UNIVERSITEIT GENT

FACULTEIT ECONOMIE
EN BEDRIJFSKUNDE

TWEEKERKENSTRAAT 2
B-9000 GENT
Tel. : 32 - (0)9 – 264.34.61
Fax. : 32 - (0)9 – 264.35.92

WORKING PAPER

Real-time parameterized expectations and the effects of government spending

Brecht Boone

Ewoud Quaghebeur

June 2017

2017/939

D/2017/7012/13

Real-time parameterized expectations and the effects of government spending

Brecht Boone and Ewoud Quaghebeur*

June 2017

Abstract

In this paper, we explore the effects of government spending in the real business cycle model where agents use a learning mechanism to form expectations. In contrast to most of the learning literature, we study learning behaviour in the original non-linear model. Following the learning interpretation of the parameterized expectations method, agents' forecast rules are approximations of the conditional expectations appearing in the Euler equation. We show that variation in agents' beliefs about the coefficients of these rules, generates time variation in the transmission of government spending shocks to the economy. Hence, our modelling approach provides an endogenous mechanism for time-varying government spending multipliers in the standard real business cycle model.

Keywords: Non-linear learning, Parameterized expectations, Fiscal policy, Time-varying multipliers

JEL Classification: E62, D83, D84, E32.

1 Introduction

The macroeconomic impact of fiscal policy depends crucially on the behavioural response of households to these policies. An important determinant of this behavioural response is the approach households apply to form expectations regarding the evolution of different endogenous, macroeconomic variables. The dominant paradigm used to model expectations in macroeconomics is the rational expectations (RE) hypothesis. According to this hypothesis, households have perfect knowledge about the structure of the model and understand the full complexities of the macro-economy. An alternative to the rational expectations hypothesis is provided by the learning literature (see e.g. Evans and Honkapohja, 2001).[1] In this literature, agents form expectations using a perceived law of motion. Over time, as new information becomes available, agents update the coefficients of their perceived law of motion.

This observation raises the question of whether the effects of government spending and the transmission thereof in the macro-economy are different in the basic Real Business Cycle (RBC) model using respectively rational expectations and a learning set-up. Indeed, it is well known that government spending multipliers generated by standard RBC and Dynamic Stochastic General Equilibrium (DSGE) models using rational expectations are typically constant. This theoretical finding is, however, not in accordance with the findings of several empirical studies. Auerbach and Gorodnichenko (2012) and Owyang et al. (2013), among others, show that the government spending multiplier is time-varying. In this paper, we start from these empirical findings and show that the introduction of a learning set-up in the standard RBC model can generate substantial time variation in the government spending multiplier.

*Ghent University, Department of Social Economics and Study Hive for Economic Research and Public Policy Analysis (SHERPPA). Correspondence to Brecht.Boone@UGent.be or Ewoud.Quaghebeur@UGent.be. Sint-Pietersplein 6, B-9000 Ghent, Belgium. Phone: +32 9 264.34.87.

[1]In this paper we consider *Euler equation* learning as put forward by Evans and Honkapohja (2001). In this approach, agents make one-step ahead forecasts. By contrast, the *infinite horizon* approach of Preston (2005) assumes that at each date agents make forecasts about variables into the infinite future. For a discussion of these two approaches see Honkapohja et al. (2013).

Several papers have already explored the effects of fiscal policy using a learning framework within an RBC or DSGE model. Evans et al. (2009), for example, study the effects of anticipated fiscal policy changes both within an endowment economy and the Ramsey model. Their assumption is that agents fully understand and anticipate the evolution of taxes but have to forecast future factor prices using a linear learning mechanism. Building on this framework, Mitra et al. (2013) generalise the analysis of Evans et al. (2009) to a stochastic environment with elastic labour supply. Gasteiger and Zhang (2014) extend the model even further by introducing distortionary taxes. Following a similar learning approach, Benhabib et al. (2014) investigate the effects of fiscal stimulus in a new Keynesian model with a zero lower bound on the nominal interest rate.

All the aforementioned papers have enriched our knowledge on the macroeconomic effects of fiscal policy. They show that these effects can be substantially different when using the learning approach instead of the rational expectations approach. However, none of these studies explicitly focuses on the evolution and level of the government spending multiplier. Furthermore, they all study learning in the linearised counterpart of the non-linear model they are using. Indeed, it is common in the learning literature that non-linear models are first linearised around the rational expectations solution before studying their dynamics under learning. Exploring learning and the link with fiscal policy in the original non-linear model has several advantages, though. First, there is no longer the need to linearise the RE model around its steady state. Furthermore, contrary to linearised models, which necessarily lead to a local stability analysis, the context of the original non-linear system allows one to provide a more global stability analysis. Last, due to the non-linearity of the system, it is more natural to use non-linear forecasting rules compared to linearised models. This way, the usefulness of non-linear forecasting rules can be studied as well. As such, one can allow for the possibility of non-linear responses from households.

In this paper, we explore the transitional effects of government spending and the behaviour of the government spending multiplier by introducing learning in the in the original non-linear RBC model. More specifically, we adopt the learning interpretation of the parameterized expectations algorithm (PEA). This algorithm was initially developed as a solution method for non-linear, stochastic models with rational expectations (see for example den Haan and Marcet, 1990; Marcet and Lorenzoni, 1999; Marcet and Marshall, 1994). The idea behind the PEA is to replace the conditional expectations in the equilibrium conditions of the model with flexible functional forms with a finite number of arguments, e.g. polynomials. However, in Marcet and Marshall (1994), the authors also give an alternative, learning interpretation to the solution of the PEA.

We find that learning in the non-linear model leads to substantial time variation in the transmission of structural shocks in the model economy. As such, this result stands in sharp contrast with the RE and PEA solutions of the model. Our set-up thus leads to time-varying government spending multipliers. Furthermore, the time variation in our set-up itself is endogenously determined. As the economic agents update their beliefs, their response to a change in government spending changes as well, leading to a different impact on the economy.

The remainder of this paper is structured as follows. Section 2 describes the model. In Section 3, the learning mechanism is outlined in detail and compared with the rational expectations solution and the PEA solution of the model. Section 4 shows how learning behaviour in our model leads to time variation in the government spending multipliers. Finally, Section 5 concludes.

2 Model

To study the transitional dynamics of fiscal policy changes, we use the standard RBC model with elastic labour supply. In this section, we briefly introduce the different components of the model.

2.1 Households

The maximization problem of the representative household consists in maximizing

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\sigma}}{1-\sigma} + b \frac{(1-n_t)^{1-\theta}}{1-\theta} \right] \tag{1}$$

subject to its budget constraint:

$$c_t + k_{t+1} = w_t n_t + (1+r_t) k_t - T_t. \tag{2}$$

In these equations, c_t represents the household's consumption, k_{t+1} denotes the capital stock, n_t is its labour supply, and w_t and r_t are the real wage and the real interest rate. The latter is equal to the rental charge on capital after depreciation ($r_t = r_t^k - \delta$). Furthermore, b is the taste for leisure, β the discount factor, σ the coefficient of relative risk aversion, and θ the inverse of the inter-temporal elasticity of substitution in leisure. T_t is the lump sum tax in period t.

