Anticipated Fiscal Policy and Adaptive Learning[*]

George W. Evans, University of Oregon
Seppo Honkapohja, University of Cambridge
Kaushik Mitra, University of St Andrews

February 18, 2007

Abstract

We consider the impact of anticipated policy changes when agents form expectations using adaptive learning rather than rational expectations. To model this we assume that agents combine limited structural knowledge with a standard adaptive learning rule. We analyze these issues using two well-known set-ups, an endowment economy and the Ramsey model. In our set-up there are important deviations from both rational expectations and purely adaptive learning. Our approach could be applied to many macroeconomic frameworks.

JEL classification: E62, D84, E21, E43

Key words: Taxation, expectations, Ramsey model

1   Introduction

The hypothesis of rational expectations has in recent decades been the main paradigm for expectation formation in macroeconomics. The literature on adaptive learning has provided support for rational expectations by showing how, in many cases, boundedly rational agents with limited knowledge can

1

converge to a rational expectations equilibrium (REE) in the long run. More specifically, provided an appropriate stability condition is satisfied, REE is an asymptotic outcome of the use of least-squares and related learning rules by private agents.

In recent years, a new literature has developed emphasizing the importance of learning for policy design. This literature has several themes. First, the conditions for stability of the REE under least-squares learning imply constraints on the policy rule parameters – not all policies consistent with an REE guarantee convergence of learning. Second, one can compare the performance of different policy rules under the learning dynamics. A third aspect is the properties of the dynamic learning paths following a change in the policy rule. This last strand is the starting point of the current paper.

Returning to rational expectations, one of its salient contributions has been the idea that agents look forward and can anticipate the effects of an announced future shift in policy. This point was developed for the perfect foresight case by Sargent and Wallace (1973b). Using the forward solution emphasized by Sargent and Wallace (1973a), the analysis of the effects of anticipated policy changes was incorporated into graduate textbooks, e.g. Blanchard and Fischer (1989), p. 219.[1] Thus far, the literature on learning and policy changes has only considered situations where the changes come as a surprise to agents, i.e. the policy change is unanticipated, and agents begin to learn the new equilibrium as data arrive after the policy change. Such changes are examined in Evans, Honkapohja, and Marimon (2001), Marcet and Nicolini (2003) and Giannitsarou (2006). To our knowledge, no paper has examined the dynamic learning paths when an announced policy change is anticipated by the agents who do not have a complete structural model of the economy.[2]

This paper proposes a learning model in which agents combine limited structural knowledge about the economy with adaptive learning for other variables that they need to forecast. An advantage of this approach is that it is flexible in terms of the extent of structural knowledge that agents are assumed to possess. Here we explore one natural case in which agents know the government budget constraint and the announced policy shift, while knowl-

---

[1]Lucas (1976), of course, also examined the impact of policy changes within a single policy regime.

[2]Evans and Ramey (1992), Evans and Ramey (1998) do combine forward-looking behavior with less than fully rational expectations However, in their approach agents know the full structural economic model and are impeded by calculation costs.

edge of the overall structure of the economy remains incomplete.

More specifically, we consider a simple competitive representative-agent endowment economy in which the government purchases output for current consumption and levies lump-sum taxes. In the baseline case, balanced-budget spending changes are assumed and we analyze the dynamics arising from permanent and temporary anticipated changes in government spending/taxation. In the simplest scenario, starting in a steady-state equilibrium we suppose that the government announces a future permanent change in government purchases and taxes. To make their consumption plans, agents need to forecast future (real) interest rates and taxes. The future path of taxes is the structural element that we provide to agents. However, agents must forecast future interest rates using only the current history of data, revising these forecasts over time in line with standard learning models.

In treating the effects of anticipated policy changes under learning, the issue of the length of the forecast horizon becomes clearly relevant. Much of the literature on convergence to REE under adaptive learning assumes that agents' behavior is based on Euler rules that require only short-term forecasts. These behavioral rules do not explicitly incorporate the intertemporal budget constraint or transversality conditions, but satisfy these conditions ex-post when learning is convergent. This is a convenient model of bounded rationality that emphasizes short-horizon decision making under limited knowledge. It is also possible to formulate a model of bounded rationality and adaptive learning in which agents explicitly use long or infinite-horizon decision rules. See Sargent (1993), Evans, Honkapohja, and Romer (1998), Bullard and Duffy (2001), and especially Preston (2005), Preston (2006).

In many cases the length of the forecast horizon is not central for the question of the stability of REE under learning for a given policy rule. However, when considering announced policy changes, it is clearly natural to have a forecasting horizon that at least includes the time of the actual policy shift. In the current paper we therefore assume that private agents adopt an infinite horizon. The results for balanced-budget spending changes are striking. In line with the agents' forward-looking behavior, there are immediate changes in the interest rate and continued evolution over time. However, the resulting dynamics of interest rates under learning are in a sharp contrast to the fully rational path. We also consider the case of repeated shifts in policy in the context of the endowment economy, and show how our approach can be extended to allow agents to incorporate information from previous episodes of policy change.

We extend our analysis to the case of debt financing and to the Ramsey model, in which the dynamics of capital provide a new aspect to the dynamics under learning and under perfect foresight. In these extensions we obtain other surprising results. In the endowment model with debt financing, agents are assumed to know the government budget constraints and to compute their expected tax liabilities using interest rate forecasts and expected levels of government debt. Numerical simulations show that there is "near Ricardian equivalence", i.e. the dynamic paths under learning with debt are very close to the paths with learning in the balanced budget case, despite the substantial differences from the perfect foresight path. Our final extension shows how to apply our approach to the Ramsey model. Simulations show that, under learning, announced future changes in government spending have immediate effects on consumption and the near-term evolution of capital. However, the paths under learning are again substantially different from perfect foresight paths, exhibiting non-monotonicities in which initial overshooting leads to corrections and cyclical convergence.

2 A Simple Endowment Economy

We consider a nonstochastic endowment economy model with exogenous output y and no money. In the adaptive learning literature it is standard to include stochastic shocks, and these could easily be included. However, our central points can be made in a nonstochastic framework and this eases the theoretical presentation and allows for clearer numerical illustrations.

There are a large number of identical households and each household maximizes their utility U_t subject to a budget constraint:

$$
\begin{aligned}
U_t &= \max \; E_t^* \{ \sum_{s=t}^{\infty} \beta^{s-t} \frac{c_s^{1-\sigma}}{1-\sigma} \} \text{ subject to} \\
b_s &= y - c_s - \tau_s + r_s b_{s-1}.
\end{aligned}
\tag{1}
$$

where $\sigma > 0$ and $0 < \beta < 1$. b_s is net assets of the representative household in period s; c_s, τ_s are their consumption and taxes in period s and r_s is the gross real interest rate factor on one-period safe loans from $s-1$ to s, determined at date $s-1$. The Euler equation for consumption is

$$
c_t^{-\sigma} = \beta r_{t+1} E_t^* c_{t+1}^{-\sigma}.
\tag{2}
$$

Here E_t^* denotes potentially subjective expectations at time t and r_{t+1} is assumed known at time t.[3]

Using forward substitution in the flow household budget constraint and imposing the transversality condition,

$$\lim_{T \to \infty} D_{s,s+T}^{-1} b_{s+T} = 0, \tag{3}$$

yields the intertemporal budget constraint (4):

$$0 = r_s b_{s-1} + \sum_{j=1}^{\infty} D_{s,s+j}^{-1} \chi_{s+j} + \chi_s, \text{ where} \tag{4}$$

$$\chi_{s+j} = y - c_{s+j} - \tau_{s+j}, \text{ for } j = 0, 1, 2, \dots \text{ and } D_{s,s+j} = \prod_{i=1}^{j} r_{s+i}.$$

(The transversality condition will be automatically satisfied for the policies that we consider.) From equation (2) we get

$$c_t = \beta^{-\frac{1}{\sigma}} r_{t+1}^{-\frac{1}{\sigma}} c_{t+1}^e(t), \tag{5}$$

where $c_{t+1}^e(t)$ is short-hand notation for $E_t^* c_{t+1}$. As there are no random shocks we assume point expectations. Forward substitution of (5) gives

$$c_t = \beta^{-\frac{s}{\sigma}} (\prod_{i=1}^{s} r_{t+i}^e(t))^{-\frac{1}{\sigma}} c_{t+s}^e(t) \text{ or}$$

$$c_{t+s}^e(t) = c_t \beta^{\frac{s}{\sigma}} (\prod_{i=1}^{s} r_{t+i}^e(t))^{\frac{1}{\sigma}} \equiv c_t \beta^{\frac{s}{\sigma}} (D_{t,t+s}^e(t))^{\frac{1}{\sigma}}. \tag{6}$$

Here $r_{t+i}^e(t)$ refers to expectations of the interest rate r_{t+i} formed at time t, i.e., $E_t^* r_{t+i}$, and $D_{t,t+s}^e(t)$ denotes the expectations of the interest rate factor $D_{t,t+s}$ formed at time t, defined as follows:

$$D_{t,t+s}^e(t) = r_{t+1} \prod_{i=2}^{s} r_{t+i}^e(t), \ s \geq 2 \text{ and } D_{t,t+1}^e(t) = r_{t+1}. \tag{7}$$

Similarly, $\tau_{t+j}^e(t)$ denotes the expected lump-sum tax in period $t + j$ forecasted at time t. Substituting (6) into (4) yields

$$0 = r_t b_{t-1} + \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1} (y - \tau_{t+j}^e(t) - c_t \beta^{\frac{j}{\sigma}} (D_{t,t+j}^e(t))^{\frac{1}{\sigma}}) + y - c_t - \tau_t,$$

---

[3]For simplicity, different consumers are assumed to have identical expectations.

where future dated variables are replaced by their expected values. Solving for current consumption c_t yields

$$c_t(1 + S_D(t)) \quad = \quad r_t b_{t-1} + y - \tau_t + S_N(t), \text{ where} \tag{8}$$

$$S_N(t) \quad \equiv \quad \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1}(y - \tau_{t+j}^e), \tag{9}$$

$$S_D(t) \quad \equiv \quad \sum_{j=1}^{\infty} \beta^{j/\sigma}(D_{t,t+j}^e(t))^{\sigma^{-1}-1}. \tag{10}$$

$S_N(t)$ is the expected present value of future expected net income and $c_t(1 + S_D(t))$ the expected present value of current and future consumption.

Note that c_t in (8) depends on $r_{t+i}^e(t), \tau_{t+j}^e(t), r_{t+1}, \tau_t,$ and $r_t b_{t-1}$. Expectations are assumed to be formed at the beginning of period t. The full description of equilibrium requires specifying government tax and spending policies. The market clearing condition in per capita terms

$$c_t + g_t = y$$

then determines c_t and r_{t+1} given the expectations of the consumers.

3 Balanced Budget for Government

The general form of the government flow budget constraint (in per capita terms) is

$$b_t + \tau_t = g_t + r_t b_{t-1}, \tag{11}$$

where b_t is government debt in the form of one-period risk-free real bonds issued in period t and g_t denotes government purchases in period t.

In this section, however, we assume that the government runs a balanced budget and bonds are in zero net supply so that

$$\tau_t = g_t, \text{ and } b_{t-1} = 0.$$

Hence, b_t in Section 2 may be interpreted as lending to other households and this must equal zero in equilibrium.

Before considering the learning scenario, we consider what happens in the benchmark rational expectations (RE) case when there is an announced policy change in government spending at some time in the future.

3.1 Benchmark RE case

At the beginning of period $t = 1$, a policy announcement is made that the level of government spending will change permanently from g_0 to g_1 at a specified date T_p in the future. The policy announcement is assumed to be credible and known to the agents with certainty. With a balanced budget, this means equivalently that there is an anticipated change in (per capita) taxes, i.e., $\tau_t = \tau_0 = g_0$ when $t < T_p$ and $\tau_t = \tau_1 = g_1$ when $t \geq T_p$. From the balanced budget and market clearing

$$
\begin{aligned}
c_t &= y - \tau_0, \text{ for } t < T_p \\
c_t &= y - \tau_1, \text{ for } t \geq T_p.
\end{aligned}
$$

The Euler equation (2) under perfect foresight gives us $r_t = \beta^{-1}$ when $t < T_p - 1$. At date $T_p - 1$ the interest rate obtained from (2) is

$$
r_{T_p}^{RE} = \beta^{-1} \left(\frac{y - \tau_1}{y - \tau_0} \right)^\sigma, \tag{12}
$$

while $r_t = \beta^{-1}$ again for all $t \geq T_p$.

Thus, if $\tau_1 > \tau_0$, the interest rate $r_{T_p}^{RE}$ jumps down for one period, after which it immediately returns to its steady state value. The intuition is straightforward. At date $T_p - 1$, since taxes are known to go up next period, next period's consumption will be lower relative to current consumption. Hence, the interest rate between $T_p - 1$ and T_p must go down to ensure that the marginal rate of substitution of consumption between $T_p - 1$ and T_p is equal to the interest rate.

