

Anticipation, learning and welfare: the case of distortionary taxation

Emanuel Gasteiger^{a,*}, Shoujian Zhang^b

^a *Instituto Universitário de Lisboa (ISCTE - IUL), BRU - IUL, Department of Economics, Av.^a das Forças Armadas, 1649-026 Lisboa, Portugal*

^b *School of Economics & Finance, University of St Andrews, Castlecliffe, The Scores, St Andrews, Fife KY169AL, UK*

A B S T R A C T

We study the impact of anticipated fiscal policy changes in a Ramsey economy where agents form long-horizon expectations using adaptive learning. We extend the existing framework by introducing distortionary taxes as well as elastic labour supply, which makes agents' decisions non-predetermined but more realistic. We detect that the dynamic responses to anticipated tax changes under learning have oscillatory behaviour that can be interpreted as self-fulfilling waves of optimism and pessimism emerging from systematic forecast errors. Moreover, we demonstrate that these waves can have important implications for the welfare consequences of fiscal reforms.

JEL classification: E32, E62, D84

Keywords: Fiscal policy, Adaptive learning, Oscillations, Welfare

*Corresponding author. Av.^a das Forças Armadas, 1649-026 Lisboa, Portugal, Tel.: +351-926-785229, emanuel.gasteiger@iscte.pt

Email addresses: emanuel.gasteiger@iscte.pt (Emanuel Gasteiger), shoujian.zhang@st-andrews.ac.uk (Shoujian Zhang)

1. Introduction

Nowadays, fiscal policy is usually accompanied by legislation and implementation lags. These lags create a non-negligible span of time between the announcement and effective date of fiscal policy changes. This gives economic agents the opportunity to anticipate tax changes. The economic literature denotes this aspect of fiscal policy *anticipated fiscal policy*.¹

When agents anticipate, their resulting actions may to some extent depend on the way they form expectations about the future. The standard assumption of expectations in economics is perfect-foresight / rational expectations (RE). This assumption might be questioned on the grounds of its unrealistically strong restrictions. One prominent deviation of RE that imposes weaker requirements on the agent's information set when making his decisions, is the learning literature (see Evans and Honkapohja, 2001, for the foundations of this approach). The main idea is that agents form expectations about the future evolution of contemporaneously unobservable variables by engaging in a kind of statistical inference, when making their economic choices.

Although the learning approach has gained significant popularity in some areas of macroeconomics, (anticipated) fiscal policy has, until recently, been neglected. A pioneering contribution to the study of anticipated fiscal policy under learning has been made by Evans et al. (2009). They demonstrate the adaptive constant gain learning approach in several deterministic economic environments including the popular Ramsey model. In this set-up, it is assumed that agents un-

¹Recently Leeper (2009, p.24ff.) has listed empirical evidence for anticipated fiscal policy/fiscal foresight and re-emphasized the relevance of expectations for sound fiscal policy. Furthermore, Leeper et al. (2013) is another good example of empirical evidence of fiscal foresight. Therein they also demonstrate the challenges for econometricians that aim to quantify the impact of fiscal policy actions and at the same time account adequately for fiscal foresight.

derstand the structure of government financing but have to forecast factor-prices on decentralized markets. Their key result is that for an anticipated balanced-budget permanent tax change the impact effects on key variables under learning are similar to the perfect foresight case, but the transition paths are remarkably different from the latter. Their result, at least with regard to the volatility of key variables' time paths may not come as a surprise. It is well known that constant gain learning causes excess volatility (see Evans and Honkapohja, 2001, p.49).²

Building on the contribution of Evans et al. (2009), we aim to generalize their analysis of anticipated fiscal policy under learning by studying an economy featuring distortionary taxes and elastic labour supply. Thus, our theoretical key contribution is to derive the dynamic paths of key variables for anticipated permanent changes in distortionary taxes in the prominent Ramsey model.³

Note that there are fundamental differences between lump-sum taxation and distortionary taxes such as labour income, capital income, or consumption tax.⁴ Furthermore, the assumption of elastic labour supply implies that endogenous variables such as factor prices as well as employment and consumption are not predetermined as in Evans et al. (2009, p.943ff.), but determined simultaneously. Thus, one can expect that changes in distortionary taxes may yield dynamics fundamentally different from the ones in the lump-sum tax case.

One of our main results supports this hypothesis. When we assume that agents use adaptive learning rules to forecast factor prices, our model predicts oscillatory dynamic responses to pre-announced permanent tax changes. More-

²In subsequent work, Evans et al. (2012) focus on Ricardian equivalence in the basic Ramsey model with anticipated fiscal policy under learning. Most important, Evans et al. (2012, p.1267ff.) formally prove that the assumption of RE is not necessary for the Ricardian equivalence result.

³Another extension is to consider stochastic set-ups. Recently, Mitra et al. (2013) have pioneered the study of anticipated lump-sum tax changes under learning in the RBC model.

⁴Find the details in Ljungqvist and Sargent (2012, p.375ff.).

over, the case of distortionary taxation is particularly different with regard to the effects on impact and the volatility throughout the transition period. The source of the oscillations are persistent systematic forecast errors, that lead agents to incorrectly anticipate the effects of pre-announced tax changes. Thus, the oscillations can be interpreted as self-fulfilling waves of optimism and pessimism.

Farmer (1999, p.141ff.) and others have argued that these waves are a feature of US data and can be replicated by modified RBC models under RE. Mitra et al. (2013) reproduce the waves as a consequence of fiscal policy changes under the assumption of adaptive learning in the standard RBC model. We demonstrate that these waves exist even in the deterministic version of this model.

Confronted with this result, we then ask, to what extent the striking differences in the dynamics between perfect foresight and learning, affect the welfare consequences of a pre-announced tax reform in the presence of several tax instruments. For this purpose, we make use of the welfare measure proposed by Lucas (1990) and also applied by Cooley and Hansen (1992) (for discrete time), which takes into account the whole transition path between the pre- and post-reform steady-states associated with initial and changed tax rates.

This welfare analysis is our second key contribution. It links the learning literature to the part of the public finance literature that is concerned with the welfare consequences of tax reforms.⁵ However, the existing literature evaluates and ranks various distortionary tax reforms according to their welfare consequences under perfect foresight, but does not consider the case of long-horizon learning. We fill this gap and illustrate that tax reforms designed to improve welfare, do

⁵See Chamley (1981) for an example of a comparative statics analysis, Judd (1987) for differences in unanticipated and anticipated changes in factor taxes, or Cooley and Hansen (1992) for a study in a stochastic set-up. Moreover, Garcia-Milà et al. (2010) have recently conducted research on welfare consequences of fiscal policy reforms in a heterogeneous agents model.

so to a much lower extent under learning compared to perfect foresight. Thus, the learning perspective on tax reforms provides fundamental different insights for benevolent policy makers.

Furthermore, our welfare analysis can be related to the Euler equation learning perspective on surprise supply-side reforms provided by Giannitsarou (2006). Despite several differences in the approach, our results in the welfare analysis are somewhat similar to the ones of Giannitsarou (2006). We will discuss the particulars below.

The remainder of the paper is organized as follows. In Section 2 we outline the economic model. We use Section 3 to detail our approach of learning and to derive the equations that govern the dynamics under learning. Section 4 provides a simple intuitive example of lump-sum tax changes. This section also contains a sensitivity analysis for some structural parameters. In Section 5 we focus on distortionary taxation and present a numerical welfare analysis of an exemplary tax reform. Section 6 concludes and points out directions for further research.

2. The model

Our economy is a version of the Ramsey economy (see Ljungqvist and Sargent, 2012, p.375ff.). k_t evolves according to the economy-wide resource constraint

$$k_{t+1} = F(k_t, n_t) - c_t - g_t + (1 - \delta)k_t, \quad (1)$$

where $F(k_t, n_t)$ is the economy's production function showing that the firm sector uses the stock of capital k_t and labour n_t as inputs to produce the single good of the economy (see Section 2.2 for the details). Output can be purchased by households, c_t , or the government, g_t , or added to k_t , which is assumed to depreciate

at a constant rate δ .