The optimality conditions of the household with respect to labour and consumption are respectively given by

$$c_t^{-\sigma} = \frac{b(1-n_t)^{-\theta}}{w_t} \tag{3}$$

and

$$\beta E_t \left\{ c_{t+1}^{-\sigma} \left[r_{t+1}^k + 1 - \delta \right] \right\} = c_t^{-\sigma}. \tag{4}$$

2.2 Firms

The representative firm produces the final good according to

$$y_t = z_t k_t^{\alpha} n_t^{1-\alpha} \tag{5}$$

where z_t evolves according to

$$z_t = z_{t-1}^{\rho_z} \exp(\varepsilon_t^Z). \tag{6}$$

with $\rho_z \in (0,1)$. In these equations, $\varepsilon_t^Z \sim \mathscr{N}(0, \sigma_Z^2)$ is an innovation in technology. Profits are given by

$$\mathscr{L} = z_t k_t^{\alpha} n_t^{1-\alpha} - w_t n_t - r_t^k k_t \tag{7}$$

and the corresponding first-order conditions with respect to labour and capital respectively are:

$$w_t = (1-\alpha) z_t k_t^{\alpha} n_t^{-\alpha}, \tag{8}$$

$$r_t^k = \alpha z_t k_t^{\alpha-1} n_t^{1-\alpha}. \tag{9}$$

2.3 Government

The fiscal government finances its expenditures on goods by levying lump sum taxes. Formally, we have

$$g_t = T_t. \tag{10}$$

Government spending g_t evolves according to

$$g_t = g_{t-1}^{\rho_g} \exp(\varepsilon_t^G), \tag{11}$$

with $\rho_g \in (0,1)$ and where $\varepsilon_t^G \sim \mathscr{N}(0, \sigma_G^2)$ is a government spending shock.

3 Real-time non-linear learning

3.1 Set-up

We assume that agents, instead of having rational expectations, form forecasts by means of a non-linear learning mechanism. In particular, agents in our model approximate the conditional expectational function in the consumption Euler equation

$$E_t\phi\left(s_{t+1}\right) = E_t\left[c_{t+1}^{-\sigma}\left(1-\delta+\alpha z_{t+1}k_{t+1}^{\alpha-1}n_{t+1}^{1-\alpha}\right)\right], \tag{12}$$

where $s_t = [c_t, k_t, z_t]$, by a parametric function $\psi\left(x_t, \gamma_{t-1}\right)$ of the state variables $x_t = [1, k_t, z_t, g_t]$, and update the parameters γ_{t-1} using a constant-gain variant of recursive least squares. We follow the parameterized expectations literature and use an exponentiated polynomial to approximate the expectational function. More precisely, we consider the following first-order polynomial in the state variables of the model

$$E_t\phi\left(s_{t+1}\right) \simeq \psi\left(x_t, \gamma_{t-1}\right) = \exp\left[\gamma_0 + \gamma_1 \log k_t + \gamma_2 \log z_t + \gamma_3 \log g_t\right]. \tag{13}$$

Agents update the vector of belief parameters γ_t in real time according to this learning rule:

$$\gamma_t = \gamma_{t-1} + \kappa S_t^{-1} x_{t-1}\left[\log\left(\phi\left(s_t\right)\right) - \log\left(\psi\left(x_{t-1}, \gamma_{t-2}\right)\right)\right], \tag{14}$$

$$S_t = S_{t-1} + \kappa\left[x_{t-1}x_{t-1}^{'} - S_{t-1}\right], \tag{15}$$

where S_t is the moment matrix for x_t and $\kappa \in (0,1)$ is the gain parameter. If the gain parameter κ would be equal to t^{-1}, equations (14)–(15) are the recursive formulas of the ordinary least squares (OLS) estimator for the coefficients γ in the log-linear specification of equation (13), i.e. $\log\left(\phi\left(s_t\right)\right) \simeq \log\left(\psi\left(x_{t-1}, \gamma\right)\right)$. The update for γ in equation (14) uses the most recent forecast error$\log\left(\phi\left(s_t\right)\right) - \log\left(\psi\left(x_{t-1}, \gamma_{t-2}\right)\right)$.

Instead of adopting the (decreasing-gain) OLS algorithm, where $\kappa = t^{-1}$, we set the gain κ to a small constant. The constant-gain case is the most relevant one for our analysis and is widely used in the adaptive learning literature (see Eusepi and Preston, 2011; Milani, 2007; Slobodyan and Wouters, 2012, for example). As argued by, for example, Sargent (1999), Cho et al. (2002), Branch and Evans (2007), Milani (2007), and Orphanides and Williams (2007), it is a natural way of accomplishing "perpetual learning" as it places a greater weight on more recent observations. Hence, the constant gain makes sure that the variation in the agents' beliefs does not die out over time.

We are now able to formulate the dynamics of our model under non-linear learning. For some initial beliefs γ_0 and an initial state vector $x_0 = [1, k_0, z_0, g_0]$, substituting the approximating function $\psi\left(x_t, \gamma_{t-1}\right)$ for the expectation function in equation (4) gives consumption $c_t = \beta\left[\psi\left(x_t, \gamma_{t-1}\right)\right]^{-1/\sigma}$, equation (3) determines labour supply n_t, and k_{t+1} follows from the resource constraint (2). The procedure is detailed in Appendix B.

3.2 Discussion

Originally, the parameterized expectations approach was used as a method to approximate the rational expectations equilibrium of non-linear stochastic dynamic models. The original parameterized expectations algorithm (PEA) starts from a large sequence of shocks $\{z_t, g_t\}_{t=0}^{T}$, computes the corresponding endogenous variables consistent with the parameterized expectations, and iterates on the vector of parameters γ until the approximation becomes sufficiently accurate. A detailed step-by-step description of the algorithm is given in Appendix A.

In this paper, we follow the suggestion of Marcet and Marshall (1994) and interpret the parameterized expectations approach as a real-time learning mechanism. Instead of holding the parameters γ constant over the stochastic simulation of the model, we let agents update them according to the learning rule above each time new information becomes available. In doing so we build on Berardi and Duffy (2015), who applied a similar learning mechanism to an optimal growth model.

Parameter	Description	Value
α	Output elasticity of capital	$1/3$
β	Discount factor	0.98
σ	Coefficient of risk aversion	1
δ	Depreciation rate	0.025
b	Taste for leisure	1.44
θ	Preference parameter	1
$\bar{g}/\bar{y}$	Steady state government expenditure to output ratio	0.2
ρ_Z	Technology shock AR(1) coefficient	0.9
ρ_G	Government spending AR(1) coefficient	0.9
σ_Z	Standard deviation of the technology disturbance ε^Z	0.01
σ_G	Standard deviation of the fiscal disturbance ε^G	0.01
κ	Gain parameter in the learning mechanism	0.02

Table 1: Calibration.