3.2 Learning under a balanced government budget

In the learning scenario, the policy announcement is assumed to take the same form as in Section 3.1. In the standard adaptive learning approach, private agents would need to forecast future taxes as well as interest rates, since both are required for to solve for the optimal level of consumption using (8)-(10). We continue to follow this approach with respect to interest rates, but take the radically different approach for taxes of assuming that agents understand the future course of taxes implied by the announced policy. In effect we are giving the agents structural knowledge of one part of the economy: the fiscal implications of the announced future change in government spending. To keep things simple we assume in this section (and in

most of the paper) that the government operates and is known to operate under a balanced-budget rule. Given this structural knowledge of the government budget constraint and the announced path of government spending, the agents can thus use $\tau_0 = g_0$, for $t < T_p$, and $\tau_1 = g_1$, for $t \geq T_p$, for their forecasts of future taxes. Of course, for simplicity we are assuming that the announced policy change is fully credible. It would be possible to relax this assumption within the general framework of our approach.

We now turn to the formal details of the adaptive learning component of the model. In forecasting interest rates we assume that agents do *not* have adequate structural information. That is, we assume that private agents do not know the complete economic structure and thus are not able to solve for the path of future interest rates using an economic model. In particular, because agents do not know that agents are identical, they cannot deduce the path that interest rates will follow and must forecast future interest rates using statistical procedures.[4] To forecast future interest rates it is assumed that private agents make forecasts using an econometric time-series model and that its parameters are estimated using past data. The forecasts are input to agent's decision rules and in each period the economy attains an equilibrium for the current period variables given the forecasts of the agents. This temporary equilibrium provides a new data point, which in the next period leads to re-estimation of the parameters, updating of the forecasts and, in turn, to a new temporary equilibrium. The sequence of temporary equilibria may generate parameter estimates that converge to a fixed point corresponding to an REE for the economy. For a general discussion of adaptive learning see Evans and Honkapohja (2001).

In our context, the economy will eventually reach a steady state after the policy has been implemented. A natural way to model learning is to assume that agents use observed interest rates to estimate their mean, which they use to forecast future interest rates. These estimates are updated over time as new data arrives. In the learning literature this is often called "steady-state learning" since the agents are estimating a mean rather than a more complex time series model. Different assumptions about the computation of the mean have been used in the literature. One possibility is the simple arithmetic mean while another is to allow for different weights on the data.

---

[4]The calculation of RE interest rates in (12) requires that agents know their consumption to be $y - \tau_i$, $i = 0, 1$, before and after the tax change. This knowledge in turn comes from common knowledge that agents are identical and there is market clearing.

It is assumed below that agents use exponential discounting of past data, an assumption commonly used in the learning literature when agents are concerned that structural changes may be occurring.

Under steady-state learning, agents use past data to estimate the mean interest rate, which is then used as the forecast of future interest rates for all horizons, i.e.

$$r^e_{t+i}(t) = r^e(t) \text{ for all } i \geq 2$$

so that from (7)

$$D^e_{t,t+s}(t) = r_{t+1} r^e(t)^{s-1}.$$

We make the assumption of steady-steady learning, even though policy-induced structural changes will be occurring in the future, because without adequate knowledge of the complete economic structure, agents have no way to forecast the pattern of future interest rates that will result from the policy change, except through past observations of interest-rates and their evolution as time progresses.[5]

Assuming zero debt, the consumption function (8) in equilibrium implies that

$$c_t = \frac{y - \tau_t + S_N(t)}{1 + S_D(t)}, \tag{13}$$

where $S_N(t)$ and $S_D(t)$ are defined in (9) and (10). It is assumed throughout that expectations satisfy $\beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1} < 1$ and $r^e(t) > 1$, so that all the relevant sums are finite. Imposing the market clearing condition $c_t = y - g_t$ and the balanced budget condition $\tau_t = g_t$ yields

$$(y - g_t) S_D(t) = S_N(t). \tag{14}$$

The expressions for $S_D(t)$ and $S_N(t)$ take the form:

$$S_D(t) = r^{\sigma^{-1}-1}_{t+1} \beta^{\sigma^{-1}} \frac{1}{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}, \tag{15}$$

---

[5]If we were to consider *repeated* policy changes, then agents could use past implementations of policy change to provide information on the time profile of the impact of an announced policy change. We consider this issue below in Section 4. Note also that the agents are assumed not to have the structural information that the interest rate will eventually return to the original steady state value β^{-1}. This assumption could be easily altered.

$$S_N(t) = \frac{y - \tau_0 + (\tau_0 - \tau_1)(r^e(t))^{t+1-T_p}}{r_{t+1}(1 - (r^e(t))^{-1})} \text{ for } t = 1, 2, ..., T_p - 1, \text{ and}$$

$$S_N(t) \equiv \frac{y - \tau_1}{r_{t+1}(1 - (r^e(t))^{-1})} \text{ for } t \geq T_p.$$

The corresponding values for r_{t+1} are

$$r_{t+1} = [\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})}\{1 - \frac{(\tau_1 - \tau_0)}{(y - \tau_0)}(r^e(t))^{t+1-T_p}\}]^\sigma \text{ for } 1 \leq t \leq T_p - 1,$$

$$r_{t+1} = [\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})}]^\sigma \text{ for } t \geq T_p.$$

3.3 Permanent changes in policy

To obtain analytical results, we specialize to the case of logarithmic utility, i.e. $\sigma = 1$. We have

$$r_{t+1} = H(r^e(t), t), \text{ where}$$

$$H(r^e(t), t) = \begin{cases} \frac{1-\beta}{\beta(1-(r^e(t))^{-1})}[1 - \frac{(\tau_1-\tau_0)}{(y-\tau_0)}(r^e(t))^{t+1-T_p}], 1 \leq t \leq T_p - 1 \\ \frac{1-\beta}{\beta(1-(r^e(t))^{-1})}, t \geq T_p. \end{cases} \quad (16)$$

We now study this system under "steady state" learning with exponential discounting. The steady-state learning algorithm is simply

$$r^e(t+1) = r^e(t) + \gamma(r_{t+1} - r^e(t)),$$

where $0 < \gamma < 1$ is referred to as the "gain" parameter that measures the extent to which past data is discounted.[6] Under this algorithm the relative weight on data j periods earlier is $(1 - \gamma)^j$, i.e. past data is discounted at rate $1 - \gamma$. The optimal choice of γ is not straightforward and is most naturally addressed in a stochastic framework, since it involves a trade-off of "filtering" and "tracking." Lower values of γ more effectively filter out random noise, while higher values of γ are better at tracking structural change. Because the optimal choice of γ in general, and in the current context, is not straightforward, we treat γ as a given parameter.[7]

---

[6] Under traditional "least-squares learning" all data receives an equal weight. This corresponds to a decreasing gain sequence $\gamma_t = 1/t$, where t is the sample size.

[7] See Evans and Ramey (2006) and Evans and Honkapohja (2001), Chapter 14, for a discussion of the choice of γ in stochastic models with structural change.

Combining the temporary equilibrium equation with the constant gain learning algorithm we arrive at the dynamic equation for expectations,

$$r^e(t+1) = r^e(t) + \gamma(H(r^e(t), t) - r^e(t)).$$

Assume that the economy is initially at the steady state. We have the following analytical results for the impact of the policy change.[8]

Proposition 1 *Assume that the economy is initially in a steady state and consider a permanent government spending and tax increase $\tau_1 > \tau_0$, which takes place in period T_p and which is announced during period 1. Under learning, for all $\gamma > 0$ sufficiently small we have:*
(a) Interest rate expectations satisfy $(1 - \gamma)^{-1} < r^e(t) < \beta^{-1}$ for all $t \geq 2$ and $\lim_{t \to \infty} r^e(t) = \beta^{-1}$.
(b) The actual temporary equilibrium interest rate r_{t+1} satisfies:
(i) $r_{t+1} < \beta^{-1}$ for $t = 1, \ldots, T_p - 1$,
(ii) $r_{t+1} > \beta^{-1}$ for $t \geq T_p$ and $\lim_{t \to \infty} r_{t+1} = \beta^{-1}$.

Thus, under learning, an announced future permanent increase in government spending leads to low interest rates before T_p, the time of the policy change, and high interest rates on and after T_p, with interest rates eventually converging back to its steady-state value. This is a very different pattern from the outcome under fully rational expectations.

Figures 1 and 2 provide numerical illustrations of the dynamics of the interest rate and interest rate expectations under learning.[9] The figures assume $\sigma = 1$, $\beta = 0.95$, $y = 6$, $\gamma = 1/50$, $g_0 = \tau_0 = 0.9$, $g_1 = \tau_1 = 1.1$ and $T_p = 20$. The learning dynamics are assumed to start at the steady state that prevails before the announced policy change, i.e., $r_1 = r^e(1) = \beta^{-1}$.

Figure 1 illustrates that the dynamics of the interest rate under learning are strikingly different from the RE dynamics. Under learning, at $t = 1$, there is an instantaneous drop in the interest rate r_2. To understand this, note that the present value of consumers' expected net future incomes $S_N(t)$ must be constant in all periods up to $T_p - 1$, since from equation (14), the left side is constant with log utility in all periods up to $T_p - 1$. The effect of higher anticipated taxes by itself is to reduce $S_N(1)$. Since interest rate expectations in period 1, $r^e(1)$, are pre-determined at the beginning of the period (assumed to be at the steady state), the current interest rate r_2 has to fall to keep $S_N(1)$ constant in temporary equilibrium.

---

[8]Proofs of propositions are in the Appendix.
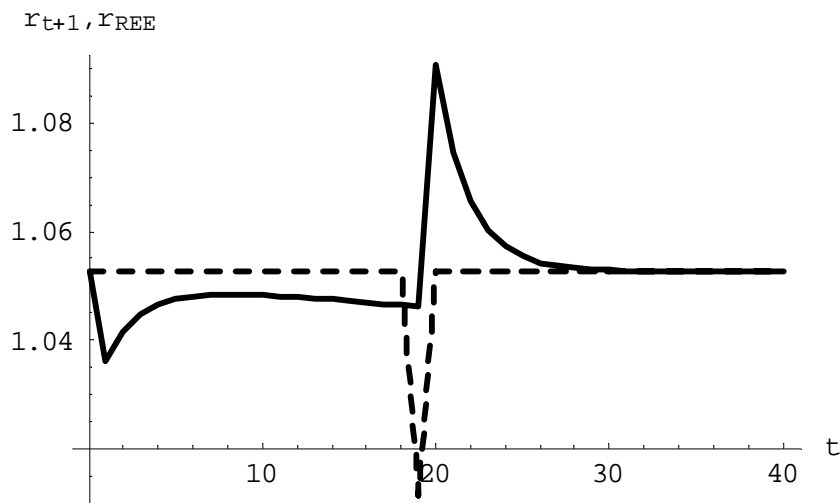[9]The numerical routines are available on request.

Figure 1: r_{t+1} under learning (solid line) and r_{REE} under rational expectations (dashed line), balanced-budget case with the permanent policy change.
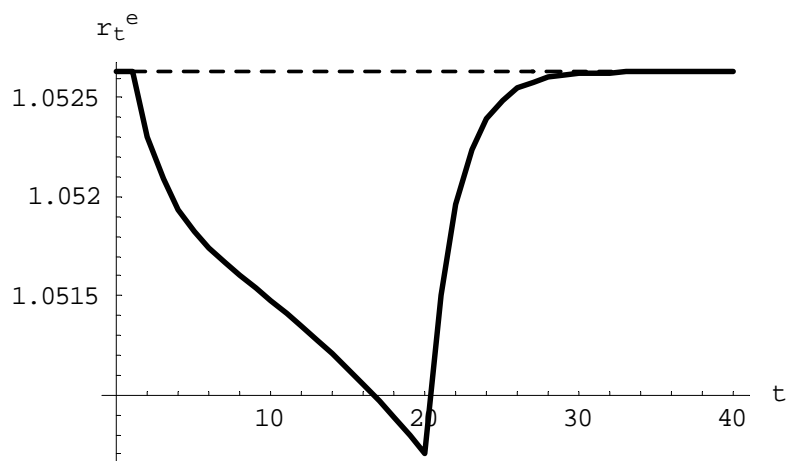

Figure 2: Interest rate expectations, balanced budget case with permanent policy change. The dashed line is the steady state value β^{-1}.