2.1. Households

We assume a continuum of households, where we normalize the size of the economy to unity and each household faces the problem

$$\max_{c_t, n_t, b_{t+1}} E_t^* \left\{ \sum_{t=0}^{\infty} \beta^t [\log(c_t) + \eta \log(\bar{L} - n_t) + \phi \log(g_t)] \right\} \text{ s.t. (2)}$$

$$k_{t+1} + \frac{b_{t+1}}{R_t} + (1 + \tau_t^c)c_t = (1 - \tau_t^l)w_t n_t + (1 - \tau_t^k)r_t k_t + (1 - \delta)k_t + b_t - \tau_t + \pi_t, (3)$$

where all variables are in per capita terms. Thus, b_{t+1} is the level of government debt holdings chosen in period t . Furthermore, r_t is the rental rate of capital and R_t is the gross real interest rate in period t . Next, τ_t^\bullet denotes a distortionary tax either on consumption, labour income or capital income.⁶ The real wage in period t is given by w_t and $(\bar{L} - n_t)$ denotes leisure. In consequence, n_t is labour supply of the household.⁷ τ_t is a per capita lump-sum tax, and $\pi_t = 0$ is the profit under perfect competition among firms. Furthermore, the parameter $\eta \geq 0$ measures the dis-utility of labour supply, ϕ measures the degree of substitution between private consumption and government spending, see Ambler and Paquet (1996), and β is the common discount rate. $E_t^*\{\bullet\}$ denotes subjective period t expectations for future values of variables.⁸

⁶We use the symbol \bullet as a placeholder throughout our analysis.

⁷We assume a log-separable utility function to obtain a closed form solution for the households' consumption function below. This approach is also used in the standard RBC literature, see e.g. Braun (1994), Christiano and Eichenbaum (1992), or Prescott (1986). One could also consider more general utility functions, but this would then in turn require a local approximation around the steady-state to derive the consumption function, see for example Mitra et al. (2013).

⁸Households apply this operator, if they do not have perfect foresight. This assumption is commonly used in the learning literature. Furthermore, note that we abstract from aggregate uncertainty, i.e. we conduct our analysis in a deterministic economy. Thus, if households do not have perfect foresight, their expectations are so-called point expectations, i.e. agents base their

Given this set-up, the household Euler condition, the usual no-arbitrage condition for capital and bonds, and the consumption leisure trade-off are respectively given by

$$\bar{c}_t^{-1} = \beta R_t E_t^* \left\{ c_{t+1}^{-1} \frac{(1 + \tau_t^c)}{(1 + \tau_{t+1}^c)} \right\}, \quad (4)$$

$$R_t = [(1 - \delta) + (1 - E_t^* \{ \tau_{t+1}^k \}) E_t^* \{ r_{t+1} \}], \quad (5)$$

$$n_t = \bar{L} - \frac{\eta(1 + \tau_t^c)c_t}{(1 - \tau_t^l)w_t}. \quad (6)$$

2.2. Firms

We assume a unit continuum of firms who compete perfectly. Each firm in each period t rents capital at given price r_t and labour at given price w_t . It produces the numeraire good with constant returns to scale production function

$$y_t = F(k_t, n_t) = A k_t^\alpha n_t^{1-\alpha}, \quad (7)$$

where $\alpha \in (0, 1)$. The optimal firm behaviour requires that

$$r_t \stackrel{!}{=} \frac{\partial y_t}{\partial k_t} = A \alpha k_t^{\alpha-1} n_t^{1-\alpha}, \quad (8)$$

$$w_t \stackrel{!}{=} \frac{\partial y_t}{\partial n_t} = A(1 - \alpha) k_t^\alpha n_t^{-\alpha}, \quad (9)$$

i.e. each production factor earns its marginal product. Finally, we have the per capita national income identity $\pi_t = y_t - r_t k_t - w_t n_t = 0$.

economic choices on the mean of their expectations, see Evans and Honkapohja (2001, p.61). In Section 3.1 below we outline our concept of learning. An important aspect of this concept is that forecasts of single variables are independent of each other. In consequence, we can assume that for any two variables X and Y it is true that $E_t^* \{XY\} = E_t^* \{X\} E_t^* \{Y\}$ holds.

2.3. Government

The government finances its expenses on goods and debt repayment by tax revenues and the issuance of new bonds in each period t ,

$$g_t + b_t = \tau_t^c c_t + \tau_t^l w_t n_t + \tau_t^k r_t k_t + \tau_t + \frac{b_{t+1}}{R_t}.$$

For the remainder, we will assume that the government operates a balanced-budget rule in each period t , thus tax revenues will fully cover expenses such that bonds are in zero net supply as a direct consequence. Thus, in each period t the government sets $g_t, \tau_t^c, \tau_t^l, \tau_t^k$ and τ_t constrained by

$$g_t = \tau_t^c c_t + \tau_t^l w_t n_t + \tau_t^k r_t k_t + \tau_t. \quad (10)$$

3. Learning and learning dynamics

3.1. Learning

The concept of learning applied herein was elaborated first in Evans et al. (2009, p.943ff.). For completeness we restate the crucial assumptions. Under learning, households are supposed to know the entire history of endogenous variables. They observe the current period value of exogenous variables and they know the state variables. Furthermore, they know the structure of the economy with regard to the fiscal policy sector. Agents understand the implications of any pre-announced policy change for the government budget constraint. They are also convinced that the intertemporal government budget constraint will always hold (see Evans et al., 2009, p.944). Notably, we assume decentralized markets for labour and capital, where agents are not in possession of perfect foresight. Actual factor prices are not observable. Thus, agents forecast factor prices $r_{t+j}^e(t)$

and $w_{t+j}^e(t)$ for $j \geq 1$, by a constant gain steady-state adaptive learning rule⁹

$$z_{t+j}^e(t) = z^e(t) = z^e(t-1) + \gamma(z_{t-1} - z^e(t-1)), \quad (11)$$

where $0 < \gamma \leq 1$ is the gain parameter.¹⁰ Our choice of this specific learning rule is motivated by two well known arguments in the learning literature. First, as Evans and Honkapohja (2001, p.332) outline, choosing a constant gain learning rule is the appropriate choice for agents, when they are aware of structural change, as in such a learning rule agents discount past data exponentially. Note that rule (11) is equivalent to $z^e(t) = \gamma \sum_{i=0}^{\infty} (1-\gamma)^i z_{t-i-1}$.¹¹ Second, the timing of the learning rule, i.e. that agents' update in period t uses data up to period $t-1$, is chosen in order to avoid simultaneity between $r^e(t)$ and r_t as well as $w^e(t)$ and w_t (see for example Evans and Honkapohja, 2001, p.51). Simultaneity in this context means that agents' expectations affect current values of aggregate endogenous variables and vice versa, which could introduce strategic behaviour.

Such a learning set-up yields a sequence of so-called *temporary equilibria*, which consist of sequences of (planned) time paths for all endogenous variables. The sequences satisfy the learning rule above, the expectation history, household and firm optimality conditions, the government budget constraint, and the economy-wide resource constraint given the exogenous variables as well as the current stock of capital in each period. These plans are revisited and potentially

⁹Here we apply the same short-hand notation as Evans et al. (2009). Thus for any variable say z , its period t expected future value in period $t+j$ derived by a learning rule may either be denoted $E_t^*\{z_{t+j}\}$ or equivalently $z_{t+j}^e(t)$. An additional notation we introduce is $z_{t+j}^p(t)$, which denotes the agent's planned choice of the variable z in period $t+j$ based on expected values formed via the learning rule (11) in period t .

¹⁰The gain parameter measures the responsiveness of the forecast to new observations, see Evans and Honkapohja (2001, p.18). Be aware that in our model the gain parameter is exogenous. See Branch and Evans (2007) for an example where agents choose the gain parameter.