3.3 Non-linear learning simulation

Before turning to the effects of government spending shocks, we compare the non-linear learning solution of the model with two conventional alternatives: the rational expectations solution and the solution provided by the original parameterized expectations algorithm (PEA).[2] Recall that in the latter case, the parameters of the approximating function are held fixed whereas in the non-linear learning model the parameters are updated over time. Given a sequence of 10,000 structural shocks, we generate a series of endogenous variables for these three different solutions.

Our assumptions on the parameter values are listed in Table 2 and are in line with the literature. The baseline value for the gain parameter, κ, is set to 0.02.[3] Given that this is a crucial parameter for the learning dynamics, we discuss the implications of different choices for κ in Section 4. The other parameters are set to values commonly used in the literature. The output elasticity of capital α is set to 1/3. According to Rogerson (2007), a reasonable range for θ in models with a macro focus is $[1,3]$. We set θ equal to 1. The coefficient of risk aversion σ equals 1 and the discount factor β is set to 0.98. The depreciation rate δ equals 0.025. The AR(1) coefficients of technology, ρ_Z, and government spending, ρ_G, are set to 0.9. The share of government spending in GDP, $\bar{g}/\bar{y}$, is set to 0.2. For the standard deviations of the technology and fiscal disturbances, σ_Z and σ_G, the value 0.01 is chosen. The parameter capturing the taste for leisure, b, is determined such that individuals work on average 1/3 of their time endowment.

It is clear from Figure 1 that the simulated series of both the learning model and the parameterized expectations algorithm are close to the rational expectations series. The figure shows the evolution of consumption, output, labour, and investment for 500 periods of the full sample.[4] The relatively small differences between the rational expectations and parameterized expectations solution, indicate that the latter is a reasonably good approximation of the former. This comforts us in the choice of the functional form of the approximating function. The simulation results also illustrate the local stability of the non-linear learning mechanism.[5] Over our long simulation horizon, the recursive estimation of the belief parameters does not drive the economy towards diverging paths. This stability is also reflected in the evolution of the belief parameters over time, depicted in Figure 2. Although the parameters can deviate from their PEA values for a sustained period of time, they remain in the neighbourhood of those

[2]We use the deterministic extended path method to find the rational expectations solution of the model.

[3]This is a value well within the range of estimates reported in the literature. According to Orphanides and Williams (2005, 2007) a gain parameter in the range 0.01–0.04 provides the best fit between the agents' forecasts in the model and the expectations data from the Survey of Forecasters. Using a similar strategy, Branch and Evans (2006) obtain a value of 0.0345. The estimate of Milani (2007) equals 0.0183 and hence lies within the same range. The estimated gain in Eusepi and Preston (2011) is 0.0029. The Bayesian estimation results from Slobodyan and Wouters (2012) provide values for κ going from 0.001 to 0.06 depending on the particular learning scheme.

[4]Investment is defined as $i_t = y_t - c_t - g_t$.

[5]Since our analysis is numerical, establishing analytical convergence results is beyond the scope of this paper. Stability results for constant-gain algorithms can be found in (inter alia) Evans and Honkapohja (2001). Loosely speaking, constant-gain learning mechanisms are locally stable only if the gain parameter is sufficiently small. We discuss the sensitivity of our results with respect to this parameter in Section 4.

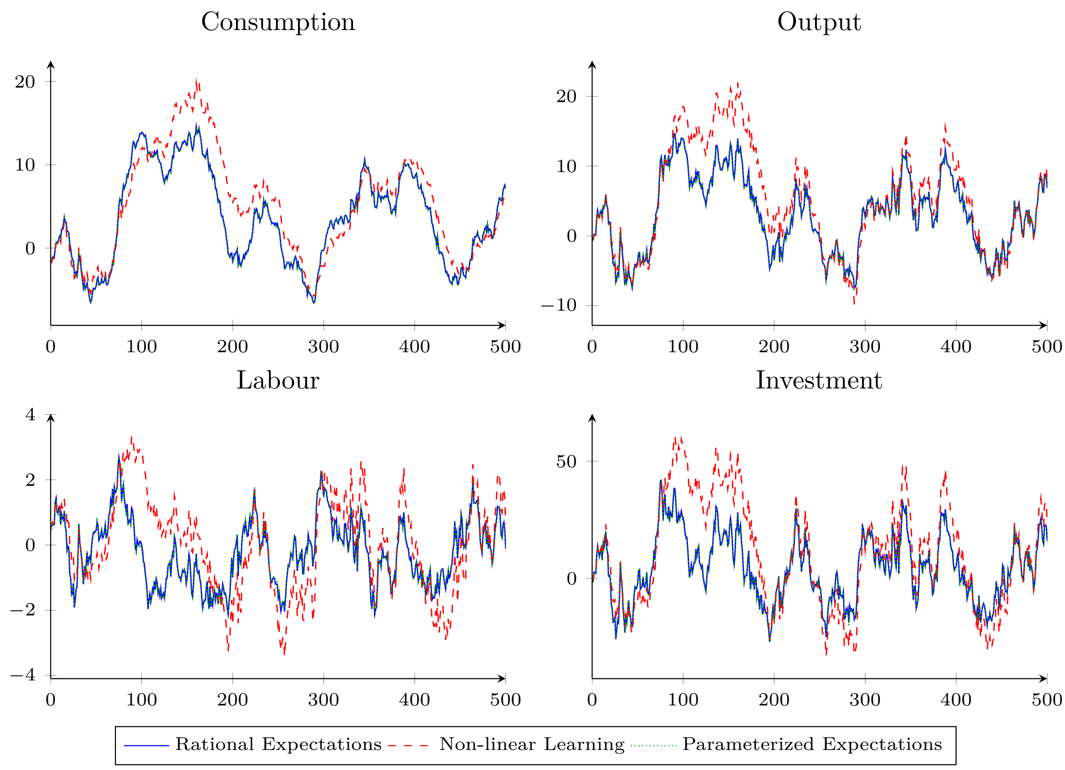

Figure 1: Simulated series of model variables under different assumptions regarding the formation of expectations. The series are measured in percentage deviations from the steady state. The horizontal axis measures time in quarters.

values.[6]

4 The effects of government spending under non-linear learning

In sharp contrast with the rational expectations and PEA solutions of the model, the non-linear learning solution generates time-variation in the transmission of structural shocks in the model economy. In this section we illustrate how this feature provides an attractive mechanism for generating variation in the government spending multipliers over time.

In standard rational expectations models, variation in the transmission mechanism of shocks is typically obtained by allowing structural parameters to vary over time. One common strategy is to estimate a fixed-parameter model on different samples and test for breaks. A downside of this strategy is that the time variation is by assumption infrequent, whereas our approach allows for time variation at a much higher frequency. Another popular strategy is to use stochastic time-varying coefficient models. In this approach, (some) model parameters are assumed to follow a stochastic process.[7] Both strategies have the disadvantage that the time variation is not endogenously determined by the model.