Figure 2 illustrates that interest rate expectations always stay below the steady state value β^{-1}, as proved in Proposition 1. They decrease monotonically from this steady state value till date T_p before rising monotonically towards β^{-1} thereafter.[10] The sharp fall in r_2 induces a gradual reduction in interest rate expectations. A lower value of $r^e(t)$ by itself acts to reduce the discount factors $D^e_{t,t+j}(t)$ and to increase $S_N(t)$, but the latter must be constant up to $T_p - 1$. The subsequent upward movements of r_{t+1} after the initial drop at date 1 compensate for this and keep $S_N(t)$ constant. In other words, lower future expectations of the interest rate increase the expected present value of future net incomes, which by itself increases planned current consumption. But since in equilibrium actual consumption cannot change in this endowment economy, the current interest rate rises to encourage savings and clear the market. Note that relatively small variations in $r^e(t)$ can have substantial effects on r_{t+1} due to the impact on present values.

In the period $T_p - 1$, just before the actual policy change, under rational expectations, there is a sudden sharp fall in $r^{RE}_{T_p}$ whereas under learning, there is only a small change in r_{T_p}. Most prominently, Figure 1 shows that the actual interest rate under learning in the subsequent period T_p spikes upwards, overshooting the steady state value before gradually converging to the steady state. Under rational expectations, on the other hand, the interest rate reverts immediately to its steady state value in period T_p. The big upward spike in r_{T_p+1} under learning also reverses the initial decline in $r^e(t)$ up to this date and induces a gradual increase in $r^e(t)$, from period $T_p + 1$ onwards, towards the steady state value in Figure 2. The latter also explains the monotonic decline in the actual interest rate towards the steady state after period T_p, seen in Figure 1, since from equation (16), for all dates $t \geq T_p$, r_{t+1} is decreasing in $r^e(t)$.[11]

Proposition 1, and Proposition 2 below, state formal results for the "small gain" case. This is the case usually examined in the theoretical literature on adaptive learning, e.g. Cho, Williams, and Sargent (2002), and empirical estimates of γ in the constant gain case are typically in the range $\gamma = 0.02$ to $\gamma = 0.05$ for quarterly macro data (see Orphanides and Williams (2005) and Branch and Evans (2006)). As discussed above, the appropriate value of

---

[10]This qualitative feature seems robust numerically but we have been unable to prove this analytically.

[11]Most qualitative features of the dynamics mentioned in our discussion are robust to variations in T_p, g_0 and g_1. However, the detailed dynamics of the actual interest rate between periods 2 and $T_p - 1$ depend on the parameters, especially on T_p.

γ will depend on the importance of random shocks to the economy as well as the size and frequency of structural change. Relatively small values of γ are usually examined because applied macroeconomic models invariably include a large role for random shocks that would need to be filtered out by agents when making forecasts.

3.4 Temporary Policy Changes

The other natural fiscal policy experiment to consider is a change in government spending that is only temporary. We consider the following policy experiment, assumed fully credible and announced at the start of period 1:

$$
g_t = \tau_t = \begin{cases}
g_0, \ t = 0, ..., T_{g1} - 1 \\
g_1, \ t = T_{g1}, ..., T_{g2} - 1 \\
g_0, \ t \geq T_{g2},
\end{cases}
$$

i.e., government spending and taxes are changed in period T_{g1} and this change is reversed at a later period T_{g2}.

The Appendix works out the details under learning. Again it is convenient to focus on the case of log utility, for which

$$
r_{t+1} = F(r^e(t), t) = \frac{1 - \beta}{\beta(1 - (r^e(t))^{-1})} \tilde{F}(r^e(t), t),
$$

where

$$
\tilde{F}(r^e(t), t) = \begin{cases}
[1 - \frac{(\tau_1 - \tau_0)}{(y - \tau_0)}(r^e(t))^{t+1-T_{g1}}(1 - r^e(t)^{T_{g1}-T_{g2}})], 1 \leq t \leq T_{g1} - 1 \\
[1 + \frac{(\tau_1 - \tau_0)}{(y - \tau_1)}(r^e(t))^{t+1-T_{g2}}], T_{g1} \leq t \leq T_{g2} - 1 \\
1, t \geq T_{g2}.
\end{cases}
$$
(17)

Under steady state learning with constant gain, the expectations are adjusted according to the algorithm

$$
r^e(t + 1) = r^e(t) + \gamma(F(r^e(t), t) - r^e(t)).
$$
(18)

Again we provide formal results for the main features of the solution under learning and then illustrate the quantitative details numerically.

Proposition 2 *Assume that the economy is initially in the steady state and consider a temporary government spending and tax increase $\tau_1 > \tau_0$, which*

14

takes place in period T_{g1} and is reversed in period $T_{g2} > T_{g1}$, with the policy announced during period 1. Under learning, for all $\gamma > 0$ sufficiently small, we have:

(a) $r^e(t) < \beta^{-1}$ in periods $t = 2, \ldots, T_{g1}$, $r^e(t) > \beta^{-1}$ for $t \geq T_{g2}$ and $\lim_{t \to \infty} r^e(t) = \beta^{-1}$.

(b) $r_{t+1} < \beta^{-1}$ for $t < T_{g1}$, $r_{T_{g1}+1} > \beta^{-1}$, $r_{t+1} < \beta^{-1}$ for $t \geq T_{g2}$ and $\lim_{t \to \infty} r_{t+1} = \beta^{-1}$.

We remark that for $T_{g1} + 1 \leq t \leq T_{g2} - 1$ the qualitative results are not theoretically clear-cut.

Under RE, with this temporary policy change, we have $r_{t+1} = \beta^{-1}$ whenever $1 \leq t \leq T_{g1} - 2$, $T_{g1} \leq t \leq T_{g2} - 2$ and for all $t \geq T_{g2}$, while at dates $T_{g1} - 1$ and $T_{g2} - 1$, interest rates are given by

$$r_{T_{g1}} = \beta^{-1}\left(\frac{y - g_1}{y - g_0}\right), r_{T_{g2}} = \beta^{-1}\left(\frac{y - g_0}{y - g_1}\right).$$

Thus, for a temporary increase in government spending the impact under RE is a double spike, with a one-period decrease in the interest rate in the period of the increase in government spending and a one-period increase in the interest rate when government spending reverts to its previous level. The reason behind these changes is the same as in Section 3.1. This solution provides the benchmark RE dynamics, with which we compare the learning dynamics. In our numerical experiment we assume that the parameter values are as before, except that $T_{g1} = 20$, $T_{g2} = 40$.

The dynamics of actual interest rates under learning and under RE, in the balanced-budget case with the temporary policy change, are illustrated in Figure 3. Figure 4 shows the corresponding interest-rate expectations. The dynamics of interest rates under learning (and under RE) until date T_{g1} are qualitatively similar to the case illustrated in Figure 1. This is because T_{g1} is the point of implementation of the first policy change here (which corresponds to T_p in Section 3.2). In particular, the interest rate under learning at date 1 falls immediately, and in period T_{g1} it overshoots the steady state value. The dynamics of interest rates after period T_{g1} tend to go in opposite directions to the dynamics before period T_{g1}. This is because the agents (correctly) anticipate that the policy change will be reversed at the later date T_{g2}. The interest rate stays above the steady state from period T_{g1} to period $T_{g2} - 1$ until it spikes downwards at T_{g2} and undershoots the steady state value.
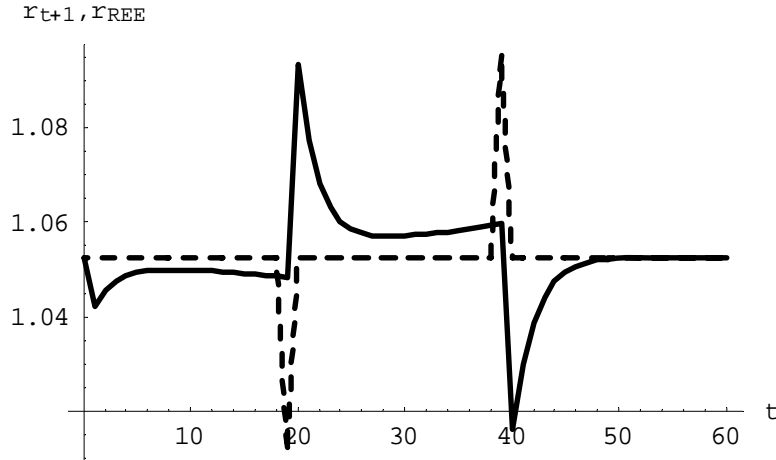
Figure 3: r_{t+1} under learning (solid line) and r_{REE} under rational expectations (dashed line), balanced-budget case with a temporary policy change.

Figure 4 shows that interest rate expectations $r^e(t)$ tend to fall until period T_{g1} before rising thereafter until date T_{g2}. The turn-around in interest rate expectations comes about in period $T_{g1} + 1$ since the interest rate at T_{g1} spikes upward (see Figure 3) which causes these expectations to rise gradually thereafter. $r^e(t)$ continues to increase till date T_{g2} before monotonically decreasing towards the steady state value. The decline in $r^e(t)$ from T_{g2} onwards arises due to the big downwards spike in the actual interest rate at date T_{g2} when the policy change is reversed.[12] Finally, the monotonic decline in $r^e(t)$ from T_{g2} onwards explains the subsequent monotonic upward movement in the actual interest rate, since the interest rate for periods $t \geq T_{g2}$ is decreasing in $r^e(t)$, as shown in (17).

The results for both permanent and temporary policy changes can be summarized as follows. Our framework provides private agents with partial structural knowledge of the economy, specifically the future course of government spending and taxes, but requires agents to use a simple statistical forecasting method for future values of other aggregate variables, specifically interest rates. In line with standard RE results in many models, we find that announced future changes in policy have immediate effects on endoge-

---

[12] We remark that $r^e(t)$ goes above the steady state value in period $T_{g1} + 2$ and stays above this value thereafter.
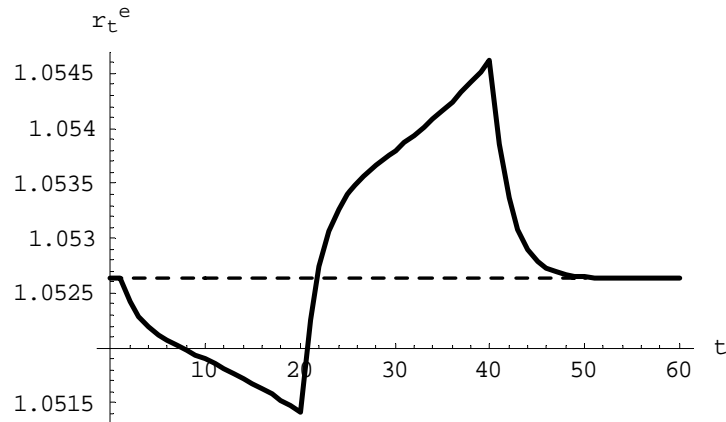
Figure 4: Interest-rate expectations, balanced budget case with the temporary policy change. The dashed line is the steady state value of β^{-1}.

nous variables. However, the resulting time paths are also strikingly different from the one those that would result from fully rational expectations: where RE leads to interest-rate spikes we find that under learning these spikes can be reversed in sign and somewhat delayed. Although these results are specific to the model examined, they show that anticipated policy changes under learning can have qualitative dynamics that are very different from RE.

4 Learning With Repeated Policy Implementations

In Section 3.4, we considered the effect of an anticipated temporary policy change under learning dynamics. The policy change considered was a once and for all change in the sense that it was never repeated. We think this type of policy change is realistic in the sense that policy changes often have unique features, so that frequent repetitions of exactly the same policy change are unlikely. Still, it is of interest to consider the implications under learning if such repetitions do occur, so that agents have the possibility of learning over time across policy experiments.

We thus now consider a situation in which the same policy change is repeated over time and in which the agents have access to data on these

17

repeated policy changes and use this data to update their interest rate path forecasts. One might expect in these circumstances that agents' behavior over time, under this type of learning dynamics, might converge to the one predicted under RE for these anticipated policy changes. We will see that this is the case, though convergence can be quite slow, with substantial qualitative differences from RE remaining for many repetitions.

To keep the analysis simple, we consider repeated one-period government spending shocks under a balanced budget. More precisely, each policy change considered by the agent is assumed to be of the form

$$g_t = \tau_t = \begin{cases} g_0, \, t = 0, ..., T_{g1} - 1 \\ g_1, \, t = T_{g1} \\ g_0, \, t > T_{g1}. \end{cases}$$

Since we are considering repeated policy changes, this specification is not fully correct for $t > T_{g1}$, but we will assume that each repetition is sufficiently far in the future that with discounting the future implementations can be ignored. We also assume logarithmic preferences throughout this section. The consumption function is as in (13) and the balanced budget assumption together with market clearing leads to (14). For notational convenience we will treat each policy implementation as announced at time $t = 1$ and implemented at time $t = T_{g1}$ and use $n = 1, 2, 3, \ldots$ to denote the number of the implementation.