¹¹The learning rule is similar to an exponential smoothing method.

altered in each period after expectations have been updated.¹²

3.2. Learning dynamics with distortionary taxation

Now, we derive the dynamic paths under learning in presence of multiple types of taxes, i.e. $\tau_t^c, \tau_t^l, \tau_t^k \in [0, 1]$ and $\tau_t \neq 0$ for all t . Equations (4)-(5) yield

$$c_t^{-1} = \beta (c_{t+1}^p(t))^{-1} \left[\frac{(1 + \tau_t^c)}{(1 + \tau_{t+1}^{c,e}(t))} \right] [(1 - \delta) + (1 - \tau_{t+1}^{k,e}(t))r_{t+1}^e(t)]$$

and forward substitution of this expression yields

$$c_{t+j}^p(t) = \beta^j D_{t,t+j}^{k,e}(t) \left[\frac{(1 + \tau_t^c)}{(1 + \tau_{t+j}^{c,e}(t))} \right] c_t, \quad (12)$$

where we define $D_{t,t+j}^{k,e}(t) \equiv \prod_{i=1}^j [(1 - \delta) + (1 - \tau_{t+i}^{k,e}(t))r_{t+i}^e(t)]$. One can think of this term as “expectations of the interest rate factor $D_{t,t+j}$ at time t ” (see Evans et al., 2009, p.933). Furthermore, notice that the consumption leisure trade-off is now given by (6). If agents plan to satisfy an adequate transversality condition

$$\lim_{T \rightarrow \infty} \left(D_{t,t+T}^{k,e}(t) \right)^{-1} k_{t+T+1}^p(t) = 0, \quad (13)$$

the inter-temporal budget constraint of a consumer is

$$(1 + \tau_t^c)c_t + \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} (1 + \tau_{t+j}^{c,e}(t))c_{t+j}^p(t) = [(1 - \delta) + (1 - \tau_t^k)r_t]k_t \\ + (1 - \tau_t^l)w_t n_t - \tau_t + \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} [(1 - \tau_{t+j}^{l,e}(t))w_{t+j}^e(t)n_{t+j}^p(t) - \tau_{t+j}^e(t)],$$

¹²As discussed in Evans et al. (2012, p.1267), our set-up can be interpreted as an “Anticipated-Utility” framework developed by Kreps (1998) and elaborated in an adaptive learning context by Preston (2005).

which by the virtue of (12) as well as (6) yields the consumption function

$$c_t = \frac{(1 - \beta)}{(1 + \eta)(1 + \tau_t^c)} [(1 - \delta) + (1 - \tau_t^k)r_t]k_t + (1 - \tau_t^l)w_t\bar{L} - \tau_t + SW_2 - ST_2 - ST_3, \text{ where} \quad (14)$$

$$SW_2 = \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} w_{t+j}^e(t)\bar{L}, \quad (15)$$

$$ST_2 = \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} \tau_{t+j}^{l,e}(t) w_{t+j}^e(t)\bar{L}, \text{ and} \quad (16)$$

$$ST_3 = \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} \tau_{t+j}^e(t) \quad (17)$$

are the expected present value of labour income, the labour income taxes and lump-sum taxes in turn.

Some deliberations regarding the policy experiment under investigation are required. We are looking at a scenario of a credible permanent (simultaneous) change in (some) taxes announced at the outset of period $t = 1$ and effective from period $t = T_p$ onwards. The dynamics under perfect foresight are standard.¹³ Under learning, given that households know the future path of taxes, they can explicitly calculate SW_2 , ST_2 , and ST_3 .¹⁴ Given a calibration of the structural parameters, we can then compute the dynamic responses for c_t and the other endogenous variables to an anticipated tax reform.

4. A simple intuitive example: lump-sum tax increase

We would like to illustrate the applied methodology for the simple case of lump-sum taxation ($\tau_t^c = \tau_t^l = \tau_t^k = 0$) for three reasons: first, the rather complex

¹³Ljungqvist and Sargent (2012, p.375ff.) illustrate the analytical derivations and numerical simulation alternatives for the perfect foresight case. We will simply make use of the Dynare toolbox throughout all calculations to compute dynamics under perfect foresight. Note that this toolbox employs linearization methods.

¹⁴See Appendices A, B, and C for details.

consumption function (14)-(17) under learning simplifies in this case and facilitates to develop some intuition; second, we want to illustrate the consequences of the introduction of elastic labour supply compared to the case of inelastic labour supply as assumed in Evans et al. (2009, p.943ff.) and its effect on the dynamic paths of the key variables such as c_t and k_t , given their calibration (see below); third, below in Subsection 4.2, we aim to present a sensitivity analysis for the very basic version of the model to deepen the intuition for the impact of the learning parameters on the dynamics.

For an anticipated lump-sum tax change under learning (14) simplifies to

$$c_t = \frac{(1 - \beta)}{(1 + \eta)} \{[(1 - \delta) + r_t]k_t + w_t \bar{L} - \tau_0 + SW_1 - ST_1\}, \quad (18)$$

where the expected present value of labour income is given by

$$SW_1 = \sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} w_{t+j}^e \bar{L} = \frac{w^e(t) \bar{L}}{r^e(t) - \delta}, \quad (19)$$

and the expected present value of lump-sum taxes is given by $ST_1 =$

$$\sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} \tau_{t+j}^e = \begin{cases} \frac{\tau_0}{r^e(t) - \delta} + (\tau_1 - \tau_0) \frac{[(1-\delta) + r^e(t)]^{t-T_p}}{1 - [(1-\delta) + r^e(t)]^{-1}} & \text{for } 1 \leq t < T_p \\ \frac{\tau_1}{r^e(t) - \delta} & \text{for } t \geq T_p. \end{cases} \quad (20)$$

4.1. Inelastic labour supply vs. elastic labour supply

We believe that it is of importance to use a model that features elastic labour supply in order to study the implications of fiscal policy reforms adequately. Completely inelastic labour supply is a quite unrealistic assumption itself and at least some moderately elastic labour supply should be considered. Moreover, inelastic labour supply implies that agents' choices of current period endogenous

variables are in fact predetermined as is pointed out in Evans et al. (2009, p.944). In order to illustrate differences in the dynamics of endogenous variables based on the assumption of inelastic and elastic labour supply, we return to the simulation exercise of Evans et al. (2009, p.943ff.).¹⁵ The calibration for this subsection is given by $A = 1.00$, $\alpha = 0.33$, $\beta = 0.95$, $\phi = 0.00$, $\delta = 0.00$, $T_p = 20$, and $\gamma = 0.10$.¹⁶

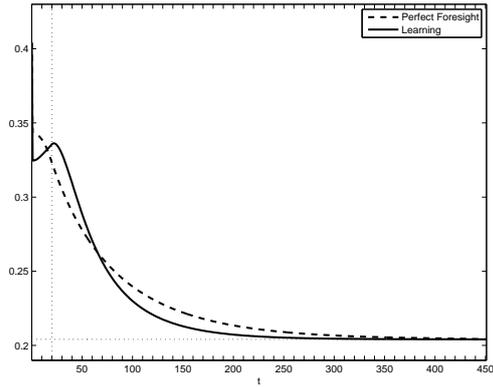
The policy experiment considered in Evans et al. (2009, p.943ff.) is a permanent increase in government purchases from $g_0 = \tau_0 = 0.9$ to $g_1 = \tau_1 = 1.1$ that is announced credibly in period $t = 1$ and will be effective from period $T_p = 20$ onwards. It is assumed that the economy is in steady-state in period $t = 0$. Simulations in Evans et al. (2009, p.943ff.) for consumption and capital are recalculated (with $\eta = 0$, $\bar{L} = 0.5182$) and displayed in Figs. 1a and 1b below. Next, Figs. 1c and 1d exhibit the dynamics for elastic labour supply (with $\eta = 2.00$, $\bar{L} = 1.00$) such that $n_0 = 0.5182$ and $g_0 = 0.9$.

Two distinct features emerge from Fig. 1. First, when we compare the dynamic paths of c_t , as well as k_t , under perfect foresight and learning, they are qualitatively different from each other no matter with or without elastic labour supply. Therefore, it may be quite important to consider learning when evaluating fiscal policies, as learning is a more realistic assumption of human behaviour from our point of view.¹⁷ Second, obviously the learning paths in Figs. 1a and

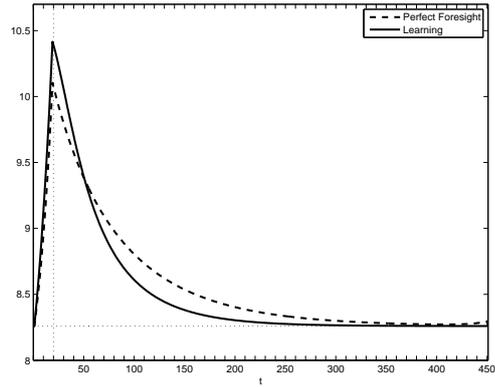
¹⁵Note that $\tau_t^c = \tau_t^l = \tau_t^k = \delta = 0$ and $\eta = 0$ imply that $n_t = \bar{L}$ (i.e. inelastic labour supply) for all t . Therefore, we are exactly in the same scenario as in Evans et al. (2009, p.943ff.).