By contrast, our approach allows for endogenous time variation of the government spending multipliers at a relatively high frequency. Figure 3 shows the impact multipliers that correspond to the series in Figure 1. The transmission mechanism that underlies the multipliers changes as agents recursively update their beliefs over time. It is well understood that expectations play a crucial role in the impact of fiscal policies, and our modelling approach highlights this expectations channel for the transmission government spending shocks.

[6]Constant gain learning implies that the parameters do not converge to a point estimate but only to a distribution around the rational expectations beliefs.

[7]Examples of the first strategy are Benati (2008) and Canova (2009); examples of the second strategy include Cogley and Sbordone (2008)

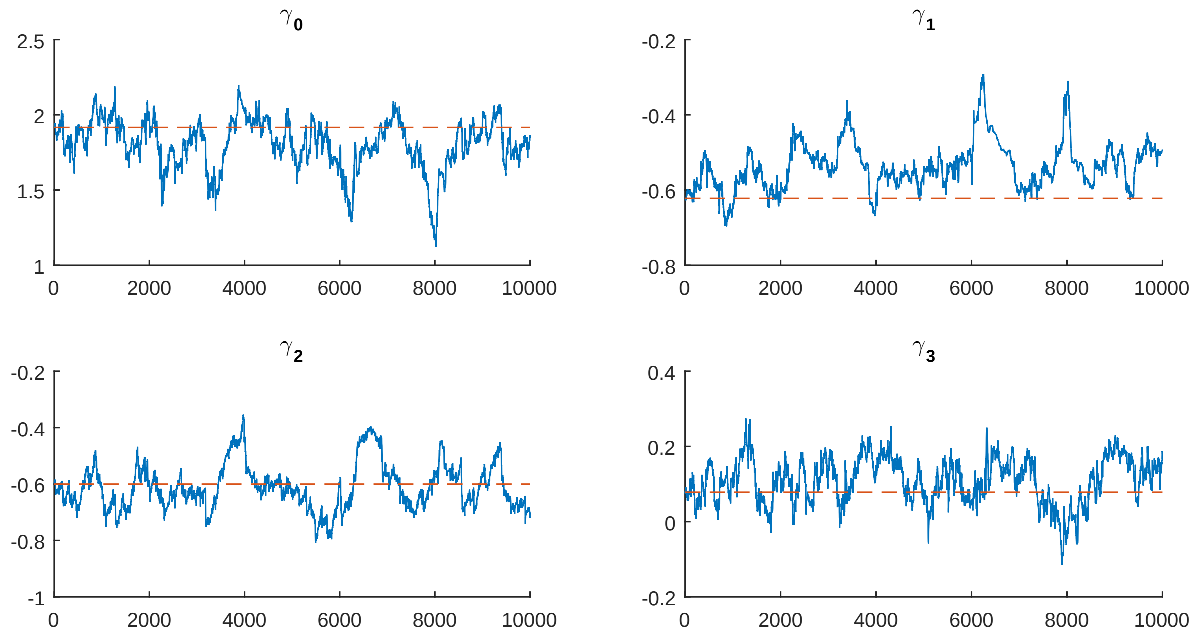

Figure 2: Evolution of the belief coefficients in the non-linear learning model. The series of solid lines correspond to the coefficients of the approximating function given by equation (13). The dashed lines represent the coefficients provided by the parameterized expectations algorithm.

Figure 3: Evolution of the impact multipliers in the non-linear learning model. The series represent the impact multipliers in non-linear learning simulation for the same 500 periods as in Figure 1.

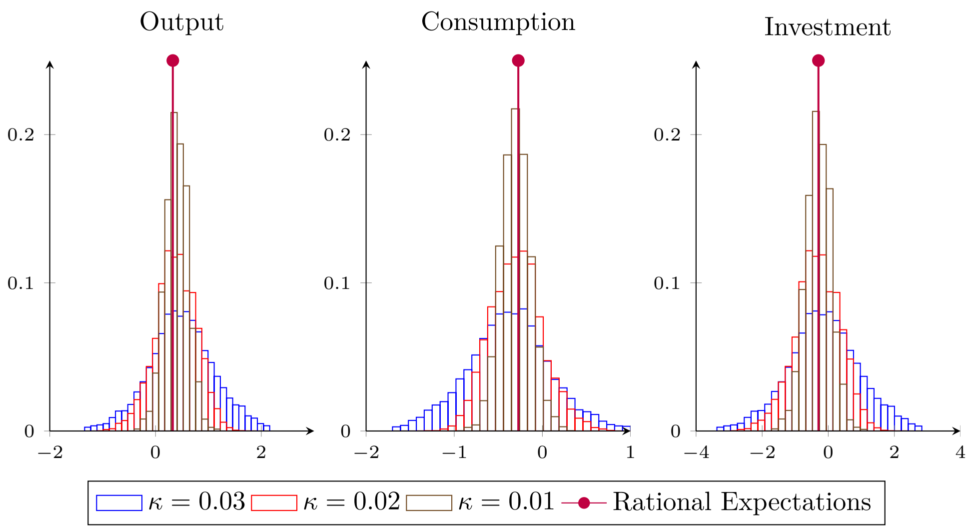

Figure 4: Distribution of impact multipliers based on 10 simulations of the non-linear learning model over 10,000 periods. The histograms represent the government spending multipliers for output, consumption, and investment in the non-linear learning model for different values of the gain parameter (κ). The purple lines represent the rational expectations multipliers. The vertical axis measures relative frequencies.

Figure 4 illustrates the time variation of the government spending multipliers generated by learning behaviour in our model. It reports the distribution of the impact multipliers for output, consumption, and investment.[8] In the rational expectations equilibrium, the multipliers are time-invariant. The multipliers for output, consumption, and investment are, respectively, 0.33, -0.27 and -0.30. The relatively small output multiplier, mainly driven by a significant crowding out of private consumption, is a standard result of the real business cycle model. Under non-linear learning behaviour, however, the multipliers vary significantly around their rational expectations values. The degree of time variation is governed by the gain parameter κ in the recursive learning rule (14) since that parameter determines the volatility of the belief parameters. Overall, learning behaviour in our model generates substantial time variation in the multipliers and a higher gain increases this variation considerably. Figure 5, on the other hand, describes the evolution of the government spending multiplier in a sample of 160 periods. This figure illustrates once more that a higher value for κ leads to more time variation. If κ is equal to 0.01, the time variation is very small. This changes drastically when κ is equal to 0.03. Even within the simple RBC model, the government spending multiplier increases from 0.15 to nearly 0.35 over a period of 40 years.

4.1 Discussion

The previous paragraph showed that our set-up leads to time-variation in the government spending multiplier. This result implies that the effectiveness of an increase in government spending varies over time. For this reason, this paragraph provides some insights into the drivers of the government spending multiplier. These determinants are of particular interest for fiscal policymakers. It allows them to assess the effectiveness of fiscal policy in the economy.