We design the learning rule so that it allows both for learning within each implementation of the policy (as earlier) and also learning across implementations. When a temporary policy change is announced we assume that agents have an initial forecast of the interest rate path, based on previous implementations. For the first implementation we will assume that this initial path is simply the steady-state constant path, in line with the earlier sections of this paper (this is discussed further below). For the nth policy implementation we denote this initial forecast path of one-period interest rates by $\hat{r}^e(t, n)$, for $t = 2, 3, \ldots$. Learning across implementations is based on an updating of these initial forecast paths using data from the last implementation, according to the algorithm

$$\hat{r}^e(t, n + 1) = \hat{r}^e(t, n) + \kappa(r_{t+1}(n) - \hat{r}^e(t, n)),$$

for $t = 1, 2, .., T_H$, where $r_{t+1}(n)$ denotes the actual interest rate path from implementation n. Here T_H denotes a forecast horizon, beyond which we

assume $r_{t+1}^e(t, n) = \beta^{-1}$ for $t > T_H$, and $0 < \kappa \le 1$ is the gain parameter across policy implementations.

We continue to allow for learning within each implementation. That is, during implementation n the expectation path $\hat{r}^e(t, n)$ is updated over time according to previous forecast errors within the implementation. We denote the resulting expectations by $r_t^e(t+1+i, n)$, which denotes the expectation at t, during implementation n, of the one-period interest-rate factor r_{t+1+i}. We assume that the within implementation updating is given by

$$
\begin{aligned}
r_t^e(t+1+i, n) &= \hat{r}^e(t+1+i, n) + \Phi(t), \text{ where} \\
\Phi(t) &= \gamma \sum_{j=1}^{t-1} (1-\gamma)^j \left(r_{t-j}(n) - \hat{r}^e(t-j-1, n) \right),
\end{aligned}
$$

where $0 < \gamma < 1$ is, as before, the gain for learning within an implementation.[13] The future expected interest rate path, during implementation n at time t, is then given by $r_{t+1+i}^e = r_t^e(t+1+i, n)$ for $i = 1, 2, 3, \ldots$.

It is easy to verify that for $n = 1$, i.e. for the first implementation, this learning scheme is identical to the one used in earlier sections, provided $\hat{r}^e(t, 1) = \beta^{-1}$ for all t. Put differently, the learning scheme analyzed in earlier sections can be viewed as assuming a flat *prior* path for interest rate forecasts, set at the mean steady-state value. This is natural, since in the first implementation of the future policy change there is no historical experience to draw upon. However, there are circumstances in which an alternative initial path for $\hat{r}^e(t, 1)$ would be plausible. This would be natural if there had been previous policy changes which, while different, were believed to provide some qualitative information about the likely path of interest rates.

An important issue in practice is the choice of gains γ and κ. As discussed in Section 3.3, this issue is more properly addressed in a stochastic setting, since there is a trade-off between the tracking ability of higher gains to reduce bias and the filtering benefits of lower gains to reduce random variability. In particular, how much the expected interest rate path should be revised in light of data from the previous policy change, i.e. the best choice of κ, will depend on the relative importance of random shocks across implementations. We leave this important issue for future research, and here illustrate the implications of learning with repeated policy changes in a nonstochastic model with given κ.

---

[13]A more general formulation would relace $\Phi(t)$ by $\Phi(t)w(i)$, where $w(i) > 0$ is a non-negative "kernel", which is decreasing in i.

19

Given interest forecasts, we determine the actual interest rate path r_{t+1} from (14) for $1 \leq t \leq T_H$. We continue to assume that $\tau_{t+s}^e = \tau_{t+s}$ i.e., agents predict taxes correctly from knowledge of the government budget constraint. From (14) we can obtain, for $1 \leq t \leq T_H - 1$, the actual interest rate path $r_{t+1} = r_{t+1}(n)$, for policy implementation n, generated by expectations r_{t+i+1}^e, namely

$$r_{t+1}(y - g_t)\frac{\beta}{1 - \beta} = y - \tau_{t+1}^e + \Delta_t, \text{ where}$$

$$\Delta_t \equiv \sum_{s=2}^{\infty}(\prod_{i=1}^{s-1}(r_{t+i+1}^e)^{-1})(y - \tau_{t+s}^e).$$

For notational convenience, we are dropping the index for implementation n and writing $r_{t+1} = r_{t+1}(n)$. Note that in contrast to earlier sections, at each t agents are now forecasting a time-varying path for future interest rates, and this is reflected in the agents' behavior and summarized by Δ_t. As shown in the Appendix, the interest rate path r_{t+1}, for $1 \leq t \leq T_H - 1$, is given by

$$r_{t+1}(y - g_t)\frac{\beta}{1 - \beta} = y - g_{t+1} + \sum_{s=2}^{T_H-1}(\prod_{i=1}^{s-1}(r_{t+i+1}^e)^{-1})(y - g_{t+s}) +$$

$$(y - g_0)(\frac{\beta^{t-1}}{1 - \beta})(\prod_{i=1}^{T_H-t}(r_{t+i+1}^e)^{-1}) \tag{19}$$

We now present a numerical illustration. We report the difference in interest rates from the RE. The RE interest rate path is given by $r_{t+1} = \beta^{-1}$ for $1 \leq t \leq T_{g1} - 2$, and for all $t \geq T_{g1} + 1$ while at dates $T_{g1} - 1$ and T_{g1}, the interest rates are

$$r_{T_{g1}} = \beta^{-1}(\frac{y - g_1}{y - g_0}), r_{T_{g1+1}} = \beta^{-1}(\frac{y - g_0}{y - g_1}).$$

The results are shown in Figure 8. To generate the Figure, we used parameter values $\sigma = 1, \beta = 0.95, y = 6, \gamma = 1/200, \kappa = 1/20, g_0 = \tau_0 = 0.9, g_1 = \tau_1 = 1.1, T_{g1} = 20$, and $T_H = 40$.

As is clear from the figure, after a sufficient number of policy iterations, the interest rate forecasts do tend to converge to the interest rate dynamics under RE with the announced one period policy change in government spending. This result seems to be robust to changes in T_{g1}, g_0, and g_1. At
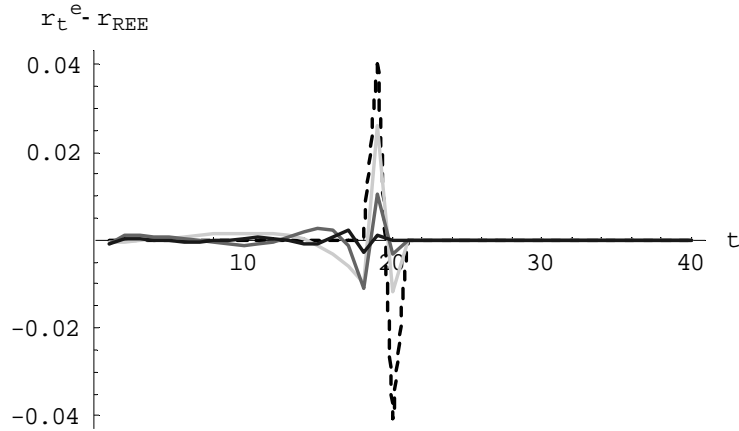
Figure 5: Difference between the RE values and the interest rate forecasts when the policy change is repeated. The differences are plotted for the second iterate (dashed line), the 25th iterate (solid light grey line), 50th iterate (solid medium grey line) and the 100th iterate (solid black line). Note the convergence towards the RE values as the number of iterations increases.

the same time it is apparent that convergence can be quite slow, especially when the gain parameters γ and κ are relatively small.[14]

On the one hand the results of this section are reassuring. Given repeated implementations of a preannounced policy change, we have shown how an adaptive learning rule will eventually arrive at RE forecasts for the interest rate path resulting from the policy change. On the other hand, our results also illustrate the extreme assumptions required for this result. We need repeated implementation of the same policy, namely a one-period balanced-budget increase in government spending, announced T_{g_1} periods in advance, with each repetition occurring after an interval of time sufficiently long for their discounted effects to be negligible. Repeated policy changes with variations in the duration of the temporary increase in government spending, or in the shape of government spending over the duration, or in the length of the announcement period T_{g_1}, would render learning across repetitions

---

[14]We remark, however, that instability can arise for larger gains. Stability depends jointly on the sizes of the two gain parameters and appears to involve a trade-off between the sizes of the two gains.

much more difficult. Repetitions that were not after long intervals could be handled, provided the interval period were constant, but variations in this interval would cause difficulties for learning. If changes in temporary government spending and the length of the announcement period were modeled as a stochastic process, asymptotic convergence to RE would require that this process be stationary and that agents fully understand the stochastic process underlying policy changes.

Realistically, many policy changes have unique features, so that the policy shift cannot plausibly be viewed as the realization of a stationary process. The conclusion we draw is the basic framework of this paper developed in Sections 2 and 3 provides a realistic way to model the impact of announced policy changes that are viewed as going beyond or fundamentally changing an established policy rule. In these circumstances, agents would naturally combine available structural information about future changes in key variables, such as tax rates, with adaptive learning for other variables such as interest rates. The current section shows how this framework can be extended to cases in which there is additional information from earlier policy changes considered relevant. If one or more earlier policy change are sufficiently similar that previous data can be used as a benchmark, the adaptive learning rule used by agents can incorporate this information using the procedures of the current section. If previous implementations have been quite different, but still have enough similarity to suggest useful qualitative information, this can be incorporated by initializing the within-implementation learning with a suitable subjectively adjusted *prior* for $\hat{r}^e(t, 1)$.

In the next two sections we illustrate how our methodology can be extended to more complicated economic set-ups. In examining these extensions we return to the Section 3 perspective of anticipated policy changes with unique features, in which any earlier policy changes are considered not to provide useful information about the likely course of interest rates and related endogenous variables.

5 Debt Financing

We now allow the government to issue debt in the form of one-period risk-free real bonds to finance its deficits. The government flow budget constraint (11) implies the following intertemporal budget constraint, assuming the transver-

sality condition (3) is satisfied.[15]

$$0 = r_s b_{s-1} + \sum_{j=1}^{\infty}(D^e_{s,s+j})^{-1}(g_{s+j} - \tau_{s+j}) + (g_s - \tau_s).$$

We now introduce a particular change in government policy that leads to debt financing.[16] Assume that $g_t = \bar{g}_0$ up to some future time $T_{p1} > 0$ and a different target level of spending $\bar{g}_1$ thereafter. The level of the lump-sum tax is assumed to change at some later time T_{p2} from a level δ_0 to a new level δ_1 such that this covers the new level of government spending $\bar{g}_1$ and the steady state interest burden on the accumulated debt $b_{T_{p2}-1}$. Thus,

$$\tau_t = \left\{ \begin{array}{l} \delta_0, \, t < T_{p2}, \\ \delta_1, \, t \geq T_{p2}, \end{array} \right. , \text{ where}$$

$$\delta_1 \equiv \bar{g}_1 + (\beta^{-1} - 1)b_{T_{p2}-1}$$

and T_{p2} is assumed to be bigger than T_{p1}. The actual level of government spending after T_{p2} is also adjusted to account for any deviations of the interest rate from its long-run steady state value. Thus,

$$g_t = \left\{ \begin{array}{l} \bar{g}_0, \, t < T_{p1} \\ \bar{g}_1, \, T_{p1} \leq t < T_{p2} \\ \bar{g}_1 - (r_t - \beta^{-1})b_{T_{p2}-1}, \, t \geq T_{p2}. \end{array} \right.$$

This policy change is announced to agents at the beginning of period 1 and they take it to be fully credible. Note that the level of government spending and the debt will stabilize in the long run as $r_t \to \beta^{-1}$.

By Ricardian equivalence, under rational expectations the dynamics of interest rates will be the same as in the balanced budget case. We take up the question of whether Ricardian equivalence will continue to hold under learning. In our fiscal set-up, private agents under learning need to forecast not only future interest rates but also the level of public debt at time $T_{p2} - 1$, since this determines subsequent taxes.

We assume that private agents forecast interest rates as in Section 3.3, and through time $T_{p2} - 1$ they assume taxes will be set at their announced

---

[15]As mentioned earlier, this will be true for the policies we consider.

[16]Naturally, one could consider various cases of debt finance and we focus on the chosen case for brevity and concreteness.

value of δ_0. Forecasts of taxes τ_t for $t \geq T_{p2}$, however, require forecasts of $b_{T_{p2}-1}$. Here we make the assumption that private agents understand the flow (per capita) budget identity (11), iterate it forward to $T_{p2} - 1$, and use the announced path of future government spending, together with their interest-rate forecasts, to estimate $b_{T_{p2}-1}$ and hence taxes after period T_{p2}.

