



Working Paper Series



This paper can be downloaded without charge from:
<http://www.richmondfed.org/publications/>



THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ■ BALTIMORE ■ CHARLOTTE

Learning about Fiscal Policy and the Effects of Policy Uncertainty

Josef Hollmayr* and Christian Matthes**

September 30, 2013
Working Paper No. 13-15

Abstract

The recent crisis in the United States has often been associated with substantial amounts of *policy uncertainty*. In this paper we ask how uncertainty about fiscal policy affects the impact of fiscal policy changes on the economy when the government tries to counteract a deep recession. The agents in our model act as econometricians by estimating the policy rules for the different fiscal policy instruments, which include distortionary tax rates.

Comparing the outcomes in our model to those under full-information rational expectations, we find that assuming that agents are not instantaneously aware of the new fiscal policy regime (or policy rule) in place leads to substantially more volatility in the short run and persistent differences in average outcomes.

JEL codes: E32, D83, E62

Keywords: DSGE, FISCAL POLICY, LEARNING

*Deutsche Bundesbank, Frankfurt am Main (e-mail: josef.hollmayr@bundesbank.de).

**Federal Reserve Bank of Richmond (email: christian.matthes@rich.frb.org). The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bundesbank, the Federal Reserve Bank of Richmond, or the Federal Reserve System. We would like to thank Klaus Adam, Tim Cogley, Michael Krause, Geert Langenus, Stéphane Moyen, Albert Marcet, Taisuke Nakata, Johannes Pfeifer, Ctirad Slavik, and Nikolai Stähler, seminar participants at the University of Hamburg and the Bundesbank as well workshop participants in Mannheim and Bucharest for helpful comments. This project started while Matthes was visiting the Bundesbank, whose hospitality is gratefully acknowledged.

1 Introduction

Partly motivated by the recent financial crisis and the subsequent recession, economists have recently placed greater emphasis on identifying uncertainty about monetary and fiscal policy as a potentially important factor determining economic outcomes, as highlighted by Baker *et al.* (2012). Natural questions seem to be how this uncertainty arises, what the exact transmission mechanism is and how this uncertainty affects equilibrium outcomes. In this paper we propose one model of fiscal policy uncertainty: an RBC-type model with distortionary taxation and government debt, in which agents act as econometricians and update their beliefs about fiscal policy every period. In our model, agents use past realizations of fiscal variables to learn what actual policy rules are in place and thus whether changes in those fiscal variables are temporary (driven by exogenous shocks) or permanent (driven by changes in the parameters of the fiscal policy rules). In our model uncertainty about fiscal policy is partly endogenous since the properties of the estimators of the fiscal policy rule coefficients employed by private agents change as the private sector's behavior changes. This behavior occurs because choice variables of the representative private agent enter the fiscal policy rules.

The task of disentangling permanent from temporary changes in fiscal policy is identified as a major source of fiscal policy uncertainty by Baker *et al.* (2012), who use an index of tax code expiration data to measure fiscal policy uncertainty¹. Figure 1 plots their index of fiscal uncertainty. Uncertainty increases substantially during the recent American Recovery and Reinvestment Act (ARRA) program, a major period of policy change, but it is very small beforehand and decreases afterward. We will use these patterns in this measure of *objective* uncertainty to inform our (learning) model of *subjective* uncertainty².

We analyze a one-time permanent change in the government spending policy rule and use Monte Carlo simulations of our model to assess how beliefs evolve and how these beliefs affect allocations. Learning leads to substantially different outcomes even though learning is quite fast: There is a substantial temporary spike in volatility under learning that is absent under full information. In addition, there are persistent average differences between the outcomes under learning and under full information. We show that investment plays a big role in creating the average differences - temporary dif-

¹They state on the associated website www.policyuncertainty.com that "Temporary tax measures are a source of uncertainty for businesses and households because Congress often extends them at the last minute, undermining stability in and certainty about the tax code."