The evolution of the government spending multipliers is mainly driven by the evolution of γ_3, the coefficient for government spending in the approximating function 13. If γ_3 is high (low), the government spending multiplier for output is high (low) as well. Intuitively, this result makes sense. If individuals attach a higher weight to government spending in their approximating function, the impact of an increase in government spending on the decision variables will also be stronger. Figure 6 illustrates this relationship for the output, consumption, and investment multipliers. If γ_3 is high, the crowding out of private consumption after the government spending shock

and Justiniano and Primiceri (2008). A more extensive overview of the existing literature goes beyond the scope of this paper.

[8]The results are based on 10 simulations of the non-linear learning model over 10,000 periods. For each simulation, the initial parameter vector is the vector provided by the parameterized expectations algorithm.

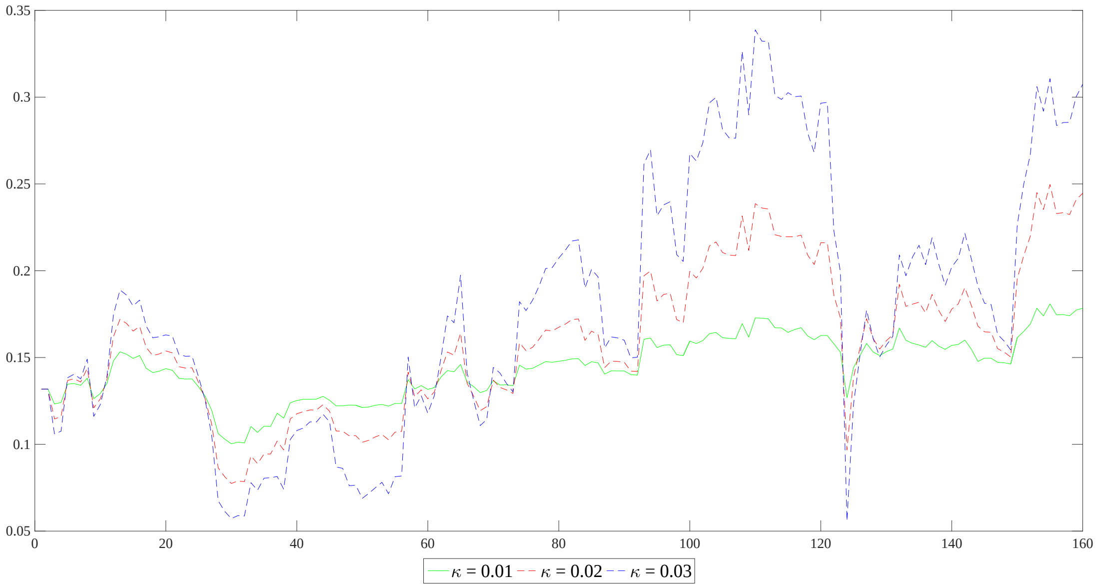

Figure 5: Evolution of the impact multipliers in a sample of 160 periods for different values of κ.

will be more severe. Hence, the consumption multiplier will be larger. Note, however, that in the context of an RBC model, a small multiplier for private consumption leads to a large multiplier for output. The output multiplier is driven by labour supply and a higher level of consumption leads, ceteris paribus, to a lower level of labour supply.

As the evolution of γ_3 is the most important driver of the government spending multiplier, the next step is to explore the drivers of γ_3. Looking at Equation (14), two important channels can be identified. The first one is the forecasting error $\left[\log\left(\phi\left(s_t\right)\right) - \log\left(\psi\left(x_{t-1}, \gamma_{t-2}\right)\right)\right]$. A positive (negative) forecasting error leads to an increase (decrease) in γ_3, everything else equal. The forecasting error itself is determined by the evolution of the technology shocks. More specifically, if $z_t > z_{t-1}$, the forecasting error is negative, whereas if $z_t < z_{t-1}$ it is positive. So, everything else equal, after a period of technological growth, the government spending multiplier is low and after a period of decreasing technology, the government spending multiplier is high. This result is, however, dependent on the level of government spending, which is the second channel. If government spending is above its steady state level, γ_3 will increase for a given level of technology, if it is below its steady state level, γ_3 will decrease for a given level of technology. Furthermore, the further government spending is from its steady state level, the stronger its influence on the evolution of γ_3 and thus the government spending multiplier.

Combined, these results lead to the following insights. The government spending multiplier is likely to be high if technology and government spending relative to its steady state level evolve in opposite direction. More specifically, the multiplier is high (increases) after (i) a considerable period of increasing technology and low government spending on the one hand, and (ii) after a considerable period of decreasing technology and high government spending on the other hand. The multiplier is low (i) after a considerable period of increasing technology and high government spending on the one hand and (ii) after a considerable period of decreasing technology and a low level of government spending on the other hand[9].

We illustrate the aforementioned results using Figure 7. In this Figure, we highlight three different episodes during which the government spending multiplier strongly changes. The corresponding levels of technology and government spending are displayed in respectively the middle and bottom panel. Furthermore, the red line in the bottom panel denotes the steady state level of government spending. Here, we discuss the evolution of the

[9]Recent contributions to the literature investigate the relationship between the size of the multiplier and the state of the business cycle – see e.g. Auerbach and Gorodnichenko (2012, 2013) and Owyang et al. (2013). We have investigated this relationship using different measures for the business cycle and found that this stylized RBC model is not adequate to uncover a structural relation. Investigating this issue in a more elaborate (more demand-driven) non-linear learning model is a promising direction for future research.

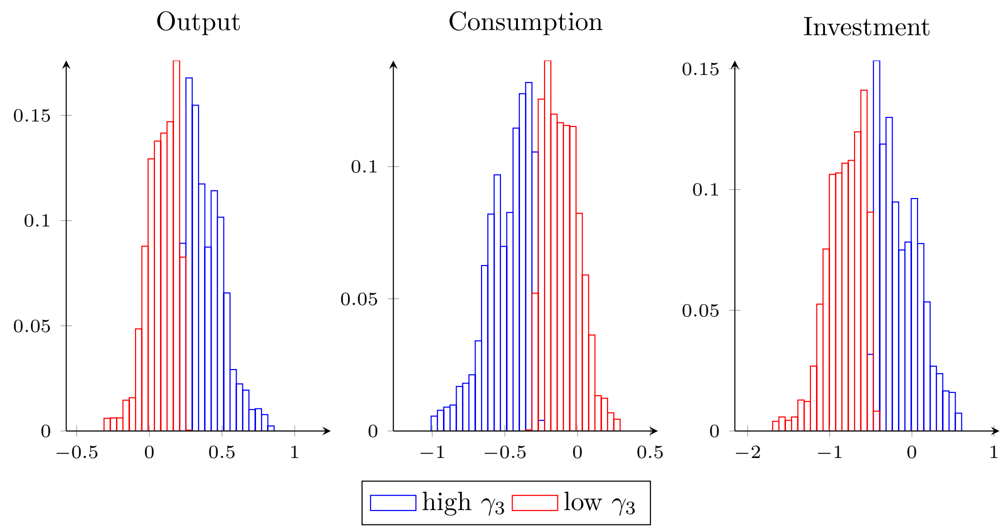

Figure 6: Distribution of impact multipliers for high and low values of belief parameter γ_3 . The histograms are based on a simulation of 10,000 periods of the non-linear learning model and represent the government spending multipliers for output, consumption, and investment. The high (low) γ_3 sub-sample contains the multipliers when γ_3 is above (below) its median value. The vertical axis measures relative frequencies.

government spending multiplier between periods 176 and 204. The other ones follow the same reasoning. At first, technology strongly increases while the level of government spending is below its steady state level. We know that this combination leads to an increase in γ_3. This increase almost directly leads to an increase in the government spending multiplier. After some time, however, technology starts to decrease. At that point, government spending is higher than the steady state level. This combination pushes the government spending multiplier upwards.