The resulting consumption function and the solution for the equilibrium interest-rate path under learning are obtained in the Appendix. For the case of logarithmic preferences we have

$$r_{t+1}[\beta(1 - \beta)^{-1}(y - g_t) - (g_t - \tau_t + r_t b_{t-1})] = \frac{y}{1 - r^e(t)^{-1}} - \tilde{Q}(t),$$

where

$$\tilde{Q}(t) = \begin{cases} \frac{\delta_0 + (\delta_1^e(t) - \delta_0)(r^e(t))^{-(T_{p2}-1-t)}}{1 - r^e(t)^{-1}} & \text{for } 1 \leq t \leq T_{p2} - 1 \\ \delta_1(1 - r^e(t)^{-1})^{-1} & \text{for } t \geq T_{p2}. \end{cases}$$

Computation of r_{t+1} for $t \geq T_{p2}$ is straightforward. However, for $1 \leq t \leq T_{p2} - 1$, it is necessary to compute $\delta_1^e(t)$ and hence $b_{T_{p2}-1}^e(t)$ as detailed in the Appendix.

It is necessary to resort to numerical computations to analyze this case of debt financing. The dynamics of the difference in actual interest rates under learning with debt finance (r_{t+1}^{debt}) and the balanced-budget case (r_{t+1}^{bb}) are illustrated in Figure 6. In the case of debt finance, $\delta_0 = g_0$, $T_{p1} = 20$ (which was the value of T_p in the balanced budget case) and $T_{p2} = 30$. The initial value of debt is assumed to be zero and other initial conditions are as in the balanced budget case illustrated in Figures 1 and 2.

The interest rate dynamics with debt under learning are almost identical to the balanced budget case (note the scale of the vertical axis in Figure 3). Thus, there is *near Ricardian equivalence* despite bounded rationality and large deviations from fully rational expectations of the interest rates. Quantitatively, there are only very small differences, with interest rates slightly lower under debt financing before the actual implementation of the change in government spending and slightly higher interest rates under debt financing after the change. In other words, the interest rate is slightly more volatile under debt financing.

The central explanation for the near Ricardian equivalence is that the agents have been endowed with full knowledge of the structure of government financing including, in particular, the way taxes and government spending
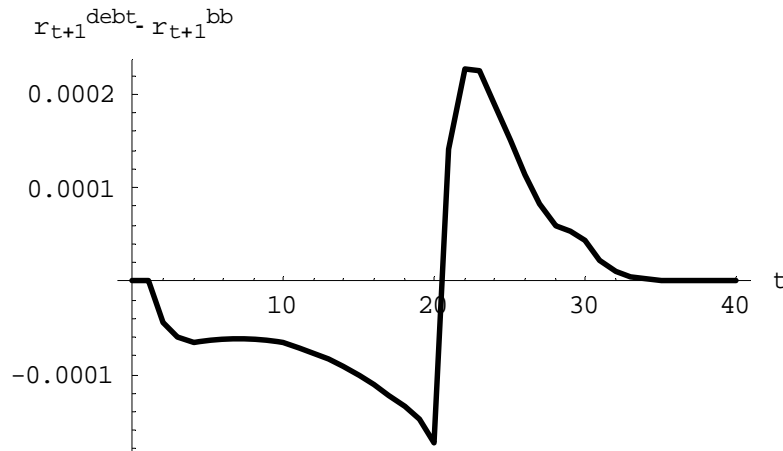
24

Figure 6: Difference between interest rates with debt financing and balanced budget, permanent policy change.

will be reconciled with the intertemporal budget constraint. Deviations of expectations from full rationality, even though these are quite large, as seen in Section 3.3, make only very small differences to the outcome with debt vs. balanced-budget financing.

6 The Ramsey Model

While the endowment economy model has been convenient for setting forth our approach to expectation formation and analyzing how its results differ from both rational expectations and purely adaptive learning, it should be clear that our approach can be applied in most macroeconomic frameworks. In this section we illustrate this point by showing how to extend the analysis to a model with capital, i.e. the discrete-time Ramsey model.

The basic economic framework is standard. The production function in intensive form is given by $y_t = f(k_t)$, where Y_t is output, K_t is capital, L_t is labour, $k_t = K_t/L_t$, $y_t = Y_t/L_t$, and $f(k) \equiv F(k, 1)$, where $F(K, L)$ is the underlying production function with constant returns to scale. For simplicity, we assume that labour is supplied inelastically and we normalize per capita labour supply to be equal to 1. There is no population and productivity

growth and no depreciation of the capital stock.

Profit maximization by firms entails that the gross interest rate r_t and the real wage rate w_t satisfy

$$r_t = 1 + f'(k_t), \tag{20}$$
$$w_t = f(k_t) - k_t f'(k_t). \tag{21}$$

Household's utility function is still given by (1) and the flow household budget constraint in per capita terms is

$$k_{t+1} = w_t + r_t k_t - c_t - \tau_t,$$

where τ_t denotes per capita lump-sum tax. For simplicity, we assume that the government budget is balanced, i.e. $\tau_t = g_t$, where g_t is real government spending on goods. The intertemporal budget constraint of the consumer is

$$0 = r_t k_t + \sum_{j=1}^{\infty} (D_{t,t+j}^e)^{-1} (w_{t+j}^e - c_{t+j}^e - \tau_{t+j}^e) + w_t - c_t - \tau_t$$

for given expectations $\{r_{t+j}^e\}, \{w_{t+j}^e\}, \{\tau_{t+j}^e\}$. The same formula holds under perfect foresight, with the equilibrium path substituted in place of the expectations. As before, we also substitute the iterated Euler equation in the form (6) to get

$$c_t \left[1 + \sum_{j=0}^{\infty} \beta^{j/\sigma} (D_{t,t+j}^e)^{\sigma^{-1}-1} \right] = r_t k_t + w_t - \tau_t + \sum_{j=1}^{\infty} (D_{t,t+j}^e)^{-1} (w_{t+j}^e - \tau_{t+j}^e), \tag{22}$$

which defines the consumption function

$$c_t = c \left(\{r_{t+j}^e\}, \{w_{t+j}^e\}, \{\tau_{t+j}^e\}, w_t, r_t k_t \right).$$

Market clearing determines k_{t+1} from

$$k_{t+1} = f(k_t) - c_t - g_t + k_t, \tag{23}$$

given the consumption function, the wage rate (21), the interest rate (20), a sequence $\{g_{t+s}\}$ together with the government budget constraint $\tau_{t+j} = g_{t+j}$, and expectations. We continue to assume that the consumers have full knowledge of the government budget constraint so that $\tau_{t+j}^e = \tau_{t+j}$.

As a benchmark, we recall the analysis of anticipated policy changes under perfect foresight.[17] The policy announcement considered is the same as in Section 3.1, i.e., at $t = 1$ it is announced that there will be a once-for-all increase in government spending from $g_t = g_0 = \tau_0$ when $t < T_p$ to $g_t = g_1 = \tau_1$ when $t \geq T_p$. Assuming that the economy is at a steady state at $t = 0$, the perfect foresight equilibrium leads to an instantaneous drop in consumption at $t = 1$ and a gradual decline thereafter such that at time T_p consumption is at the saddle path associated with the new steady state. After time T_p there is further monotonic decline to the new steady state in accordance with the saddle-point stable dynamics. The capital stock initially increases monotonically from $t = 1$ till period T_p before declining monotonically thereafter towards the new steady state. The dynamics of the interest and wage rates mirror the dynamics of capital in view of equations (20) and (21).

Returning to the analysis of learning, we consider how the dynamics of c_t and k_t evolve when agents must forecast the future paths $\left\{ r^e_{t+j} \right\}, \left\{ w^e_{t+j} \right\}$ as they have much less knowledge of the economy than under perfect foresight. We will assume that forecasts of r_{t+j} and w_{t+j} are done using steady state learning with constant gain, i.e. agents' forecasts are given by

$$
\begin{aligned}
r^e_{t+j}(t) &= r^e(t), & (24) \\
w^e_{t+j}(t) &= w^e(t) & (25)
\end{aligned}
$$

as in Section 3.2. This assumption is a plausible simplified approach in the current setting for the following reason.

More generally, agents' perceptions might take into account the hypothesis that r_t, w_t depend on the capital stock k_t in view of equations and (20)-(21) and that the capital stock evolves gradually, i.e. depends on its lagged value k_{t-1} as in (23). In other words, agents might estimate the parameters κ_r, φ_r, κ_w, φ_w, κ_k, φ_k of the linear relations $r_t = \kappa_r + \varphi_r k_t$, $w_t = \kappa_w + \varphi_w k_t$ and $k_t = \kappa_k + \varphi_k k_{t-1}$.[18] However, in the current nonstochastic setting, it is not possible to statistically estimate both the intercept and slope parameters of these relations because of emerging perfect multicollinearity as the economy converges to the steady state. In the nonstochastic setting agents can only

---

[17]The details are well explained in Ljungqvist and Sargent (2003), Chapter 11. See in particular Figure 11.3.1 and Section 11.6.

[18]Note that the linearized saddle-path dynamics of the RE equilibrium take this form for specific values of the parameters κ_r, φ_r, κ_w, φ_w, κ_k, φ_k.

estimate the long-run steady state values $\bar{r}$ and $\bar{w}$.[19] For a discussion of the subtle differences of learning in deterministic and stochastic models see Evans and Honkapohja (1998).

Coming back to the analysis, we also assume that current consumption by the agents is determined before the value of r_{t+1} is known, in other words $r^e_{t+1}(t) = r^e(t)$, and so consumption is thus fully determined by expectations and current variables w_t, τ_t and $r_t k_t$, which are predetermined.[20] Then equation (23) determines k_{t+1} (and hence r_{t+1}) as a residual. One could also adopt the alternative formulation whereby consumption depends on the actual value of r_{t+1} and in this case k_{t+1} and r_{t+1} are simultaneously determined by equations (23), (22) and (20).

In what follows, we assume that the production function is Cobb-Douglas, i.e. $f(k) = k^\alpha, 0 < \alpha < 1$, and that the utility function is logarithmic, i.e. $\sigma = 1$. Using the details in the Appendix, we have for $1 \le t < T_p$,

$$c_t = (1 - \beta)\left[k_t + k_t^\alpha + \frac{w^e(t)}{r^e(t) - 1} - \frac{\tau_0 r^e(t)}{r^e(t) - 1} + (\tau_0 - \tau_1)\frac{r^e(t)^{t-T_p}}{1 - r^e(t)^{-1}}\right] \quad (26)$$

and for $t \ge T_p$ we have

$$c_t = (1 - \beta)\left[k_t + k_t^\alpha + \frac{w^e(t)}{r^e(t) - 1} - \frac{\tau_1 r^e(t)}{r^e(t) - 1}\right]. \quad (27)$$

The dynamics of expectations are given by

$$
\begin{aligned}
r^e(t+1) &= r^e(t) + \gamma(r_t - r^e(t)), \\
w^e(t+1) &= w^e(t) + \gamma(w_t - w^e(t)).
\end{aligned}
$$

and the initial conditions are given by the "old" steady state. The Appendix shows that the jump in consumption in period 1 is

$$c_1 - c_0 = (\tau_0 - \tau_1)\beta^{T_p - 1}, \quad (28)$$

---

[19] In principle, they can also estimate the steady state value $\bar{k}$, which is not needed in the forecasting. Alternatively, if the intercepts are known, the slope coefficients can be estimated, see Section 4.8.1 of Evans and Honkapohja (2001).

[20] This assumption differs from what was assumed for the endowment economy in Section 3.2, where agents' consumption is a function of the actual r_{t+1}, which is determined as part of the temporary equilibrium. For the endowment economy the adopted assumption is the only natural formulation.

whereas the jump in consumption under perfect foresight is

$$(\tau_0 - \tau_1)(\lambda_1^{-1})^{T_p - 1},$$

where λ_1 is the unstable root in the perfect foresight dynamics.[21] It can be shown that $\lambda_1^{-1} < \beta$, so that the fall in consumption is greater under learning than under perfect foresight dynamics. This is confirmed by the numerical simulations below. More generally, our aim is to compare the paths of c_t and k_t under steady state learning and under perfect foresight using numerical simulations.

The computation of the paths under learning is straightforward and we can just generalize the procedures in earlier settings. For perfect foresight dynamics the new (though well-known) feature is that, with durable capital, the dynamics are computed using saddle-path stability conditions of the perfect foresight paths. The computation of perfect foresight path can either be done directly from the nonlinear setting using shooting-algorithms or using a linear approximation to the model to compute the backward-forward solution. See Ljungqvist and Sargent (2003), Chapter 11 for a discussion of these approaches. For our analysis we opt for the linearization technique.[22]

We now compare the dynamics of consumption and capital under steady state learning and under perfect foresight. The following parameter values are assumed in the simulations: $T_p = 20$, $\gamma = 0.1$, $\beta = 0.95$, $\alpha = 0.3$, $g_0 = 0.9$ and $g_1 = 1.1$. Figure 7 presents the results for consumption. The solid curve gives the dynamics of consumption under learning and the dashed curve under perfect foresight. (The horizontal dashed line indicates the new steady state.) The perfect foresight path verifies the features asserted earlier: an initial negative jump and a gradual decline afterwards towards the new steady state. Under learning the initial negative jump is bigger than under under perfect foresight and the path is non-monotonic; consumers want to "correct" the big negative jump and consumption starts to increase, though eventually after the actual policy change consumption starts to decline. Another notable feature in the simulation is that the learning path actually undershoots the new steady state value.