¹⁶Note that we do not fully agree with the calibration of Evans et al. (2009), but we will stick to their calibration in this subsection to keep our results comparable. The basic reason for this disagreement is the combination of parameters $\beta = 0.95$ and $T_p = 20$. These parameter choices imply that a government, which in reality is usually in charge of a legislation period of four to six years, may announce a tax policy change that will be effective in 20 years' time. From our perception of political execution and our confidence in fiscal policy makers' ability to commit, this appears to be unrealistic in most cases.

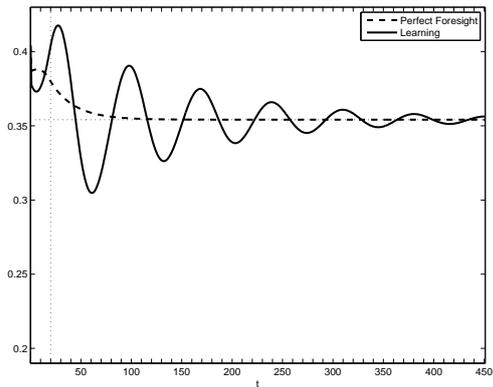
¹⁷This is the core message of Evans et al. (2009).



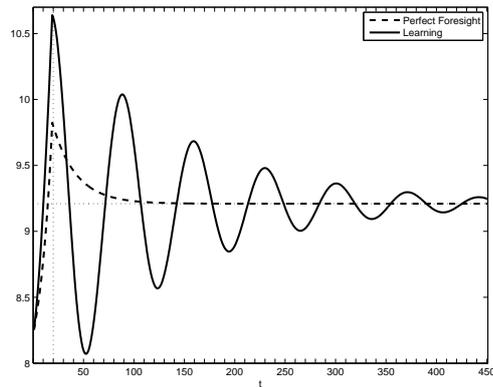
a. Consumption with inelastic n_t



b. Capital with inelastic n_t



c. Consumption with elastic n_t



d. Capital with elastic n_t

Fig. 1. Consumption (a) and capital (b) dynamics under learning (solid curve) and perfect foresight (dashed curve) with inelastic labour supply as in Evans et al. (2009, p.943ff.) as well as consumption (c) and capital (d) dynamics under learning (solid curve) and perfect foresight (dashed curve) with elastic labour supply. The dotted horizontal line indicates the (new) steady-state, the dotted vertical line indicates period T_p .

1b for inelastic labour supply are strikingly different to the ones under elastic labour supply in Figs. 1c and 1d. In particular, elastic labour supply yields more volatility in the time paths of c_t and k_t (as well as other variables in the model) compared to the inelastic labour supply case. In fact, the variables oscillate around their steady-state until they converge to it.

In order to develop some intuition for the transition dynamics illustrated in Fig. 1, we start with the familiar case of *perfect foresight*.¹⁸ Given *inelastic labour supply*, agents respond to the credible pre-announced tax increase with a decrease in c_t on impact. The policy change means a negative wealth effect to the agents, as expected future taxes increase. Due to their perfect foresight, they precisely quantify this negative wealth effect. In the subsequent period, this increases k_t beyond its steady-state equilibrium value. Out of the steady-state equilibrium, the transitional dynamics lead to a further decrease in c_t and a temporary investment boom that increases k_t throughout the pre-implementation period. The rise in k_t affects marginal products and increases w_t , decreases r_t , and increases y_t . Once the policy change becomes effective in T_p , the levels of c_t and k_t are consistent with the new saddle-path and thereafter decrease monotonically to the new steady-state equilibrium. The evolution of w_t , r_t , and y_t is reversed and they converge monotonically to their respective steady-state values.

Also under *learning*, agents anticipate a negative wealth effect, and decrease c_t on impact. However, they do so to a lower extent, as they are not able to precisely quantify the negative wealth effect. The reason is that their expectations about factor prices are predetermined point estimates. In fact, on impact they project the initial steady-state values into the entire future. Thereby they are too pessimistic about $w^e(t)$ and too optimistic about $r^e(t)$. The latter induces

¹⁸Our intuition is developed along the lines of Mitra et al. (2013) and the references therein.

them to underestimate the present value of future taxes, as one can see from (20). However, the negative impact effect on c_t causes an investment boom that yields an increase in k_t during the pre-implementation period. A higher k_t also yields a higher level of y_t . Moreover, via marginal products w_t increases and r_t decreases.

Now, these changes in factor prices, in turn affect c_t via two distinct channels, as (18) makes clear. First, there is a direct effect on c_t . A higher w_t can increase c_t and a lower r_t can decrease c_t . However, the net effect in our case is positive. Second, there is a self-fulfilling channel. The changes in actual factor prices trigger a learning process due to systematic forecast errors. Rule (11) leads to upward revisions of $w^e(t)$ and downward revisions of $r^e(t)$. The consumption function (18) makes clear that this leads to an increase in the sum of expected future labour income and taxes, where the increase in the former outweighs the increase in the latter. Thus, c_t also increases via this channel in the pre-implementation period.

When the increase in lump-sum taxes becomes effective in period T_p , the consumption function (18) indicates that it is almost fully anticipated by that date. Only the self-fulfilling channel, due to its persistence, is weakly driving up c_t . Thus, the tax increase mostly materializes in a drop in k_t via (1).

Throughout the post-implementation period, y_t decreases due to the shrinking k_t . In addition, w_t decreases and r_t increases. Over time, this yields downward revisions of $w^e(t)$ and upward revisions of $r^e(t)$. This reverses the evolution of c_t and decreases it. In the long run, the self-fulfilling channel gains importance. If agents manage to learn the new steady-state factor prices by continuous upward and downward revisions of factor prices, the economy eventually converges to the new steady-state.

In principle, the convergence is characterized by a sequence of waves of opti-

mism and pessimism about factor prices and other variables in consequence. The only reason, why those waves are not very apparent in the learning paths in Figs. 1a-1b is the particular calibration, as we will argue in Subsection 4.2 below.

In the case of *elastic labour supply* under *perfect foresight*, the dynamics are qualitatively similar to the ones with inelastic labour supply, which is well known from the standard Ramsey model. However, under *learning*, elastic labour supply has two major consequences. First, the impact effect on c_t is smaller, as $\eta > 0$ lowers the initial drop via (18). Second, actual factor prices and expectations about factor prices affect both c_t and n_t . The latter two are substitutes, and therefore the evolution of n_t is now subject to two opposing influences. Everything else equal, first, n_t tends to the opposite direction of c_t . Thus, given the drop in c_t on impact, households increase n_t , as can be seen from (6). Second, n_t in turn, now affects actual factor prices. On impact, the increase in n_t not only increases r_t and y_t , but decreases w_t , which in turn lowers n_t .

Lower c_t and larger y_t again create an investment boom before T_p . The intuition throughout the pre-implementation, implementation and post-implementation period is similar to the case of inelastic labour supply. However, now actual factor prices are exposed to two potentially opposing influences via k_t and n_t .

Nevertheless, the self-fulfilling channel is quantitatively more important right away and oscillations are amplified. The reason is that actual factor prices now react on impact, and expectations therefore already react in the period afterwards. In particular, $w^e(t)$ are first lowered, and, once the investment boom materializes, upward revised even in the pre-implementation period. Likewise, $r^e(t)$ are first increased and subsequently downward revised, even before T_p . These movements in expectations also explain, why c_t decreases further after T_p , before it recovers.

In sum, allowing for elastic labour supply results in amplified oscillations

starting at an earlier date. However, the oscillations are a feature of the samples that agents utilize to learn the steady-state values of factor prices and their resulting persistent and systematic forecast errors.

4.2. Sensitivity analysis

Compared to the previous literature, the learning approach herein introduces two additional structural parameters. One is γ , the gain parameter and a second one is the implementation date T_p . Therefore, we are interested in how these parameters affect the dynamics of the economy.