4.2 Choice of the approximating function: linear versus non-linear learning

In the baseline simulation, we assume that agents approximate the expectation in the consumption Euler equation (cf. equation (12)) by the following non-linear function

$$\psi\left(x_t, \gamma_{t-1}\right) = \exp\left[\gamma_0 + \gamma_1 \log k_t + \gamma_2 \log z_t + \gamma_3 \log g_t\right].$$

In this section, we compare this "non-linear" learning approach with "linear learning" where agents use the following linear approximating function

$$\psi\left(x_t, \gamma_{t-1}\right) = \gamma_0 + \gamma_1 k_t + \gamma_2 z_t + \gamma_3 g_t.$$

As for the "non-linear" learning case, the parameters γ_{t-1} are updated using the constant-gain recursive least squares formulas (14)–(15) but the updating term between square brackets in equation (14) now becomes $[\phi\left(s_t\right) -\psi\left(x_{t-1}, \gamma_{t-2}\right)]$.

When comparing the linear with the non-linear learning approach, two interesting observations can be made. First, non-linear learning outperforms linear learning in terms of forecasting performance, especially when (large) shocks drive the economy far away from its steady state. To illustrate this, Table 2 reports the root-mean square error (RMSE) of the forecasts under linear and non-linear learning. In a "high volatility" simulation, when structural shocks to the economy are larger than in the baseline simulation, non-linear learning improves the forecasting performance relative to linear learning. In the baseline simulation, however – when the structural shocks are relatively small – the forecasting performance of non-linear learning is (approximately) equal to linear learning. Hence, when shocks are relatively large, non-linear learning leads to better forecasting. To provide some intuition for this

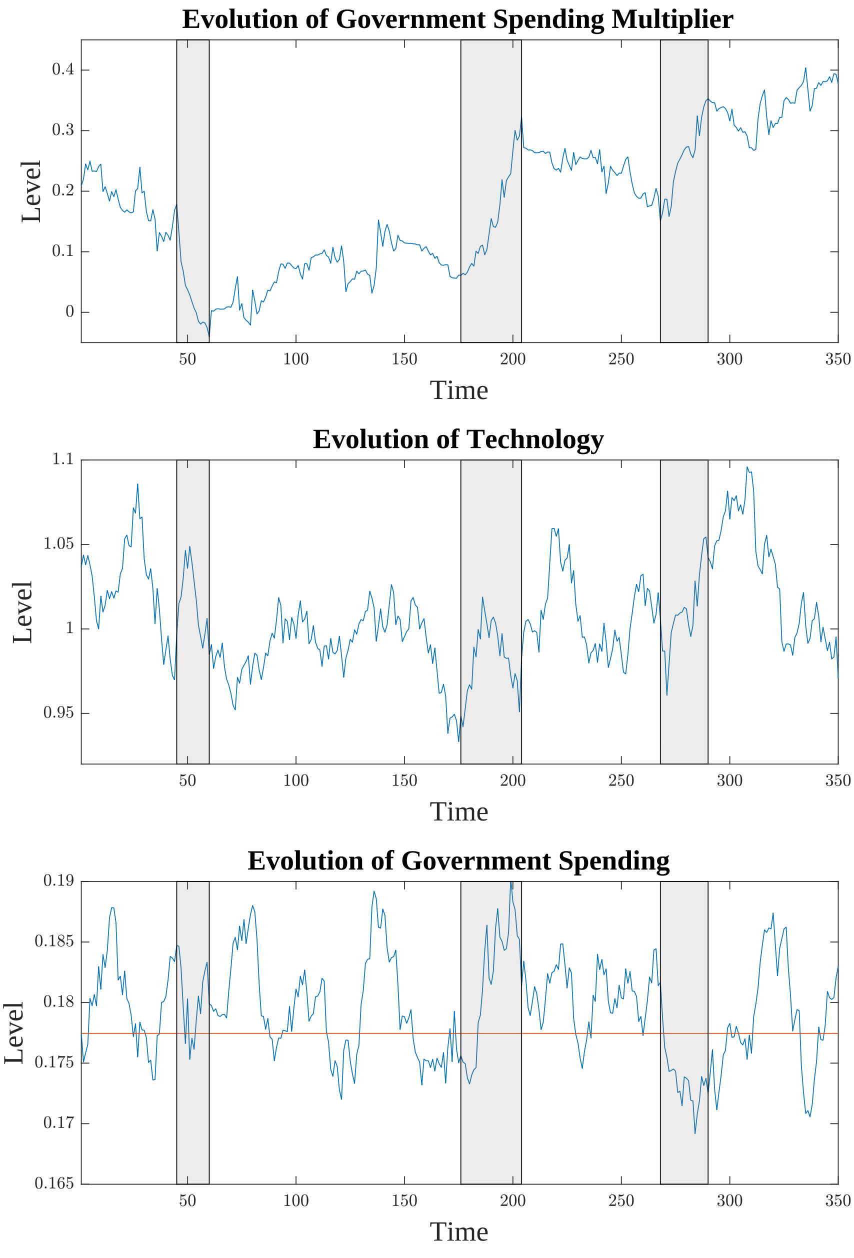

Figure 7: Endogenous drivers of the government spending multiplier

Learning scheme	Root-mean-square error	
	Baseline ($\sigma_Z = \sigma_G = 0.01$)	High volatility ($\sigma_Z = \sigma_G = 0.025$)
Non-linear learning	0.010	0.026
Linear learning	0.010	0.034

Table 2: Forecasting performance of linear and non-linear learning. Root-mean-squared errors are calculated over 10,000 simulation periods.

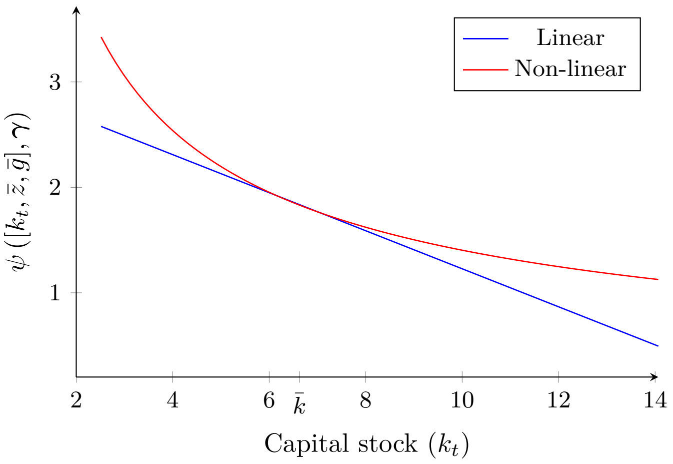

Figure 8: Linear and non-linear approximating function for different values of the capital stock. Government spending and technology are fixed at their steady state values and γ equals the coefficient vector of the parameterized expectations algorithm.

result, Figure 8 plots the linear and non-linear approximating function for different values of the capital stock. For values close to the steady state ($k = \bar{k}$), the approximation by the linear and non-linear function is very similar. For capital (and government spending) far away from the steady state, the linear and non-linear approximation differ quite substantially. In both cases, agents can update the coefficients of the approximating function to improve their forecasting, but the results presented in Table 2 show that non-linear learning outperforms linear learning in terms of RMSE.