---

[21] Analytical details of the Ramsey model in the perfect foresight case are standard, see Ljungqvist and Sargent (2003), Chapter 11, and are available on request.

[22] The shooting algorithm involves fixing an arbitrary long but finite horizon at which the numerically calculated path is close to the steady state. The approximate PF paths will eventually diverge because of numerical approximations about the switch point and the saddle path at the time of the actual policy change.
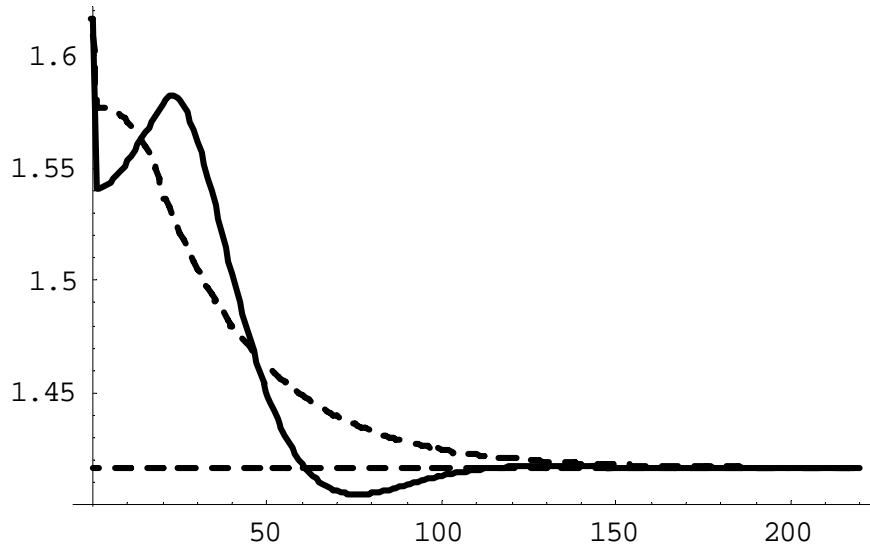
Figure 7: Consumption dynamics under learning (solid curve) and perfect foresight (dashed curve). The dashed straight line is the new steady state for consumption.

The next figure shows the dynamics of the capital stock under learning and under perfect foresight, the solid and dashed curves, respectively. (The horizontal dashed line again indicates the new steady state.) For a good part the dynamics of the adjustment of capital are qualitatively similar under learning and perfect foresight. A notable difference is that under learning the capital stock seems to undershoot the new steady state value, whereas the decline towards the new steady state is monotonic under perfect foresight.

7 Concluding Remarks

This paper has developed a methodology for studying the effects of anticipated policy changes on the dynamics of learning paths. We have examined this issue using models that are well behaved in the sense that adaptive or statistical learning rules converge asymptotically to rational expectations in stationary environments. A disadvantage of purely adaptive learning is that it is not forward-looking. Our framework overcomes this limitation by assuming that agents combine structural information about future policy with
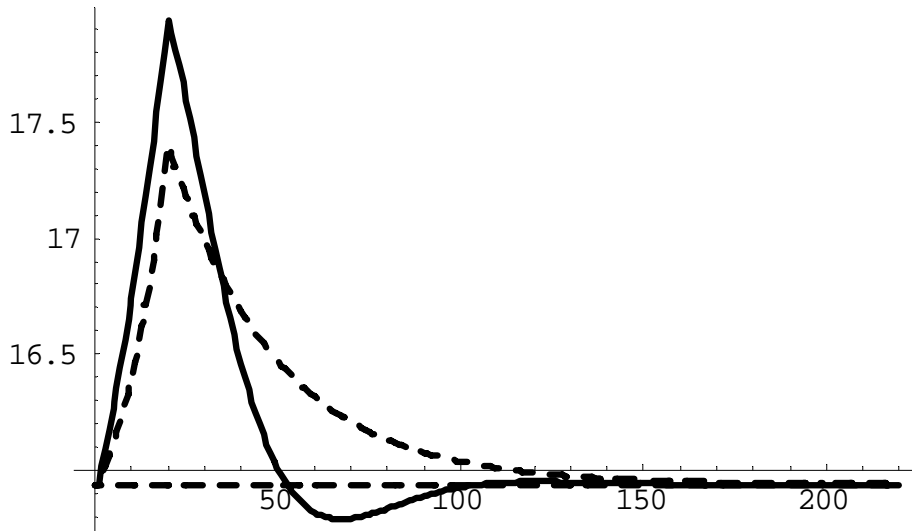
Figure 8: Dynamics of k_t under learning (solid curve) and perfect foresight (dashed curve). The dashed straight line is the new steady state of k_t.

adaptive learning for the remaining variables they are required to forecast.

To develop our analysis we employed simple models of fiscal policy. In these frameworks it is plausible to assume that credible plans for the paths of government policy variables are structural information that agents can incorporate into their forecasting procedures. Likewise, government budget constraints are plausible candidates for structural information that agents exploit in making their forecasts.

The cases we have examined are natural formulations of the general idea that, depending on the context, there may be structural knowledge that is available and used even when agents have much more imperfect information about the determinants of aggregate variables than presumed by the rational expectations hypothesis. Cases of partial structural knowledge lead to modifications of both the way agents make forecasts and their behavioral rules when learning. In our context, agents incorporated the future changes in government fiscal variables into behavioral rules with long horizons, combining the structural tax forecasts with adaptive learning rules for forecasting interest rates and wages.

Our analysis focused on a few very stylized economies in order to study

31

these ideas in a simple setting. Quite clearly, similar ideas can be exploited in more complex models. Some direct extensions of the models in this paper would also be important. In particular, extending the modelling to settings with random shocks has priority for future research. This is because formulation of statistical learning rules is more straightforward in stochastic settings. An example of this is the Ramsey model in which a stochastic framework would allow agents to forecast interest rates and wages by estimating and updating standard statistical models. Randomness also arises naturally in connection with policy announcements. One case is uncertainty about the date at which the policy change is implemented. A second case of uncertainty about policy is the imperfect credibility of future policy announcements.

Appendix: Derivations and Proofs

A   Endowment Economy

A.1   Permanent changes in policy

Under learning we have

$$
\begin{aligned}
S_D(t) &= \sum_{j=1}^{\infty} \beta^{j/\sigma} \{ r_{t+1} r^e(t)^{j-1} \}^{\sigma^{-1}-1} = r_{t+1}^{\sigma^{-1}-1} \beta^{\sigma^{-1}} \sum_{j=1}^{\infty} \{ \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1} \}^{j-1} \\
&= r_{t+1}^{\sigma^{-1}-1} \beta^{\sigma^{-1}} \frac{1}{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}. \tag{29}
\end{aligned}
$$

For $t = 1, 2, ..., T_p - 1$, we evaluate $S_N(t)$ as

$$
\begin{aligned}
S_N(t) &\equiv y \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1} - \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1} \tau_{t+j}^e \\
&= \frac{y r_{t+1}^{-1}}{1 - r^e(t)^{-1}} - \left\{ \tau_0 r_{t+1}^{-1} \sum_{j=1}^{T-1} (r^e(t))^{1-j} + \tau_1 r_{t+1}^{-1} \sum_{j=T}^{\infty} (r^e(t))^{1-j} \right\} \tag{30} \\
&= \frac{y r_{t+1}^{-1}}{1 - r^e(t)^{-1}} - \tau_0 r_{t+1}^{-1} \frac{1 - (r^e(t))^{1-T}}{1 - (r^e(t))^{-1}} - \tau_1 r_{t+1}^{-1} \frac{(r^e(t))^{1-T}}{1 - (r^e(t))^{-1}},
\end{aligned}
$$

where T is the number of periods until T_p, i.e., $T = T_p - t$. Note that the computation of $S_N(t)$ above assumes that $\tau_t = \tau_0$ when $t < T_p$ and $\tau_t = \tau_1$ when $t \geq T_p$. Hence, for $1 \leq t \leq T_p - 1$,

$$
S_N(t) \equiv \frac{y - \tau_0 + (\tau_0 - \tau_1)(r^e(t))^{t+1-T_p}}{r_{t+1}(1 - (r^e(t))^{-1})}.
$$

We compute r_{t+1} from the market clearing condition for $1 \leq t \leq T_p - 1$,

$$
r_{t+1}^{\sigma^{-1}-1} r_{t+1} = \frac{y - \tau_0 + (\tau_0 - \tau_1)(r^e(t))^{t+1-T_p}}{(1 - (r^e(t))^{-1})(y - \tau_0)\beta^{\sigma^{-1}}} (1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}),
$$

which yields the solution for r_{t+1}, when $1 \leq t \leq T_p - 1$, as

$$
r_{t+1} = \left[\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})} \left\{ 1 - \frac{(\tau_1 - \tau_0)}{(y - \tau_0)} (r^e(t))^{t+1-T_p} \right\} \right]^{\sigma}.
$$

On the other hand, for $t \geq T_p$, we have

$$S_N(t) \equiv \frac{y - \tau_1}{r_{t+1}(1 - (r^e(t))^{-1})}$$

and the corresponding r_{t+1} for, $t \geq T_p$, is

$$r_{t+1} = [\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})}]^{\sigma}.$$

Proof of Proposition 1:

First consider periods $1, \ldots, T_p - 1$. We have

$$r^e(t+1) = G(r^e(1), \gamma, t) \text{ where}$$

$$G(r^e(1), \gamma, t) = r^e(1) + \gamma \sum_{i=1}^{t} \left(H(r^e(i), i) - r^e(i) \right).$$

Initially we have $r^e(1) = \beta^{-1}$. Clearly $G(\beta^{-1}, 0, t) = \beta^{-1}$. From equation (16) we also have $H(\beta^{-1}, i) < \beta^{-1}$ for all $1 \leq i \leq T_p - 1$. This implies that

$$\frac{\partial}{\partial \gamma} G(\beta^{-1}, 0, t) = \sum_{i=1}^{t} \left(H(\beta^{-1}, i) - \beta^{-1} \right) < 0.$$

This proves (a) for periods $t = 1, \ldots, T_p - 1$.

Next, consider period T_p:

$$r^e(T_p + 1) = (1 - \gamma)r^e(T_p) + \gamma r_{T_p+1}$$
$$= (1 - \gamma)r^e(T_p) + \gamma \frac{1 - \beta}{\beta} \frac{r^e(T_p)}{r^e(T_p) - 1},$$

where we have used the formula for r_{T_p+1} in (16). Introducing temporary notation $x = r^e(T_p)$ we have $r^e(T_p + 1) < \beta^{-1}$ iff

$$(1 - \gamma)x + \frac{1 - \beta}{\beta} \frac{x}{x - 1} < \beta^{-1}.$$

As $x \geq 1$ by rearranging the terms we have that $r^e(T_p + 1) < \beta^{-1}$ iff

$$(1 - \gamma)x(x - 1) + \gamma \frac{1 - \beta}{\beta} x < \beta^{-1}(x - 1), \text{ i.e.}$$
$$(1 - \gamma)(x - \beta^{-1})(x - (1 - \gamma)^{-1}) < 0.$$

34

Thus, $r^e(T_p + 1) < \beta^{-1}$ iff $(1-\gamma)^{-1} < x < \beta^{-1}$, which is true as, by the earlier arguments,

$$x = r^e(T_p) = G(\beta^{-1}, \gamma, T_p - 1) < \beta^{-1}$$

for $\gamma > 0$ sufficiently small. For $t \geq T_p + 2$ we can repeat the argument. To complete the proof for part (a) we note that for $r^e(t) < \beta^{-1}$ we have

$$\frac{r^e(t+1)}{r^e(t)} = (1-\gamma) + \gamma \frac{1-\beta}{\beta} \frac{1}{r^e(t) - 1} > 1.$$

This implies that $r^e(t+1)$ is a strictly increasing sequence and clearly its limit must be β^{-1}.

