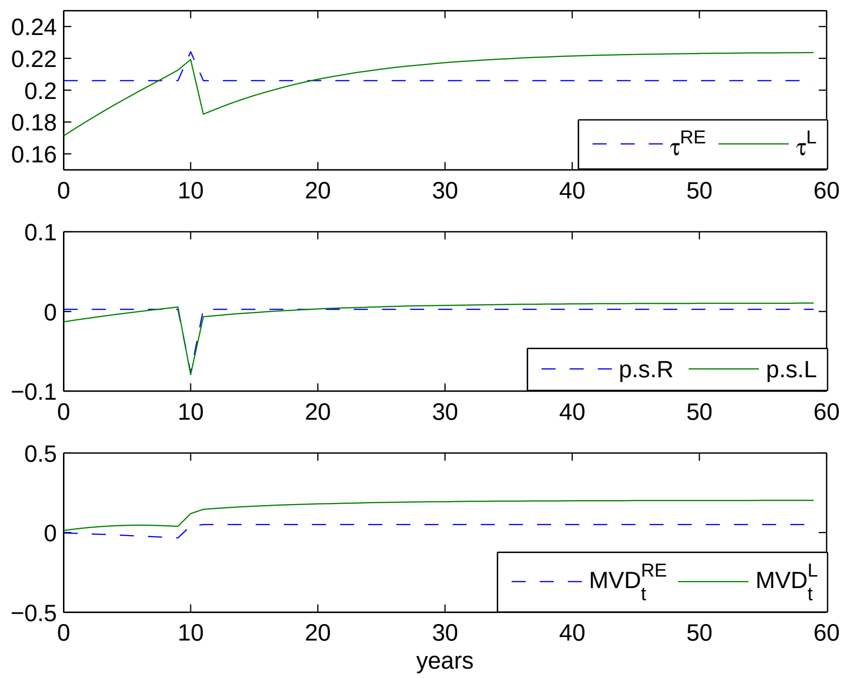

Figure 9: Taxes, primary surplus and debt under RE and under learning dynamics

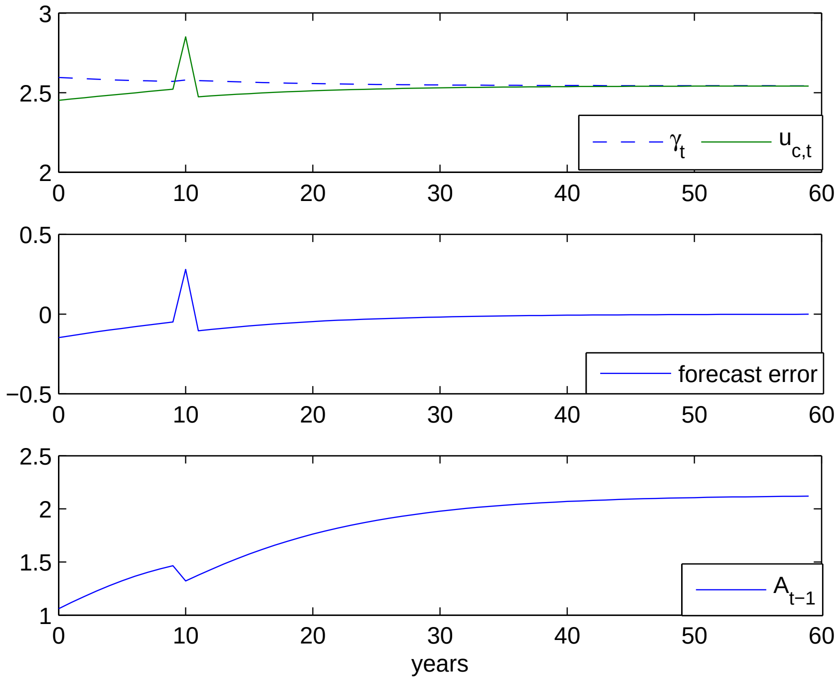

Figure 10: Forecast errors, history and non convergence to the RE values