4.2.1. Sensitivity analysis for the gain parameter

There exists no consensus for the choice of γ in the learning literature. We are only aware of a single empirical estimate provided by Milani (2007, p.2074) for quarterly frequency and this is $\gamma = 0.0183$.¹⁹ But a reason to be cautious to use the estimate of Milani (2007, p.2074) is that it is based on a dataset containing output, inflation and the nominal interest rate, whereas in our setting agents forecast the rental rate of capital and the real wage. Next, Milani (2007, p.2074) mentions that a range of $\gamma \in [0.01, 0.03]$ is commonly used. Evans and Honkapohja (2009, p.154) note a range of $\gamma \in [0.01, 0.06]$ as known estimates.

Below we will present sensitivity of the dynamics under learning for $\gamma \in \{0.01, 0.08, 0.10\}$.²⁰ We do so for the original numerical analysis of Evans et al. (2009, p.943ff.) ($\bar{L} = 1.00$, $\eta = 0.00$), as in this case, there is inelastic labour supply and we can focus solely on the possible fluctuations introduced by varying the gain parameter γ .²¹

¹⁹This number indicates that agents use approximately $1/\gamma \approx 55$ quarters of data.

²⁰Values in the range $\gamma \in [0.002, 0.01]$ do not alter the conclusions in this subsection.

²¹Note that the two thick lines in Figs. 2a and 2b exactly replicate the Figs. 8 and 9 in Evans et al. (2009, p.943ff.).

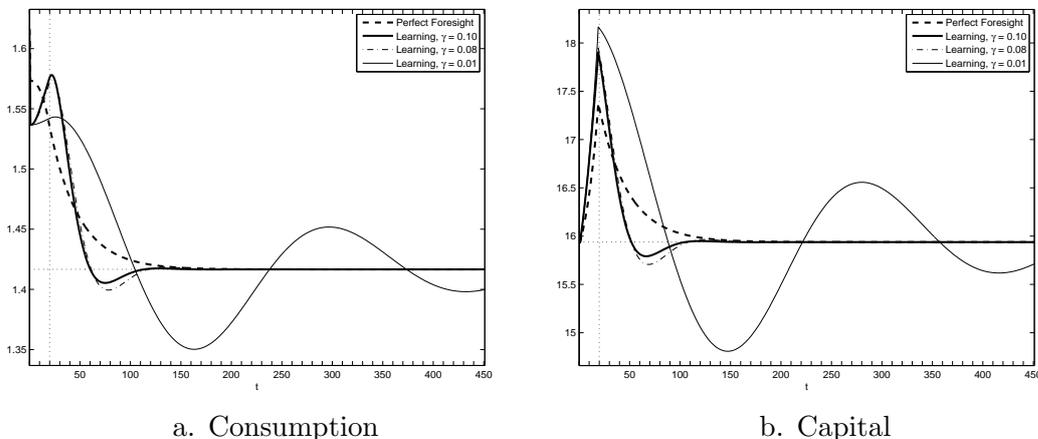


Fig. 2. Consumption (a) and capital (b) dynamics under learning and perfect foresight with inelastic labour supply as in Evans et al. (2009, p.943ff.) for alternative values of γ . The dotted horizontal line indicates the (new) steady-state, the dotted vertical line indicates period T_p .

In Fig. 2a we observe that the smaller γ , the smaller the increase in c_t until T_p (after the initial drop). Furthermore, as we recognize from Fig. 2b, the smaller γ , the larger the increase in k_t until T_p . However, in both Fig. 2a and 2b, we observe that, after T_p , with decreasing γ the dynamics fluctuate around the steady-state with increasing amplitude and convergence slows down. These observations are partly at odds with what Evans and Honkapohja (2001, p.332) report for the statistical literature: “a larger gain is better at tracking changes but at the cost of a larger variance”. In our case it holds, that a smaller gain yields larger volatility.²²

Inspection of the learning rule (11) explains this fact. The term $(1 - \gamma)^i$ indicates that a smaller γ , means stronger discounting of past observations.

²²Evans et al. (2011, p.23) confirm this observation with us in a similar model under educative learning. Both papers study a deterministic multivariate economy with capital accumulation, where agents form point expectations over an infinite-horizon. Exactly this combination of features distinguishes them from the existing literature, which reports the opposite result (e.g. see the references in Evans et al., 2011, p.23). Our observation may be the starting point for providing a general explanation as the result of future research.

Thus, with a smaller γ , agents have more confidence in their initial expectations. Throughout the pre-implementation period, a smaller γ means less optimism about SW_1 and approximately similar pessimism about ST_1 . In consequence, the temporary increase of c_t is lower and the temporary investment boom peaks at a higher level of k_t . Moreover, a smaller γ yields slower convergence in the post-implementation period, as the expectational errors are more persistent.

4.2.2. Sensitivity analysis for the implementation date

Next, we examine sensitivity of dynamics under learning for various implementation dates, in particular $T_p \in \{3, 10, 20\}$.

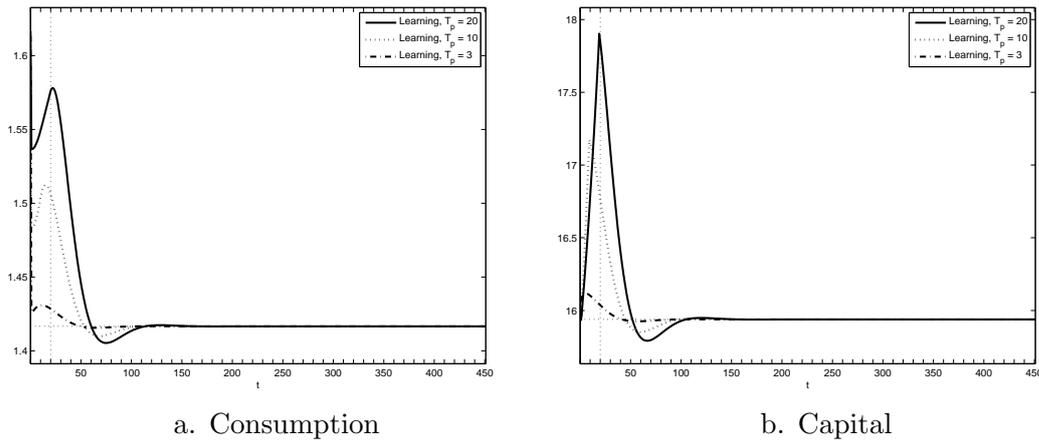


Fig. 3. Consumption (a) and capital (b) dynamics under learning with inelastic labour supply as in Evans et al. (2009, p.943ff.) for alternative values of T_p and $\gamma = 0.10$. The dotted horizontal line indicates the (new) steady-state, the dotted vertical line indicates period $T_p = 20$.

In Fig. 3a we observe that the shorter the distance between the announcement date and T_p , the higher the initial drop in c_t and the lower the increase in c_t until T_p thereafter. Next, in Fig. 3b we observe that with decreasing T_p , the peaks in k_t become smaller, but the learning dynamics remain qualitatively similar.

Inspection of SW_1 and ST_1 makes clear that T_p has only a direct impact on

ST_1 . If T_p decreases, agents are less optimistic about ST_1 and reduce c_t more on impact. With lower T_p agents are also less optimistic about ST_1 throughout the pre-implementation, thus the temporary recovery in c_t is dampened, the investment boom ends earlier, and convergence speeds up. However, the nature of dynamics is not seriously affected.

In order to summarize, there are three important insights from the analysis above. First, there are at least qualitative differences between the case of inelastic labour supply ($\eta = 0$) and elastic labour supply ($\eta > 0$). Therefore, if one regards the latter assumption as more realistic, a model that allows for elastic labour supply is a more appropriate framework to study anticipated fiscal policy under learning. Second, our sensitivity analysis suggests that the choice of the gain parameter γ and the implementation date T_p does not affect the nature of transition paths so we consider ourselves free to choose any of the values considered in the sensitivity analysis.²³ Finally and most notably, we observed a qualitative difference in the dynamics under learning compared to the dynamics under perfect foresight. The former appear to be much more volatile than the latter. This stylized fact motivates the quantification and comparison of welfare effects of anticipated fiscal policy reforms under learning and RE below.