Second, the linear learning approach also generates time variation in the government spending multipliers. Figure 9 shows how the impact multipliers for output, consumption, and investment again fluctuate around the rational expectations multipliers, but their values are less dispersed. Under non-linear learning, the updating of the belief coefficients generate more variation in the consumption response after the government spending shock. Consequently, the effects on output and investment will also vary more.

5 Conclusion

In this paper, we explore the transitional dynamics following fiscal policy changes in a stochastic macroeconomic framework where agents use adaptive learning to update non-linear forecast rules to form expectations. Several papers have already studied the effect of fiscal policy using a learning framework (see e.g. Evans and Honkapohja, 2009; Gasteiger and Zhang, 2014; Benhabib et al., 2014). All of these papers, however, use linear forecast rules. In this paper, we apply a different approach. Following, inter alia, Marcet and Marshall (1994) and Berardi and Duffy (2015), we interpret the parameterized expectations algorithm (PEA) as a real-time learning mechanism used by agents to update their expectations over time in the original non-linear model.

Our main contribution is to study the dynamics of government spending in the standard RBC model where

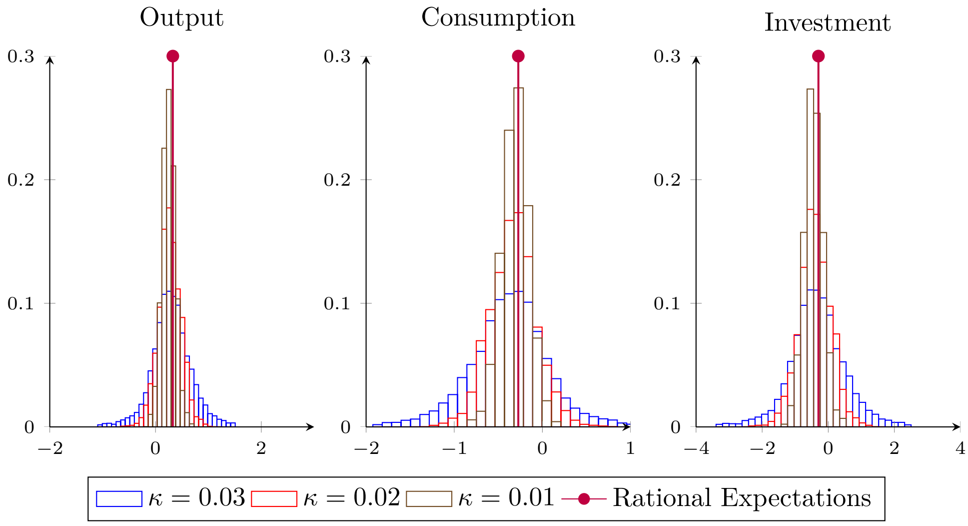

Figure 9: Distribution of impact multipliers based on 10 simulations over 10,000 periods of the model with a linear approximating function. The histograms represent the government spending multipliers for output, consumption, and investment for different values of the gain parameter (κ). The purple lines represent the rational expectations multipliers. The vertical axis measures relative frequencies.

agents use this non-linear learning mechanism to form expectations. To the best of our knowledge, ours is the first study to compare the transitional dynamics resulting from this framework with the dynamics under rational expectations.

In our non-linear learning set-up, the effects of government spending shocks vary substantially over time. We have shown that this variation is endogenously driven by agents' expectations about the future. The resulting fluctuations in the government spending multipliers are in marked contrast to the time-invariant multipliers of the rational expectations solution of the model. We also show that the forecasting performance of the learning mechanism is better if agents use a non-linear approximating function instead of a linear one. In particular in the context of relatively high structural shocks which drive the economy far away from its steady state, using a non-linear approximating function to form expectations is advantageous.

Our findings provide several avenues for future research. First, it may be very fruitful to go beyond the standard RBC model. Understanding how learning affects the variability of fiscal multipliers in more elaborate models is an important topic. It may, for example, shed further light on the dependency of multipliers on the state of the business cycle and the stance of monetary policy. Second, the non-linear learning model provides a natural framework for studying the consequences of a structural change in fiscal policy. Comparing the dynamics of this learning model with the rational expectations dynamics is, in our view, a promising direction for further research.

Acknowledgements

We thank Freddy Heylen, Gert Peersman, Punnoose Jacob, Céline Poilly, Raf Wouters, and the members of the Macro Research Group of Ghent University for valuable suggestions. Brecht Boone is a Ph.D.-Fellow of the Research Foundation - Flanders and acknowledges support from the Foundation.

References

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Auerbach, A. J. and Gorodnichenko, Y. (2013). Fiscal Multipliers in Recession and Expansion. In *Fiscal Policy after the Financial Crisis*, NBER Chapters, pages 63–98. National Bureau of Economic Research, Inc.

Benati, L. (2008). Investigating inflation persistence across monetary regimes. *The Quarterly Journal of Economics*, 123(3):1005–1060.

Benhabib, J., Evans, G. W., and Honkapohja, S. (2014). Liquidity traps and expectation dynamics: Fiscal stimulus or fiscal austerity? *Journal of Economic Dynamics and Control*, 45:220 – 238.

Berardi, M. and Duffy, J. (2015). Real-Time, Adaptive Learning Via Parameterized Expectations. *Macroeconomic Dynamics*, 19(02):245–269.

Branch, W. and Evans, G. W. (2007). Model uncertainty and endogenous volatility. *Review of Economic Dynamics*, 10(2):207–237.

Branch, W. A. and Evans, G. W. (2006). A simple recursive forecasting model. *Economics Letters*, 91(2):158–166.

Canova, F. (2009). What Explains The Great Moderation in the U.S.? A Structural Analysis. *Journal of the European Economic Association*, 7(4):697–721.

Cho, I.-K., Williams, N., and Sargent, T. J. (2002). Escaping Nash Inflation. *Review of Economic Studies*, 69(1):1–40.

Cogley, T. and Sbordone, A. M. (2008). Trend inflation, indexation, and inflation persistence in the new keynesian phillips curve. *American Economic Review*, 98(5):2101–26.

den Haan, W. J. and Marcet, A. (1990). Solving the Stochastic Growth Model by Parameterizing Expectations. *Journal of Business & Economic Statistics*, 8(1):31–34.