To prove part (b) we first note that $H(\beta^{-1}, t) < \beta^{-1}$ and hence by continuity $r_{t+1} = H(r^e(t), t) < \beta^{-1}$ for $t = 2, \ldots, T_p - 1$ when γ is sufficiently small. For $t \geq T_p$ we have

$$r_{t+1} = \frac{1-\beta}{\beta} \frac{1}{1 - (r^e(t))^{-1}} > \beta^{-1}$$

as $r^e(t) < \beta^{-1}$ holds for all $t \geq 2$ by part (a). Using the same formulae, we also have $\lim_{t\to\infty} r_{t+1} = \beta^{-1}$ as $\lim_{t\to\infty} r^e(t) = \beta^{-1}$ by part (a) of the Proposition. Q.E.D.

A.2 Temporary policy changes

The consumption function is as in (13) and $S_D(t)$ is as in (15). For $1 \leq t \leq T_{g1} - 1$, we have

$$S_N(t) \equiv y \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1} - \sum_{j=1}^{\infty} (D_{t,t+j}^e(t))^{-1} \tau_{t+j}^e = \frac{yr_{t+1}^{-1}}{1 - r^e(t)^{-1}}$$

$$- \tau_0 r_{t+1}^{-1} \sum_{j=1}^{T_1-1} (r^e(t))^{1-j} - \tau_1 r_{t+1}^{-1} \sum_{j=T_1}^{T_2-1} (r^e(t))^{1-j} - \tau_0 r_{t+1}^{-1} \sum_{j=T_2}^{\infty} (r^e(t))^{1-j}$$

$$= \frac{yr_{t+1}^{-1}}{1 - r^e(t)^{-1}} - \tau_0 r_{t+1}^{-1} \frac{1 - (r^e(t))^{1-T_1}}{1 - (r^e(t))^{-1}} -$$

$$\tau_1 r_{t+1}^{-1} (r^e(t))^{1-T_1} \frac{1 - (r^e(t))^{T_1-T_2}}{1 - (r^e(t))^{-1}} - \tau_0 r_{t+1}^{-1} \frac{(r^e(t))^{1-T_2}}{1 - (r^e(t))^{-1}}.$$

where $T_1 = T_{g1} - t$, and $T_2 = T_{g2} - t$, are the number of periods remaining until T_{g1} and T_{g2} respectively.

For $T_{g1} \leq t \leq T_{g2} - 1$,we have

$$S_N(t) = \frac{yr_{t+1}^{-1}}{1 - r^e(t)^{-1}} - \tau_1 r_{t+1}^{-1} \frac{1 - (r^e(t))^{1-T_2}}{1 - (r^e(t))^{-1}} - \tau_0 r_{t+1}^{-1} (r^e(t))^{1-T_2} \frac{1}{1 - (r^e(t))^{-1}}$$

and finally for $t \geq T_{g2}$,we get

$$S_N(t) = (y - \tau_0) r_{t+1}^{-1} (1 - r^e(t)^{-1})^{-1}.$$

As in Appendix A.1, we can solve for the interest rate r_{t+1}. For $1 \leq t \leq T_{g1} - 1$, this is given by

$$r_{t+1} = [(\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})})\{1 - \frac{(\tau_1 - \tau_0)}{(y - \tau_0)}(r^e(t))^{t+1-T_{g1}}(1 - r^e(t)^{T_{g1}-T_{g2}})\}]^\sigma.$$

When $T_{g1} \leq t \leq T_{g2} - 1$, the interest rate r_{t+1} is given by

$$r_{t+1} = [(\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})})\{1 + (\frac{\tau_1 - \tau_0}{y - \tau_1})r^e(t)^{t+1-T_{g2}}\}]^\sigma,$$

and finally for $t \geq T_{g2}$, we have

$$r_{t+1} = [\frac{1 - \beta^{\sigma^{-1}} r^e(t)^{\sigma^{-1}-1}}{\beta^{\sigma^{-1}}(1 - (r^e(t))^{-1})}]^\sigma.$$

Proof of Proposition 2:

To prove part (a), we first note that for $2 \leq t \leq T_{g1}$

$$r^e(t + 1) = \tilde{G}(r^e(1), \gamma, t), \text{ where}$$

$$\tilde{G}(r^e(1), \gamma, t) = r^e(1) + \gamma \sum_{i=1}^{t} (F(r^e(i), i) - r^e(i)),$$

using equation (18). In this expression

$$F(\beta^{-1}, t) = \beta^{-1}[1 - a_1 \beta^{T_{g1}-t-1}(1 - \beta^{T_{g2}-T_{g1}}) < \beta^{-1},$$

where $a_1 = \frac{\tau_1 - \tau_0}{y - \tau_0}$, and

$$\frac{\partial}{\partial \gamma} \tilde{G}(\beta^{-1}, 0, t) < 0.$$

36

As in Proposition 1, it follows that $r^e(t) < \beta^{-1}$ for $\gamma > 0$ sufficiently small.

To prove the second assertion, we compute

$$\left.\frac{\partial r^e(T_{g2})}{\partial \gamma}\right|_{\gamma=0} = \sum_{i=1}^{T_{g2}-1} [r_{i+1} - r^e(i)]\Big|_{\gamma=0} \tag{31}$$

where $r^e(i) = \beta^{-1}$ at $\gamma = 0$,

$$r_{i+1}|_{\gamma=0} = \beta^{-1}[1 - a_1(\beta^{-1})^{i+1-T_{g1}}(1 - \beta^{T_{g2}-T_{g1}})$$

for $i = 1, \ldots, T_{g1} - 1$, and

$$r_{T_{g1}+i}\big|_{\gamma=0} = \beta^{-1}[1 + a_2\beta^{T_{g2}-T_{g1}-i}]$$

for $i = 1, \ldots, T_{g2} - T_{g1}$. Here $a_2 = \frac{\tau_1-\tau_0}{y-\tau_1}$. Computing the sum (31) we get

$$\left.\frac{\partial r^e(T_{g2})}{\partial \gamma}\right|_{\gamma=0} = -\beta^{-1}a_1\frac{(1 - \beta^{T_{g2}-T_{g1}})(1 - \beta^{T_{g1}-1})}{1 - \beta}$$
$$+ \beta^{-1}a_2\beta^{T_{g2}-T_{g1}-1}\frac{1 - (\beta^{-1})^{T_{g2}-T_{g1}}}{1 - \beta^{-1}}.$$

Since $\tau_1 > \tau_0$ we have $a_1 < a_2$ and therefore to show that $\left.\frac{\partial r^e(T_{g2})}{\partial \gamma}\right|_{\gamma=0} > 0$ it is sufficient that

$$\beta^{T_{g2}-T_{g1}-1}\frac{\beta}{\beta - 1}(1 - (\beta^{-1})^{T_{g2}-T_{g1}}) - \frac{(1 - \beta^{T_{g2}-T_{g1}})(1 - \beta^{T_{g1}-1})}{1 - \beta}$$
$$= (1 - \beta)^{-1}[\beta^{T_{g1}-1} - \beta^{T_{g2}-1}] > 0,$$

which holds since $\beta < 1$ and $T_{g2} > T_{g1}$. The limit result follows an argument analogous to the one used in the proof of Proposition 1. This completes the proof of part (a).

To prove part (b), it is seen from above that $F(\beta^{-1}, t)$ is bounded strictly below β^{-1} for $t < T_{g1}$. It then follows by continuity that $F(r^e(t), t) < \beta^{-1}$ for all $\gamma > 0$ sufficiently small. For $t = T_{g1}$ we have

$$r_{T_{g1}+1} = \frac{1 - \beta}{\beta(1 - (r^e(t))^{-1})}[1 + a_2(r^e(t))^{T_{g1}+1-T_{g2}}] > \beta^{-1}$$

37

since $r^e(T_{g1}) < \beta^{-1}$. For $t \geq T_{g2}$ we get

$$r_{t+1} = \frac{1-\beta}{\beta(1-(r^e(t))^{-1})} < \beta^{-1}$$

since $r^e(t) > \beta^{-1}$. Finally, the limit result follows from the convergence of expectations to the steady state value. Q.E.D.

B Repeated Policy Learning

We have

$$\Delta_t \equiv \sum_{s=2}^{T_H-1} (\prod_{i=1}^{s-1} (r^e_{t+i+1})^{-1})(y - \tau^e_{t+s}) + \sum_{s=T_H}^{\infty} (\prod_{i=1}^{s-1} (r^e_{t+i+1})^{-1})(y - \tau^e_{t+s}).$$

The second summation in Δ_t can be simplified as follows

$$\sum_{s=T_H}^{\infty} (\prod_{i=1}^{s-1} (r^e_{t+i+1})^{-1})(y - \tau^e_{t+s}) = (y - g_0) \sum_{s=T_H}^{\infty} (\prod_{i=1}^{s-1} (r^e_{t+i+1})^{-1}) =$$

$$(y - g_0) \sum_{s=T_H}^{\infty} (\prod_{i=1}^{T_H-t} (r^e_{t+i+1})^{-1})(\prod_{i=T_H-t+1}^{s-1} \beta) =$$

$$(y - g_0)(\prod_{i=1}^{T_H-t} (r^e_{t+i+1})^{-1})(\sum_{s=T_H}^{\infty} \beta^{s+t-T_H-1}) = (y - g_0)(\prod_{i=1}^{T_H-t} (r^e_{t+i+1})^{-1})(\frac{\beta^{t-1}}{1-\beta}).$$

Combining with market clearing yields equation (19).

C Debt financing

We begin by noting that the actual evolution of debt is as follows:

$$\begin{aligned}
b_t &= \bar{g}_0 - \delta_0 + r_t b_{t-1}, \text{ for } t = 1, \ldots, T_{p1} - 1 \\
b_t &= \bar{g}_1 - \delta_0 + r_t b_{t-1}, \text{ for } t = T_{p1}, \ldots, T_{p2} - 1 \\
b_t &= b_{T_{p2}-1}, \text{ for } t \geq T_{p2}.
\end{aligned}$$

Note that the level of debt is stabilized at $b_{T_{p2}-1}$ from date T_{p2} onwards.

Let $b_s^e(t)$ denote the time t forecast of b_s. As bonds are pre-determined the one-step ahead expectations of bonds satisfy $b_{t+1}^e(t) = b_{t+1}$. To forecast the evolution of debt at time t we have for $t = 1, \ldots, T_{p1} - 2$:

$$
\begin{aligned}
b_{t+1}^e(t) &= b_{t+1} = \bar{g}_0 - \delta_0 + r_{t+1} b_t, \\
b_{t+2}^e(t) &= \bar{g}_0 - \delta_0 + r^e(t) b_{t+1}, \\
b_s^e(t) &= \bar{g}_0 - \delta_0 + r^e(t) b_{s-1}^e(t), \text{ for } s = t+3, \ldots, T_{p1} - 1, \\
b_s^e(t) &= \bar{g}_1 - \delta_0 + r^e(t) b_{s-1}^e(t), \text{ for } s = T_{p1}, \ldots, T_{p2} - 1,
\end{aligned}
$$

noting that any date t, both b_t and b_{t+1} are known.

At $T_{p1} - 1 \leq t \leq T_{p2} - 3$ we have

$$
\begin{aligned}
b_{t+1}^e(t) &= \bar{g}_1 - \delta_0 + r_{t+1} b_t, \\
b_{t+2}^e(t) &= \bar{g}_1 - \delta_0 + r^e(t) b_{t+1}, \\
b_s^e(t) &= \bar{g}_1 - \delta_0 + r^e(t) b_{s-1}^e(t), \text{ for } s = t+3, \ldots, T_{p2} - 1.
\end{aligned}
$$

At $t = T_{p2} - 2$ we have

$$
b_{T_{p2}-1}^e(T_{p2} - 2) = \bar{g}_1 - \delta_0 + r_{T_{p2}-1} b_{T_{p2}-2}.
$$

Using $b_{T_{p2}-1}^e(t)$, agents forecast taxes by

$$
\begin{aligned}
\tau_s^e(t) &= \bar{g}_1 + (\beta^{-1} - 1) b_{T_{p2}-1}^e(t), \text{ for } s \geq T_{p2} \text{ and } t = 1, \ldots, T_{p2} - 2. \\
\tau_s^e(t) &= \tau_s = \bar{g}_1 + (\beta^{-1} - 1) b_{T_{p2}-1}, \text{ for } s \geq T_{p2} \text{ and } t \geq T_{p2} - 1.
\end{aligned}
$$

One can use recursive substitution to solve further for the above formulas. For example, when $1 \leq t \leq T_{p1} - 2$, we get

$$
\begin{aligned}
b_{t+3}^e(t) &= \bar{g}_0 - \delta_0 + r^e(t) b_{t+2}^e(t) = \bar{g}_0 - \delta_0 + r^e(t)(\bar{g}_0 - \delta_0 + r^e(t) b_{t+1}) \\
&= (\bar{g}_0 - \delta_0)(1 + r^e(t)) + (r^e(t))^2 b_{t+1}, \\
b_{t+4}^e(t) &= \bar{g}_0 - \delta_0 + r^e(t) b_{t+3}^e(t) \\
&= (\bar{g}_0 - \delta_0)(1 + r^e(t) + (r^e(t))^2) + (r^e(t))^3 b_{t+1}.
\end{aligned}
$$