5. The case of distortionary taxation

Herein we are utilizing our derivations from Subsection 3.2 to conduct an exemplary tax reform. In particular, we consider a credible pre-announced cut in τ_t^k and adjust τ_t^l such that the new steady-state government revenue is the same as before the policy change. The change is effective in $T_p = 8$, which will

²³In particular, in the subsequent analysis, we will choose $\gamma = 0.08$ and $T_p = 8$.

correspond to a duration of 2 years.²⁴ Moreover, we assume that the budget is balanced in each period and that $g_t = \bar{g} = 0.25$. Thus, we require τ_t to adjust throughout the transition. The calibrations used for the tax reform herein are $A = 1.00$, $\alpha = 0.33$, $\beta = 0.985$, $\eta = 4.00$, $\phi = 0.30$, $\delta = 0.025$, $T_p = 8$, $\gamma = 0.08$, and $\bar{L} = 0.15$. The calibrations of the preference parameters are guided by Mitra et al. (2013), which follow the RBC literature, see e.g. Heijdra (2009) or King and Rebelo (1999).²⁵

We evaluate the tax reform by welfare measures, which are computed following the approach of Cooley and Hansen (1992, p.301ff.). Intuitively speaking, we compute the increase in consumption that an individual would require to be as well off as under the equilibrium allocation without taxes. We express that number in percentage of output. Non-zero pre-reform tax rates, as chosen according to Table 1, lead to distortions. These distortions cause a pre-reform steady-state welfare loss of 4.75%.

Table 1

Tax rate settings and simulation results for the tax reform.

	τ_t	τ_t^k	τ_t^l	τ_t^c	\mathcal{W}^P	\mathcal{W}^L
Pre-reform	15.96%	50.00%	23.00%	5.00%	4.75%	4.75%
Post-reform	15.96%	25.00%	31.80%	5.00%	1.16%	4.52%
Difference	0.00%	-25.00%	+8.80%	0.00%	-3.59%	-0.23%

The initial tax rates are the same as in Cooley and Hansen (1992, p.305) except for the consumption tax. The latter is chosen as in Giannitsarou (2007, p.1433). \mathcal{W}^P and \mathcal{W}^L measure the welfare losses under perfect foresight and adaptive learning respectively. The pre-reform steady-state welfare losses are calculated according to Appendix D. The computation of post-reform welfare losses, which account for the entire transition, is detailed in Appendix E.

²⁴This reform is comparable to the ones of Cooley and Hansen (1992), Giannitsarou (2006), Judd (1987), or Lucas (1990). According to this literature, such a reform yields a welfare improvement under the RE hypothesis. The learning perspective, despite its empirical support (see for example Milani, 2011), is unknown for an anticipated reform.

²⁵Even though Mitra et al. (2013) consider a more general utility function, they calibrate it to coincide with a log-separable utility function for their numerical simulations.

The particular tax reform is a credible pre-announced cut in τ_t^k from 50% to 25% and an adequate increase in τ_t^l by 8.8%. Fig. 4 indicates that the dynamics are familiar to ones observed for the simple example discussed above. Note that the new steady-state c_t is higher and n_t is lower. This can be explained by the less severe distortions due to a lower τ_t^k . Next, Table 1 reveals that the reform yields a considerable welfare gain if agents have perfect foresight. The welfare loss decreases by 3.59% to 1.16%, when the whole transition is taken into account. However, the welfare loss under learning only decreases by 0.23% to 4.52%. This indicates that the reform does not yield large welfare improvements when agents lack perfect foresight.

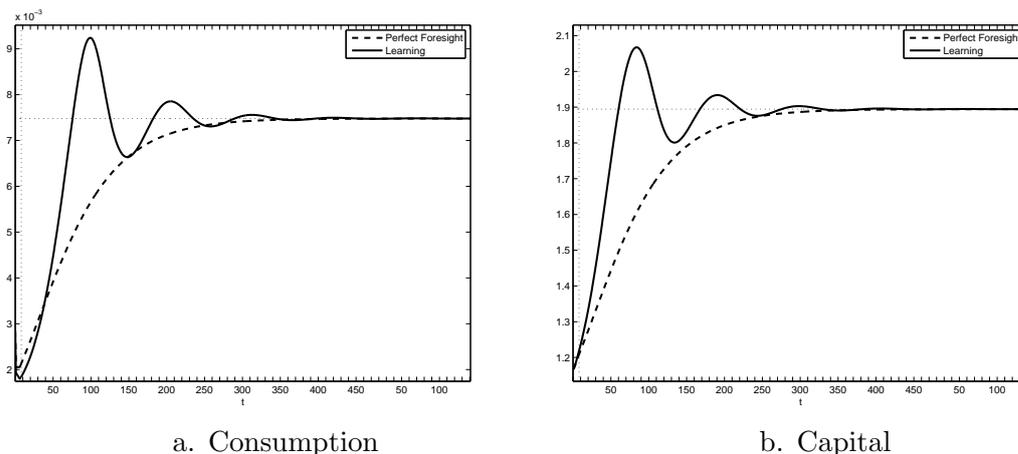


Fig. 4. Consumption (a) and capital (b) dynamics under learning (solid curve) and perfect foresight (dashed curve). The dotted horizontal line indicates the (new) steady-state, the dotted vertical line indicates T_p .

Some intuition for the transitory dynamics and the welfare loss can be developed. Under *perfect foresight*, it is well known that, on impact, the pre-announced cut in τ_t^k yields to a drop in c_t and therefore to an increase in n_t . This increases r_t as well as y_t , and decreases w_t on impact. Agents anticipate that saving is more attractive in the future and that this will have a positive net effect on

wealth. Therefore agents start accumulating k_t right away. The responses of non-predetermined variables to the simultaneously pre-announced increase in τ_t^l are in line with these impact effects. Subsequent accumulation of k_t yields increases in w_t as well as y_t , and a decrease in r_t . The monotonic evolution of variables continues until the saddle-path is reached in T_p . Thereafter, c_t and k_t monotonically increase to the new steady-state. The increase in c_t yields a decrease in n_t , that in turn decreases r_t and increases w_t . The sustained growth in k_t causes y_t to grow as well until the new steady-state is reached. Finally, constant government expenditures imply that there is no kink in the evolution of k_t in T_p .

Under *learning*, agents again make mistakes, when anticipating the effects of the policy change. They underestimate the positive net effect on wealth. However, as a feature of their sample and the learning rule (11), agents start forming either too optimistic or too pessimistic expectations about factor prices via the self-fulfilling channel, which has similar consequences throughout the pre- and post-implementation period as discussed above in Section 4. Thus, we observe waves of optimism and pessimism until the economy eventually converges to the new steady-state. Exactly those waves generated by systematic forecast errors cause excessive volatility of the economy and in turn are the source for the low welfare gain of the tax reform.

In sum, the tax reform indicates that the resulting welfare improvements of an anticipated tax reform might be much smaller in magnitude under learning compared to its improvements under perfect foresight. Notably, our results herein are somewhat similar to the findings in Giannitsarou (2006). This is remarkable, as Giannitsarou (2006) analyses unanticipated decreases in τ_t^k accompanied by a surprise exogenous technological shock in a RBC model. Furthermore, she

models adaptive learning following the “Euler Equation Approach”, whereas we utilize the “Long-Horizon Approach”, see Preston (2005, p.96) for the details. Finally, Giannitsarou (2006) conditions the welfare measure on the initial shock to technology. Nevertheless, Giannitsarou (2006) finds similar welfare implications: from an adaptive learning viewpoint the welfare gains from a supply-side reform are smaller, than what is known from the viewpoint of the RE hypothesis.

6. Conclusion

We demonstrate within the Ramsey model that the responses to anticipated permanent tax changes when agents learn are remarkably different compared to their counterparts under perfect foresight. The learning dynamics appear to oscillate around the steady-state to which they converge slowly. Thus, the economy exhibits more volatility under learning.

We argue that the observed oscillations are related to expectational errors. These systematic forecast errors are caused by the anticipated permanent tax change and lead agents to incorrectly quantify the effects of tax reforms on their life-time wealth. The persistence of the expectational errors in the learning rule of the agents is the fundamental reason for agents to alternately be either too optimistic or too pessimistic about the consequences of a pre-announced tax change.

These learning dynamics, despite the simplicity of the model, have the potential to capture important features of the empirical evidence on how the economy responds to policy changes. However, we believe that future research in this area needs to come up with more empirical evidence on whether or how agents learn about fiscal policy.