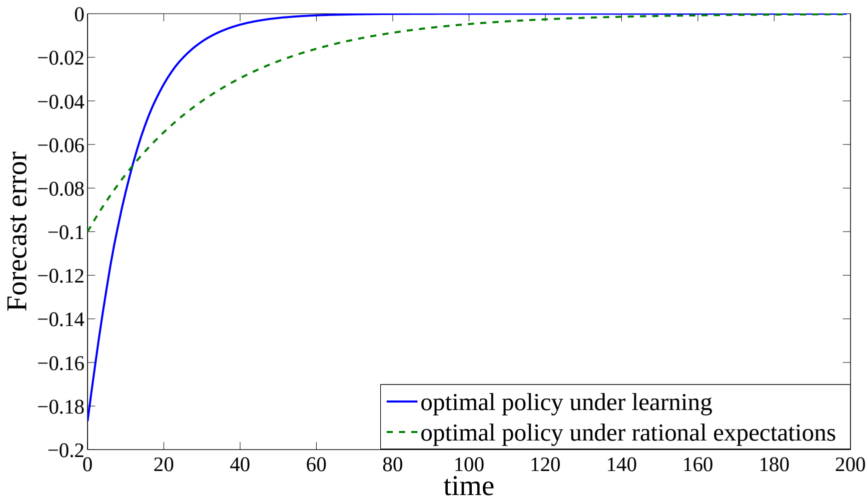

Figure 11: Forecast Error

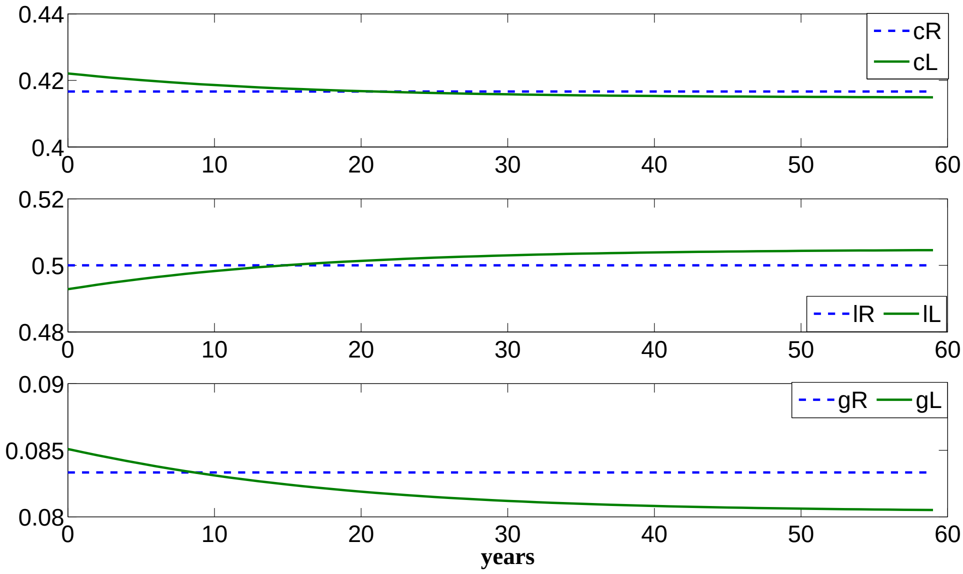

Figure 12: Consumption and leisure under RE and under learning dynamics