²The subjective measure of fiscal policy uncertainty used in Baker *et al.* (2012), a measure of disagreement among professional forecasts of fiscal spending, shows a similar pattern around the introduction of the ARRA program.

ferences in investment between the learning and full information environments have long-lasting effects via the capital stock. The uncertainty about government spending induces uncertainty about the steady state of other variables such as GDP and debt, which in turn influences uncertainty about the steady state of other fiscal policy instruments, even though the coefficients of those policy rules are tightly (and correctly) estimated. Thus, even though we only consider changing a small subset of the fiscal policy coefficients, this uncertainty creeps into other fiscal variables.

To check for robustness, we consider various assumptions about the agents' information set and their preferences as well as an alternative change in fiscal policy. Our qualitative results remain unchanged throughout.

We are far from being the first to model fiscal policy in an environment in which agents adaptively learn about the economy. Papers such as that of Eusepi and Preston (2011) encompass both monetary and fiscal policy, but have a smaller set of fiscal policy instruments (in particular no distortionary taxation). We instead choose to focus on fiscal policy alone, leaving the interesting issue of fiscal and monetary policy interaction for future work. We do, however, have a larger set of fiscal policy instruments. Giannitsarou (2006) does feature distortionary taxation and is interested in issues similar to ours, but does not feature government debt, which we include in order to be able to view the current policy debate in the United States through the lens of our model. Mitra *et al.* (2012) focus on the question of anticipated versus unanticipated changes in fiscal policy when agents are learning, but they only study the case of lump-sum taxation.

What sets our model apart is the way agents form their beliefs about the stance of fiscal policy. In contrast to the previously mentioned papers, our agents know the structure of the economy, except for the policy rules followed by the fiscal authority. Our paper thus follows the approach laid out in Cogley *et al.* (2011), who study a model of monetary policy. Firms and households in our model estimate the coefficients of the policy rules using the Kalman filter and incorporate both these beliefs and all cross-equation restrictions coming from knowledge of the structure of the economy when making their decisions. Knowledge of the timing of the policy change is incorporated by agents into their estimation problem by using a time-varying covariance matrix for the parameters. Furthermore, the agents in our model are aware that the government budget constraint has to hold. Thus they estimate policy rules for all but one fiscal policy instrument, with the beliefs about the last policy instrument being determined by the period-by-period government budget constraint.

In our model, agents are uncertain not only about future fiscal policy, but also about the policy rules currently in place. Papers such as Davig *et al.* (2010) instead model

the fiscal policy rule coefficients as being governed by a discrete state Markov chain, which is observable to private agents. Thus agents in those environments know the policy rule coefficients in place in the current period.

Another strand of the literature that studies uncertainty³ (or risk) about future fiscal policy is represented by Born and Pfeifer (2013) and Fernandez-Villaverde *et al.* (2011), who study stochastic volatility in the innovations of otherwise standard fiscal policy rules. The view of uncertainty encoded in the latter two papers is quite different from both our approach as well as the approach used by Davig *et al.* (2010) and similar papers: In our model, agents are uncertain as to how the government systematically sets its fiscal policy instruments (both currently and in the future), whereas in Born and Pfeifer (2013) and Fernandez-Villaverde *et al.* (2011) agents are uncertain as to how important (i.e. volatile) the random component of fiscal policy will be in the future. Davig *et al.* (2010), Born and Pfeifer (2013) and Fernandez-Villaverde *et al.* (2011) use full-information rational expectations models, whereas our approach encodes a form of bounded rationality common in the learning literature (the anticipated utility approach of Kreps (1998)), which sets the approaches further apart. The anticipated utility approach we use abstracts from precautionary behavior driven by model uncertainty on behalf of the private agents. Our results can thus be viewed as representing a bound on the difference between full information and limited information approaches.

2 Model

Our model is a simplified version of Leeper *et al.* (2010). It is a real model of a closed economy without habits and other frictions. The only deviation from the simplest possible RBC model (Kydland and Prescott (1982)) is the rich fiscal sector with distortionary taxation, government spending, and transfers. First-order conditions and the complete log-linearized model may be found in the Appendix.

³When we talk about uncertainty, we do *not* mean Knightian uncertainty. For a study of (optimal) fiscal policy when agents face Knightian uncertainty, see Karantounias (2013). A corresponding analysis of optimal fiscal policy when agents are learning is provided by Caprioli (2010). Both papers use a smaller set of fiscal policy instruments than we do.