Apart from this, our sensitivity analyses show that a smaller gain parameter

leads to higher volatility in the presence of policy change. This result is at odds with conventional wisdom about the link between the gain parameter and the dynamic responses in the literature and requires further investigation as well.

Finally, an exemplary anticipated capital tax reform indicates that the magnitude of welfare improvements appears to be substantially lower under the assumption of learning compared to the case of perfect foresight. This can be explained by the oscillations under learning. Thus, the learning perspective on the tax reform considered herein has a more general implication for benevolent policy makers. Tax reforms may not lead to a considerable aggregate welfare improvement compared to the status quo. In consequence, other criteria, like the distributional consequences of tax reforms under the assumption of heterogeneity, may become more important in the design of such tax reforms.

Acknowledgements

We are indebted to Alejandro Cuñat, Robert Kunst and Gerhard Sorger for sound supervision and patience. We are also grateful to the editors and two anonymous referees, Szabolcs Deák, Seppo Honkapohja, Torsten Lisson, Albert Marcet, Kaushik Mitra, Paul Pichler, and Monika Turyna as well as the participants at the 4th RGS Doctoral Conference in Economics, the 2011 Annual Conference of the Scottish Economic Society and the Quantitative Economics Doctorate (QED) Jamboree 2011 for valuable comments. Gasteiger acknowledges financial support from Fundação para a Ciência e a Tecnologia (PEst-OE/EGE/UI0315/2011). All remaining errors are the responsibility of the authors.

Appendix A. Derivation of SW_2

We start from (15), i.e. $\sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} w_{t+j}^e(t) \bar{L}$. Next, we recall the definition of $D_{t,t+j}^{k,e}(t)$. Given the learning rule (11) we get

$$D_{t,t+j}^{k,e}(t) = \begin{cases} \prod_{i=1}^j [(1-\delta) + (1-\tau_0^k)r^e(t)] = [(1-\delta) + (1-\tau_0^k)r^e(t)]^j & \text{for } \tau_{t+j}^{k,e}(t) = \tau_0^k \\ \prod_{i=1}^j [(1-\delta) + (1-\tau_1^k)r^e(t)] = [(1-\delta) + (1-\tau_1^k)r^e(t)]^j & \text{for } \tau_{t+j}^{k,e}(t) = \tau_1^k. \end{cases} \quad (21)$$

Thereafter, we split this infinite sum into

$$\begin{aligned} SW_2 &= \bar{L} \left[\sum_{j=1}^{T-1} \frac{1}{D_{t,t+j}^{k,e}(t)} w^e(t) + \sum_{j=T}^{\infty} \frac{1}{D_{t,t+j}^{k,e}(t)} w^e(t) \right] \\ &= \bar{L} \left[\sum_{j=1}^{T-1} ([(1-\delta) + (1-\tau_0^k)r^e(t)]^j)^{-1} w^e(t) + \sum_{j=T}^{\infty} ([(1-\delta) + (1-\tau_1^k)r^e(t)]^j)^{-1} w^e(t) \right], \end{aligned}$$

or

$$\begin{aligned} SW_2 &= \frac{w^e(t) \bar{L}}{[(1-\delta) + (1-\tau_0^k)r^e(t)]} \sum_{j=0}^{T-2} \left([(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} \right)^j \\ &\quad + \frac{w^e(t) \bar{L}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]} \sum_{j=T-1}^{\infty} \left([(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} \right)^j. \end{aligned}$$

Now we exploit the properties of a geometric series, e.g. $\sum_{j=m}^n f^j = \frac{f^{n+1} - f^m}{f-1}$ for some constant f , and derive

$$\begin{aligned} SW_2 &= \frac{w^e(t) \bar{L}}{[(1-\delta) + (1-\tau_0^k)r^e(t)]} \left(\frac{1 - [(1-\delta) + (1-\tau_0^k)r^e(t)]^{1-T}}{1 - [(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1}} \right) \\ &\quad + \frac{w^e(t) \bar{L}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]} \left(\frac{[(1-\delta) + (1-\tau_1^k)r^e(t)]^{1-T}}{1 - [(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1}} \right). \end{aligned}$$

When we respect the timing outlined in the experiment above, we get (15)

$$SW_2 = \begin{cases} \frac{w^e(t)\bar{L}}{[(1-\tau_0^k)r^e(t)-\delta]} + w^e(t)\bar{L} \left[\frac{[(1-\delta)+(1-\tau_1^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_1^k)r^e(t)]^{-1}} - \frac{[(1-\delta)+(1-\tau_0^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_0^k)r^e(t)]^{-1}} \right] & \text{for } 1 \leq t < T_p \\ \frac{w^e(t)\bar{L}}{[(1-\tau_1^k)r^e(t)-\delta]} & \text{for } t \geq T_p. \end{cases}$$

Appendix B. Derivation of ST_2

Starting from (16), i.e. $\sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}} \tau_{t+j}^{l,e}(t) w_{t+j}^e(t) \bar{L}$, for (21) and $\tau_{t+j}^{l,e}(t)$ is either given by τ_0^l or τ_1^l , we may again split the infinite sum into

$$\begin{aligned} ST_2 &= w^e(t)\bar{L} \times \\ &\quad \left[\sum_{j=1}^{T-1} \left([(1-\delta) + (1-\tau_0^k)r^e(t)]^j \right)^{-1} \tau_0^l + \sum_{j=T}^{\infty} \left([(1-\delta) + (1-\tau_1^k)r^e(t)]^j \right)^{-1} \tau_1^l \right], \\ &= \frac{\tau_0^l w^e(t)\bar{L}}{[(1-\delta) + (1-\tau_0^k)r^e(t)]} \sum_{j=0}^{T-2} \left([(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} \right)^j \\ &\quad + \frac{\tau_1^l w^e(t)\bar{L}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]} \sum_{j=T-1}^{\infty} \left([(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} \right)^j. \end{aligned}$$

Now, the properties of the geometric series allow us to rewrite this as

$$\begin{aligned} ST_2 &= \frac{\tau_0^l w^e(t)\bar{L}}{[(1-\delta) + (1-\tau_0^k)r^e(t)]} \left(\frac{[(1-\delta) + (1-\tau_0^k)r^e(t)]^{1-T} - 1}{[(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} - 1} \right) \\ &\quad + \frac{\tau_1^l w^e(t)\bar{L}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]} \left(\frac{-[(1-\delta) + (1-\tau_1^k)r^e(t)]^{1-T}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} - 1} \right). \end{aligned}$$

Given the timing outlined above, we get (16)

$$ST_2 = \begin{cases} \frac{\tau_0^l w^e(t)\bar{L}}{[(1-\tau_0^k)r^e(t)-\delta]} + w^e(t)\bar{L} \left[\frac{\tau_1^l [(1-\delta)+(1-\tau_1^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_1^k)r^e(t)]^{-1}} - \frac{\tau_0^l [(1-\delta)+(1-\tau_0^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_0^k)r^e(t)]^{-1}} \right] & \text{for } 1 \leq t < T_p \\ \frac{\tau_1^l w^e(t)\bar{L}}{[(1-\tau_1^k)r^e(t)-\delta]} & \text{for } t \geq T_p. \end{cases}$$

Appendix C. Derivation of ST_3

Starting from (17), i.e. $\sum_{j=1}^{\infty} \frac{1}{D_{t,t+j}^{k,e}} \tau_{t+j}^e(t)$, given (21) and $\tau_{t+j}^e(t)$ is either τ_0 or τ_1 , we again split the infinite sum into

$$\begin{aligned} ST_3 &= \left[\sum_{j=1}^{T-1} \left([(1-\delta) + (1-\tau_0^k)r^e(t)]^j \right)^{-1} \tau_0 + \sum_{j=T}^{\infty} \left([(1-\delta) + (1-\tau_1^k)r^e(t)]^j \right)^{-1} \tau_1 \right], \\ &= [(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} \left[\sum_{j=0}^{T-2} \left([(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} \right)^j \tau_0 \right] \\ &\quad + [(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} \left[\sum_{j=T-1}^{\infty} \left([(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} \right)^j \tau_1 \right]. \end{aligned}$$