Eusepi, S. and Preston, B. (2011). Expectations, learning, and business cycle fluctuations. *American Economic Review*, 101(6):2844–72.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

Evans, G. W. and Honkapohja, S. (2009). Learning and macroeconomics. *Annual Review of Economics*, 1(1):421–451.

Evans, G. W., Honkapohja, S., and Mitra, K. (2009). Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics*, 56(7):930–953.

Gasteiger, E. and Zhang, S. (2014). Anticipation, learning and welfare: the case of distortionary taxation. *Journal of Economic Dynamics and Control*, 39:113 – 126.

Honkapohja, S., Mitra, K., and Evans, G. W. (2013). Notes on agents' behavioral rules under adaptive learning and studies of monetary policy. In Sargent, T. and Vilmunen, J., editors, *Macroeconomics at the Service of Public Policy*, chapter 4, pages 63–79. Oxford University Press, Oxford, United Kingdom.

Justiniano, A. and Primiceri, G. E. (2008). The time-varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–41.

Marcet, A. and Lorenzoni, G. (1999). Computational Methods for the Study of Dynamic Economies. Number 9780198294979 in OUP Catalogue, chapter The Parameterized Expectations Approach Some Practical Issues, pages 143–171. Oxford University Press.

Marcet, A. and Marshall, D. A. (1994). Solving nonlinear rational expectations models by parameterized expectations: Convergence to stationary solutions. *Federal Reserve Bank of Minneapolis Discussion Paper*, (91).

Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics*, 54(7):2065–2082.

Mitra, K., Evans, G. W., and Honkapohja, S. (2013). Policy change and learning in the RBC model. *Journal of Economic Dynamics and Control*, 37(10):1947–1971.

Orphanides, A. and Williams, J. C. (2005). The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control*, 29(11):1927–1950.

Orphanides, A. and Williams, J. C. (2007). Robust monetary policy with imperfect knowledge. *Journal of Monetary Economics*, 54(5):1406–1435.

Owyang, M. T., Ramey, V. A., and Zubairy, S. (2013). Are Government Spending Multipliers Greater during Periods of Slack? Evidence from Twentieth-Century Historical Data. *American Economic Review*, 103(3):129–34.

Preston, B. (2005). Learning about Monetary Policy Rules when Long-Horizon Expectations Matter. *International Journal of Central Banking*, 1(2).

Rogerson, R. (2007). Taxation and market work: is Scandinavia an outlier? *Economic Theory*, 32(1):59–85.

Sargent, T. J. (1999). *The Conquest of American Inflation*. Princeton University Press.

Slobodyan, S. and Wouters, R. (2012). Learning in an estimated medium-scale DSGE model. *Journal of Economic Dynamics and Control*, 36(1):26–46.

Appendices

A Parameterized expectations algorithm

Under the parameterized expectations algorithm (PEA) the vector of coefficients γ of the approximating function $\psi(\cdot,\gamma)$ is calculated with the following iterative procedure.

Step 1 Draw a large sequence of shocks $\left\{\varepsilon_t^G, \varepsilon_t^Z\right\}_{t=0}^T$ and compute $\{g_t, z_t\}_{t=0}^T$ as defined in (6) and (11).

Step 2 Choose an initial guess γ_0 and an initial value k_0 for the capital stock.

Step 3 At iteration $i \in \{0, \ldots, i_{max}\}$, use γ_i to generate the endogenous variables $\{c_t(\gamma_i), n_t(\gamma_i), k_{t+1}(\gamma_i)\}_{t=0}^T$. In particular, (i) substituting the approximating function $\psi(x_t,\gamma_i)$ for the expectation function in equation (4) gives consumption $c_t = \beta\left[\psi(x_t,\gamma_i)\right]^{-1/\sigma}$, (ii) the first-order condition (3) determines labour supply n_t, and (iii) k_{t+1} follows from the resource constraint (2).

Step 4 Use the data for $t = 0,\ldots,T-1$ to run the non-linear least squares regression of $c_{t+1}^{-\sigma}\left(1-\delta+\alpha z_{t+1}k_{t+1}^{\alpha-1}n_{t+1}^{1-\alpha}\right)$ on the approximating function $\psi(x_t,\cdot)$ to obtain an estimate $\hat{\gamma}$.

Step 5 Use this estimate to update the guess for γ according to

$$\gamma_{i+1} = (1-\mu)\gamma_i + \mu\hat{\gamma}, \tag{16}$$

where $(0,1]$ is a damping parameter.

Step 6 Apply steps 3-5 iteratively until the convergence criterion $\sum_k \left|\gamma_{i+1}^k - \hat{\gamma}^k\right| < \tau$ is met, where k is the number of parameters in γ.

In our experiments, we set $T = 10{,}000$, $i_{max} = 200$, $\mu = 0.8$, and $\tau = 1\times10^{-5}$. The initial value for the capital stock is the steady state: $k_0 = \bar{k}$.

B Learning algorithm

To simulate the learning model, we follow the following steps.

Initialisation

1. Draw a sequence of shocks $\left\{\varepsilon_t^G, \varepsilon_t^Z\right\}_{t=0}^S$ and compute $\{g_t, z_t\}_{t=0}^S$ as defined in (6) and (11).
2. Choose an initial guess γ_0 and an initial value k_0 for the capital stock.

Simulation Simulate the model S periods forward using the following scheme.

1. Approximate the conditional expectational function in the consumption Euler equation

$$E_t\phi(s_{t+1}) = E_t\left[c_{t+1}^{-\sigma}\left(1-\delta+\alpha z_{t+1}k_{t+1}^{\alpha-1}n_{t+1}^{1-\alpha}\right)\right], \tag{17}$$

by the parametric function $\psi\left(x_t, \gamma_{t-1}\right)$ of the state variables $x_t = [1, k_t, z_t, g_t]$.

2. Calculate the corresponding endogenous variables determined by the following system

$$c_t = \beta \left[\psi\left(x_t, \gamma_{t-1}\right)\right]^{-1/\sigma}, \tag{18}$$

$$c_t^{-\sigma} = \frac{b\left(1-n_t\right)^{-\theta}}{w_t}, \tag{19}$$

$$w_t = (1-\alpha) z_t k_t^{\alpha} n_t^{-\alpha} \tag{20}$$

$$y_t = z_t k_t^{\alpha} n_t^{1-\alpha}, \tag{21}$$

$$i_t = y_t - c_t - g_t, \tag{22}$$

$$k_{t+1} = (1-\delta) k_t + i_t. \tag{23}$$

3. Update the vector of parameters γ_{t-1} according to

$$\gamma_t = \gamma_{t-1} + \kappa R_t^{-1} x_{t-1} \left[\log\left(\phi\left(s_t\right)\right) - \log\left(\psi\left(x_{t-1}, \gamma_{t-2}\right)\right)\right], \tag{24}$$

$$R_t = R_{t-1} + \kappa \left[x_{t-1} x_{t-1}^{'} - R_{t-1}\right]. \tag{25}$$