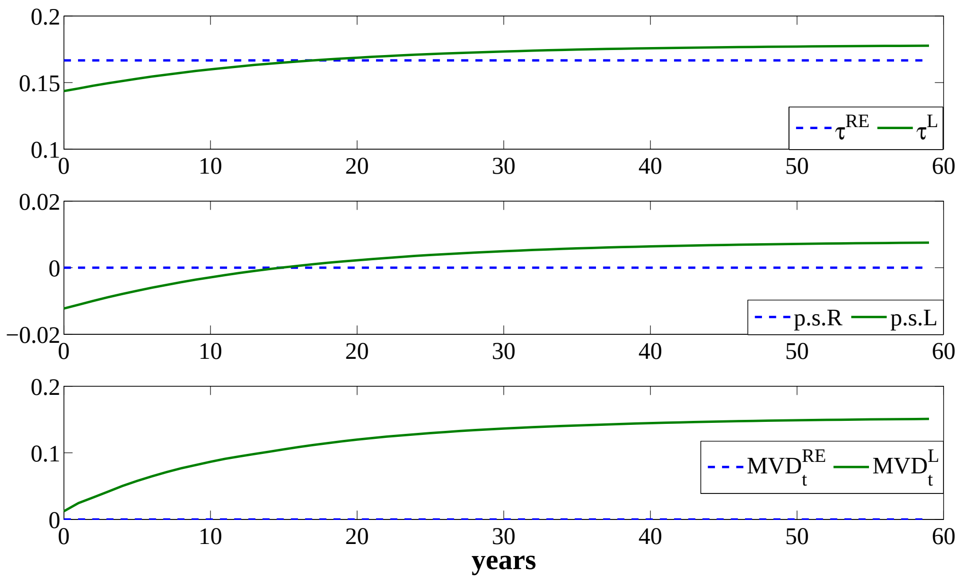

Figure 13: Taxes, primary surplus and debt under RE and under learning dynamics

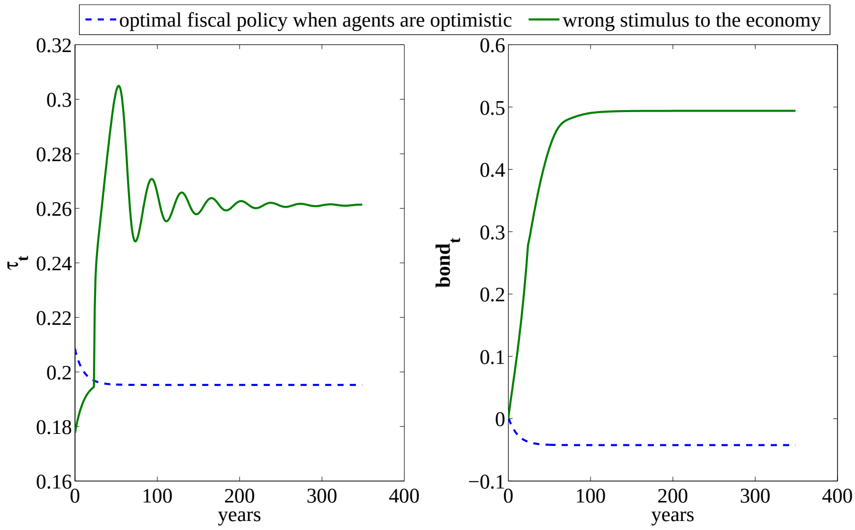

Figure 14:

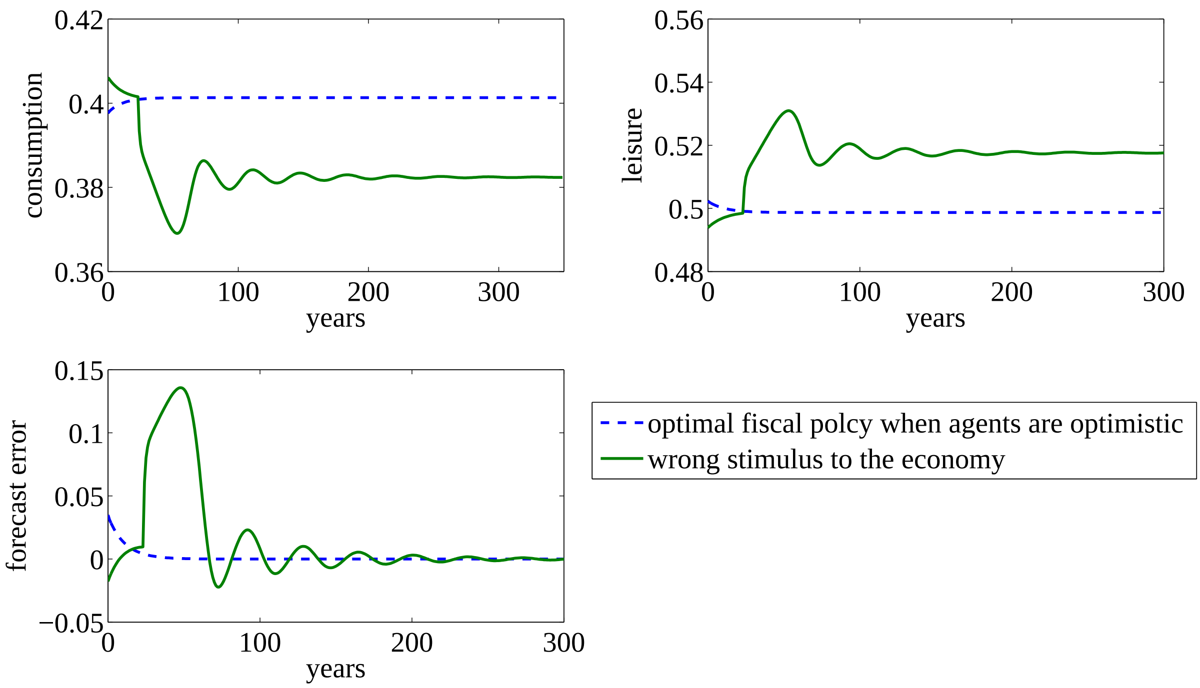

Figure 15:

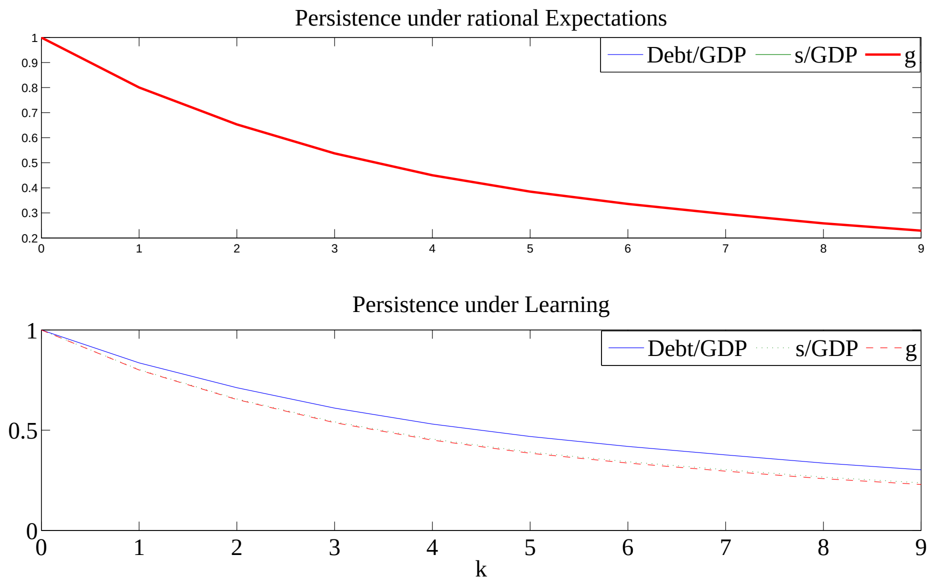

Figure 16: Top Panel: RE framework; Bottom Panel: Learning framework

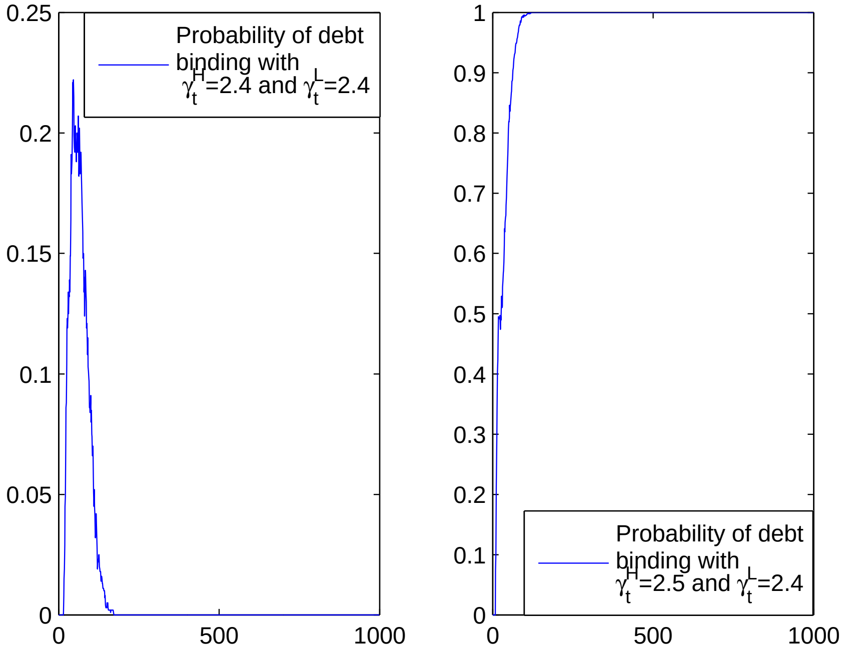

Figure 17: Comparison of probabilities of binding debt limits