Given the properties of geometric series we can rewrite the latter as

$$\begin{aligned} ST_3 &= [(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} \left(\frac{[(1-\delta) + (1-\tau_0^k)r^e(t)]^{1-T} - 1}{[(1-\delta) + (1-\tau_0^k)r^e(t)]^{-1} - 1} \tau_0 \right) \\ &\quad + [(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} \left(\frac{-[(1-\delta) + (1-\tau_1^k)r^e(t)]^{1-T}}{[(1-\delta) + (1-\tau_1^k)r^e(t)]^{-1} - 1} \tau_1 \right). \end{aligned}$$

Once more, given the timing outlined above, we get (17)

$$ST_3 = \begin{cases} \frac{\tau_0}{[(1-\tau_0^k)r^e(t)-\delta]} + \left[\frac{[(1-\delta)+(1-\tau_1^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_1^k)r^e(t)]^{-1}} \tau_1 - \frac{[(1-\delta)+(1-\tau_0^k)r^e(t)]^{t-T_p}}{1-[(1-\delta)+(1-\tau_0^k)r^e(t)]^{-1}} \tau_0 \right] & \text{for } 1 \leq t < T_p \\ \frac{\tau_1}{[(1-\tau_1^k)r^e(t)-\delta]} & \text{for } t \geq T_p. \end{cases}$$

Appendix D. Computing welfare changes: comparative statics

We follow the approach of Cooley and Hansen (1992, p.301ff.). Their measure of welfare change for a given policy change is derived by solving

$$U_0 = \log[c_1(1+x^\bullet)] + \eta \log[1-n_1] + \phi \log[g_1] \quad (22)$$

for x^\bullet .²⁶ U_0 is the utility a household obtains in the steady-state without any tax. c_1 and n_1 are the values of consumption and employment at the new steady-state after the tax change either under perfect foresight or learning. It follows that

$$x^\bullet = \frac{\exp(U_0)}{c_1(1 - n_1)^\eta g_1^\phi} - 1. \quad (23)$$

Thus, in general, we need to solve for x^P for the perfect foresight dynamics and another x^L for the dynamics under learning.²⁷ Given x^\bullet we can calculate

$$\bar{W} = \frac{\Delta C}{y_1} = \frac{x^\bullet c_1}{y_1}, \quad (24)$$

where ΔC is the restoration value of consumption, which in our case may be interpreted as the total change in consumption required to restore a household to the level of utility obtained under the allocation associated with zero taxes. y_1 is the level of output at the new steady-state.

Appendix E. Computing welfare changes: transition measure

Again we follow the approach of Cooley and Hansen (1992, p.301ff.) based on Lucas (1990). Their measure of welfare change accounting for transition given a policy change is derived by solving

$$\sum_{t=1}^{\infty} \beta^t \{ \log[c_t(1 + x^\bullet)] + \eta \log[1 - n_t] + \phi \log[g_t] - U_0 \} \stackrel{!}{=} 0 \quad (25)$$

for x^P under perfect foresight and x^L under learning. c_t , n_t , g_t , and y_t are period t consumption, employment, government spending and output respectively, either

²⁶ x^\bullet is either x^P under perfect foresight or x^L under learning.

²⁷This must yield the same $x^P = x^L$ both under perfect-foresight and under learning, but this number may be useful to compare different policy experiments.

under perfect foresight or learning. It follows that

$$x^\bullet = \left[\frac{\exp\left(\frac{\beta}{(1-\beta)}[U_0 - \beta^T U_1]\right)}{\prod_{t=1}^T c_t^{\beta^t} \times \prod_{t=1}^T (1 - n_t)^{\eta\beta^t} \times \prod_{t=1}^T g_t^{\phi\beta^t}} \right]^{\frac{(1-\beta)}{\beta}} - 1. \quad (26)$$

Hereby T is the terminal period of the simulation. For the $T < t \leq \infty$ it is assumed that agents' period utility is approximately given by U_1 , i.e. all variables involved are close to their new steady-state value. Thus, T needs to be sufficiently large such that x^\bullet does no longer change significantly. Given x^\bullet we can calculate

$$\mathcal{W}^\bullet = \frac{\sum_{t=1}^T \beta^t \{x^\bullet c_t\}}{\sum_{t=1}^T \beta^t \{y_t\}}, \quad (27)$$

which will be reported as \mathcal{W}^P for the perfect foresight dynamics and as \mathcal{W}^L for the dynamics under learning.

References

- Ambler, S., Paquet, A., 1996. Fiscal spending shocks, endogenous government spending, and real business cycles. *Journal of Economic Dynamics and Control* 20, 237–256.
- Branch, W.A., Evans, G.W., 2007. Model uncertainty and endogenous volatility. *Review of Economic Dynamics* 10, 207–237.
- Braun, R.A., 1994. Tax disturbances and real economic activity in the postwar United States. *Journal of Monetary Economics* 33, 441–462.
- Chamley, C., 1981. The welfare cost of capital income taxation in a growing economy. *The Journal of Political Economy* 89, 468–496.
- Christiano, L.J., Eichenbaum, M., 1992. Current real business cycle theory and aggregate labor market fluctuations. *American Economic Review* 82, 430–450.
- Cooley, T.F., Hansen, G.D., 1992. Tax distortions in a neoclassical monetary economy. *Journal of Economic Theory* 58, 290–316.

- Evans, G.W., Guesnerie, R., McGough, B., 2011. Eductive stability in real business cycle models. Mimeo .
- Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.
- Evans, G.W., Honkapohja, S., 2009. Robust learning stability with operational monetary policy rules, in: Schmidt-Hebbel, K., Walsh, C.E. (Eds.), Monetary Policy Under Uncertainty and Learning. Central Bank of Chile, Santiago, pp. 145–170.
- Evans, G.W., Honkapohja, S., Mitra, K., 2009. Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics* 56, 930–953.
- Evans, G.W., Honkapohja, S., Mitra, K., 2012. Does Ricardian equivalence hold when expectations are not rational? *Journal of Money, Credit and Banking* 44, 1259–1283.
- Farmer, R.E.A., 1999. *Macroeconomics of Self-fulfilling Prophecies*, second ed. MIT Press, Cambridge, MA.
- Garcia-Milà, T., Marcet, A., Ventura, E., 2010. Supply side interventions and redistribution. *The Economic Journal* 120, 105–130.
- Giannitsarou, C., 2006. Supply-side reforms and learning dynamics. *Journal of Monetary Economics* 53, 291–309.
- Giannitsarou, C., 2007. Balanced budget rules and aggregate instability: The role of consumption taxes. *The Economic Journal* 117, 1423–1435.
- Heijdra, B.J., 2009. *Foundations of Modern Macroeconomics*, second ed. Oxford University Press, Oxford.
- Judd, K.L., 1987. The welfare cost of factor taxation in a perfect-foresight model. *The Journal of Political Economy* 95, 675–709.
- King, R.G., Rebelo, S.T., 1999. Resuscitating real business cycles, in: Taylor, J.B., Woodford, M. (Eds.), *Handbook of Macroeconomics*, vol. 1. Elsevier, Amsterdam, pp. 927–1007.
- Kreps, D.M., 1998. Anticipated utility and dynamic choice, in: Jacobs, D.P., Kalai, E., Kamien, M.I. (Eds.), *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press, Cambridge, UK.
- Leeper, E.M., 2009. Anchoring fiscal expectations. *Reserve Bank of New Zealand Bulletin* 72, 17–42.
- Leeper, E.M., Walker, T.B., Yang, S.C.S., 2013. Fiscal foresight and information flows. *Econo-*

- metrica 81, 1115–1145.
- Ljungqvist, L., Sargent, T.J., 2012. *Recursive Macroeconomic Theory*, third ed. MIT Press, Cambridge, MA.
- Lucas, Jr., R.E., 1990. Supply-side economics: an analytical review. *Oxford Economic Papers* 42, 293–316.
- Milani, F., 2007. Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics* 54, 2065–2082.
- Milani, F., 2011. Expectation shocks and learning as drivers of the business cycle. *The Economic Journal* 121, 379–401.
- Mitra, K., Evans, G.W., Honkapohja, S., 2013. Policy change and learning in the RBC model. *Journal of Economic Dynamics and Control* 37, 1947–1971.
- Prescott, E.C., 1986. Theory ahead of business-cycle measurement. *Carnegie-Rochester Conference Series on Public Policy* 25, 11–44.
- Preston, B., 2005. Learning about monetary policy rules when long-horizon expectations matter. *International Journal of Central Banking* 1, 81–126.