Continuing in this fashion, we inductively obtain for all $1 \leq t \leq T_{p1} - 2$,

$$
\begin{aligned}
b_{T_{p1}-1}^e(t) &= (\bar{g}_0 - \delta_0)(1 + r^e(t) + \ldots + (r^e(t))^{T_{p1}-3-t}) + (r^e(t))^{T_{p1}-2-t} b_{t+1} \\
&= (\bar{g}_0 - \delta_0) \frac{1 - (r^e(t))^{T_{p1}-2-t}}{1 - r^e(t)} + (r^e(t))^{T_{p1}-2-t}(\bar{g}_0 - \delta_0 + r_{t+1} b_t).
\end{aligned}
$$

Thereafter continuing in the same fashion for $1 \le t \le T_{p2} - 2$, we have

$$
\begin{aligned}
b^e_{T_{p1}}(t) &= (\bar{g}_1 - \delta_0) + r^e(t) b^e_{T_{p1}-1}(t), \\
b^e_{T_{p1}+1}(t) &= (\bar{g}_1 - \delta_0)(1 + r^e(t)) + (r^e(t))^2 b^e_{T_{p1}-1}(t), \\
b^e_{T_{p2}-1}(t) &= (\bar{g}_1 - \delta_0)(1 + r^e(t) + ... + (r^e(t))^{T_{p2}-1-T_{p1}}) + (r^e(t))^{T_{p2}-T_{p1}} b^e_{T_{p1}-1}(t).
\end{aligned}
$$

The last line simplifies as below, so that for $1 \le t \le T_{p1} - 2$, we finally have

$$
b^e_{T_{p2}-1}(t) = (\bar{g}_1 - \delta_0) \frac{1 - (r^e(t))^{T_{p2}-T_{p1}}}{1 - r^e(t)} + (r^e(t))^{T_{p2}-T_{p1}} b^e_{T_{p1}-1}(t).
$$

We now move to the case when $T_{p1} - 1 \le t \le T_{p2} - 3$. Using the same approach as before we get

$$
\begin{aligned}
b^e_{t+2}(t) &= \bar{g}_0 - \delta_0 + r^e(t) b_{t+1}, \\
b^e_{t+3}(t) &= (\bar{g}_1 - \delta_0)(1 + r^e(t)) + (r^e(t))^2 b_{t+1}, \\
b^e_{T_{p2}-1}(t) &= (\bar{g}_1 - \delta_0)(1 + r^e(t) + ... + (r^e(t))^{T_{p2}-t-3}) + (r^e(t))^{T_{p2}-t-2} b_{t+1} \\
&= (\bar{g}_1 - \delta_0) \frac{1 - (r^e(t))^{T_{p2}-t-2}}{1 - r^e(t)} + (r^e(t))^{T_{p2}-t-2} b_{t+1},
\end{aligned}
$$

where $b_{t+1} = \bar{g}_1 - \delta_0 + r_{t+1} b_t$.

Finally, for $t = T_{p2} - 2$, we have

$$
b^e_{T_{p2}-1}(T_{p2} - 2) = \bar{g}_1 - \delta_0 + r_{T_{p2}-1} b_{T_{p2}-2} = b_{T_{p2}-1},
$$

which is a known quantity at $T_{p2} - 2$.

Given the tax and interest rate forecasts $\tau^e_{t+j}(t)$ and $r^e(t)$, consumption demand at time t is given by (8), which may be rewritten as

$$
c_t = \frac{r_t b_{t-1} + y - \tau_t + \sum_{j=1}^{\infty}(D^e_{t,t+j}(t))^{-1} y - Q(t)}{1 + S_D(t)}; \text{ where} \tag{32}
$$

$$
Q(t) \equiv \sum_{j=1}^{\infty}(D^e_{t,t+j}(t))^{-1} \tau^e_{t+j}(t). \tag{33}
$$

$S_D(t)$ can be simplified to (29), and the infinite sum in (32) simplifies to the first term in (30).

The tax forecasts in equation (33), $Q(t)$, are given by

1. For $t \geq T_{p2} - 1$,

$$\tau_{t+j}^e(t) = \tau_{t+j} = \bar{g}_1 + (\beta^{-1} - 1)b_{T_{p2}-1} \equiv \delta_1$$

for all $j \geq 1$ since $b_{T_{p2}-1}$ is known at time $t \geq T_{p2} - 1$. Hence,

$$Q(t) = \delta_1 r_{t+1}^{-1}(1 - r^e(t)^{-1})^{-1}.$$

2. For $1 \leq t \leq T_{p2} - 2$,

$$Q(t) \;=\; \sum_{j=1}^{T_{p2}-1-t} (D_{t,t+j}^e(t))^{-1}\delta_0 + \sum_{j=T_{p2}-t}^{\infty} (D_{t,t+j}^e(t))^{-1}\delta_1^e(t),$$

$$\delta_1^e(t) \;=\; \bar{g}_1 + (\beta^{-1} - 1)b_{T_{p2}-1}^e(t),$$

where $\delta_1^e(t) = \tau_s^e(t)$ for $s \geq T_{p2}$, so that

$$Q(t) \;=\; \delta_0 r_{t+1}^{-1}\Big(\frac{1 - (r^e(t))^{-(T_{p2}-1-t)}}{1 - r^e(t)^{-1}}\Big) + \delta_1^e(t)r_{t+1}^{-1}\Big(\frac{(r^e(t))^{-(T_{p2}-1-t)}}{1 - r^e(t)^{-1}}\Big)$$

$$\equiv\; r_{t+1}^{-1}\tilde{Q}(t), \text{ where } \tilde{Q}(t) \equiv \frac{\delta_0 + (\delta_1^e(t) - \delta_0)(r^e(t))^{-(T_{p2}-1-t)}}{1 - r^e(t)^{-1}}.$$

Combining market clearing with (32) yields the interest rate solution from

$$r_{t+1}^{\sigma^{-1}}\beta^{\sigma^{-1}}(1-\beta^{\sigma^{-1}}r^e(t)^{\sigma^{-1}-1})^{-1}(y-g_t) = (g_t-\tau_t+r_t b_{t-1})r_{t+1}+\frac{y}{1-r^e(t)^{-1}}-\tilde{Q}(t).$$

D   Ramsey model under learning

We need to compute the infinite sums involved in the right hand side of (22) to evaluate consumption. Using (24)-(25), we get

$$\sum_{j=1}^{\infty}(D_{t,t+j}^e)^{-1}(w_{t+j}^e - \tau_{t+j}^e) \;=\; SW - ST_1, \text{ where}$$

$$SW \;=\; \frac{w^e(t)}{r^e(t) - 1}, t \geq 1,$$

$$ST_1 \;=\; \begin{cases} \frac{\tau_0}{r^e(t)-1} + (\tau_1 - \tau_0)\frac{r^e(t)^{t-T_p}}{1-r^e(t)^{-1}}, & 1 \leq t < T_p, \\ \frac{\tau_1}{r^e(t)-1}, & t \geq T_p. \end{cases}$$

Furthermore, we assume that $f(k) = k^\alpha, 0 < \alpha < 1$, and that $\sigma = 1$. The infinite sum on the left hand side of (22), $S_D(t)$, equals $\beta(1-\beta)^{-1}$ from (29) when $\sigma = 1$. Hence, (26) follows. In addition, we have

$$
\begin{aligned}
r_t &= 1 + \alpha k_t^{\alpha-1}, w_t = (1-\alpha)k_t^\alpha, \\
\bar{k} &= \left(\frac{1-\beta}{\alpha\beta}\right)^{\frac{1}{\alpha-1}},
\end{aligned}
$$

where $\bar{k}$ is the steady state value of capital. For given level of government spending g, steady state value of consumption is $\bar{c} = (\bar{k})^\alpha - g$.

We now compute the jump in consumption under learning dynamics. As before, the dynamics start at the steady state so that $k_0 = k_1 = \bar{k}$, $r_1 = 1 + \alpha\bar{k}^{\alpha-1}$, and $w_1 = (1-\alpha)\bar{k}^\alpha$. In addition, the initial estimates satisfy $r^e(0) = 1 + \alpha\bar{k}^{\alpha-1}$ and $w^e(0) = (1-\alpha)\bar{k}^\alpha$. These also imply that $r^e(1) = r^e(0) = 1 + \alpha\bar{k}^{\alpha-1}$ and $w^e(1) = w^e(0) = (1-\alpha)\bar{k}^\alpha$.

At $t = 1$, the value of consumption from (26) is given by

$$
\begin{aligned}
c_1 &= (1-\beta)[\bar{k} + \bar{k}^\alpha + \frac{(1-\alpha)\bar{k}^\alpha}{\alpha\bar{k}^{\alpha-1}}] - (1-\beta)\tau_0(\frac{1+\alpha\bar{k}^{\alpha-1}}{\alpha\bar{k}^{\alpha-1}}) + \\
&\quad (1-\beta)(\tau_0 - \tau_1)\frac{(1+\alpha\bar{k}^{\alpha-1})^{2-T_p}}{\alpha\bar{k}^{\alpha-1}}.
\end{aligned}
$$

Using the steady state value of $\bar{k}$, it is easy to show that $1 + \alpha\bar{k}^{\alpha-1} = \beta^{-1}$ and $\alpha\bar{k}^{\alpha-1} = \beta^{-1}(1-\beta)$. With these values, one can show that the first term in the expression above for c_1 simplifies to $\bar{k}^\alpha$, the second term to $-\tau_0$, and the third term to $(\tau_0 - \tau_1)\beta^{T_p-1}$. This finally proves that under learning we have

$$
c_1 = \bar{k}^\alpha - \tau_0 + (\tau_0 - \tau_1)\beta^{T_p-1} = c_0 + (\tau_0 - \tau_1)\beta^{T_p-1},
$$

so that the jump in consumption is given by (28).

References

BLANCHARD, O. J., AND S. FISCHER (1989): *Lectures on Macroeconomics.* MIT Press, Cambridge, Mass.

BRANCH, W., AND G. W. EVANS (2006): "A Simple Recursive Forecasting Model," *Economic Letters*, 91, 158–166.

BRUNNER, K., AND A. MELTZER (eds.) (1976): *The Phillips Curve and the Labor Market. Carnegie-Rochester Conferences in Public Policy, Vol. 1.* North-Holland, Amsterdam.

BULLARD, J., AND J. DUFFY (2001): "Learning and Excess Volatility," *Macroeconomic Dynamics*, 5, 272–302.

CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

EVANS, G. W., AND S. HONKAPOHJA (1998): "Economic Dynamics with Learning: New Stability Results," *Review of Economic Studies*, 65, 23–44.

——— (2001): *Learning and Expectations in Macroeconomics.* Princeton University Press, Princeton, New Jersey.

EVANS, G. W., S. HONKAPOHJA, AND R. MARIMON (2001): "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules," *Macroeconomic Dynamics*, 5, 1–31.

EVANS, G. W., S. HONKAPOHJA, AND P. ROMER (1998): "Growth Cycles," *American Economic Review*, 88, 495–515.

EVANS, G. W., AND G. RAMEY (1992): "Expectation Calculation and Macroeconomic Dynamics," *American Economic Review*, 82, 207–224.

——— (1998): "Calculation, Adaptation and Rational Expectations," *Macroeconomic Dynamics*, 2, 156–182.

——— (2006): "Adaptive Expectations, Underparameterization and the Lucas Critique," *Journal of Monetary Economics*, 53, 249–264.

GIANNITSAROU, C. (2006): "Supply-Side Reforms and Learning Dynamics," *Journal of Monetary Economics*, 53, 291–309.

LJUNGQVIST, L., AND T. J. SARGENT (2003): *Recursive Macroeconomic Theory, Second edition*. MIT Press, Cambridge, MA.

LUCAS, JR., R. E. (1976): "Econometric Policy Evaluation: a Critique," in Brunner and Meltzer (1976), pp. 19–46.

MARCET, A., AND J. P. NICOLINI (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93, 1476–1498.

ORPHANIDES, A., AND J. C. WILLIAMS (2005): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations," *Journal of Economic Dynamics and Control*, 29, 1927–1950.

PRESTON, B. (2005): "Learning about Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, 1, 81–126.

———— (2006): "Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy," *Journal of Monetary Economics*, 53, 507–535.

SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

SARGENT, T. J., AND N. WALLACE (1973a): "Rational Expectations and the Dynamics of Hyperinflation," *International Economic Review*, 14, 429–472.

———— (1973b): "The Stability of Models of Money and Growth with Perfect Foresight," *Econometrica*, 41, 1043–1048.