2.1 Households

Households are expected utility maximizers⁴ with the instantaneous utility function of the representative household taking the following form:

$$U_t = \frac{C_t^{1-\sigma}}{1-\sigma} - \frac{L_t^{1+\phi}}{1+\phi} \quad (1)$$

with consumption C_t and labor L_t . Each period households can choose to consume, save in the form of government bonds (B_t) or invest (I_t) in the capital stock (K_t) that they hold. Therefore they maximize the infinite sum of discounted utility under the following constraints:

$$C_t(1 + \tau_t^C) + B_t + I_t = W_t L_t(1 - \tau_t^L) + (1 - \tau_t^K)R_t^K K_{t-1} + R_{t-1}B_{t-1} + Z_t \quad (2)$$

$$K_t = (1 - \delta)K_{t-1} + I_t \quad (3)$$

where the first constraint is the budget constraint of the household and the latter is the law of motion for capital. The household's income stems from working at the wage W_t , gains from renting out capital R_t^K and interest payments on their savings at the rate R_t . Z_t represents lump-sum transfers or taxes. τ_t^i with $i = K, L, C$ denotes the various tax rates that the government levies on capital, labor and consumption.

2.2 Firms

The production function is of the standard Cobb-Douglas type:

$$Y_t = \exp(A_t)K_{t-1}^\alpha L_t^{1-\alpha} \quad (4)$$

where Y_t denotes the output produced with a certain level of technology A_t , capital K_t and labor L_t . Technology follows an AR(1) process. The exogenous process for technology is an AR(1):

$$A_t = \rho_a A_{t-1} + \epsilon_t^A \quad (5)$$

⁴This statement extends up to their beliefs about fiscal policy rule coefficients, which they treat as fixed when making their decisions. We use an *anticipated utility* assumption, which is common in the literature on adaptive learning. It is described in detail in the section that elaborates on our learning algorithm.

2.3 Government

The government in this setup only consists of the fiscal branch. The government budget constraint is given by:

$$B_t = B_{t-1}R_{t-1} - R_t^K K_t \tau_t^K - W_t L_t \tau_t^L - C_t \tau_t^C + G_t + Z_t \quad (6)$$

We follow Leeper *et al.* (2010) in the choice of right-hand side variables for the policy rules, except that we make time t fiscal policy instruments functions of time $t - 1$ endogenous variables. This assumption simplifies our learning algorithm, which we discuss later. Given the lags in fiscal policy decision-making, this assumption does not seem overly strong⁵.

Government Expenditure:

$$\log(G_t) = G_c - \rho_{g,y} \log(Y_{t-1}) - \rho_{g,b} \log(B_{t-1}) + \epsilon_t^G \quad (7)$$

Transfers:

$$\log(Z_t) = Z_c - \rho_{z,y} \log(Y_{t-1}) - \rho_{z,b} \log(B_{t-1}) + \epsilon_t^Z, \quad (8)$$

Consumption Tax Rate Rule:

$$\log(\tau_t^C) = \tau_c^c + \epsilon_t^C \quad (9)$$

Labor Tax Rate Rule:

$$\log(\tau_t^L) = \tau_c^l + \rho_{L,y} \log(Y_{t-1}) + \rho_{L,b} \log(B_{t-1}) + \epsilon_t^L \quad (10)$$

Capital Tax Rate Rule:

$$\log(\tau_t^K) = \tau_c^k + \rho_{K,y} \log(Y_{t-1}) + \rho_{K,b} \log(B_{t-1}) + \epsilon_t^K \quad (11)$$

In contrast to Leeper *et al.* (2010) we simplify the model and do not assume that the innovations to the tax rates are contemporaneously correlated.

The firms and households in our model know the form of the policy rules described above, but they do not know the coefficients, which they have to estimate. They also know that the government budget constraint has to hold in every period.

⁵For a discussion of the link between simple fiscal policy rules like the ones employed here and optimal fiscal policy, see Kliem and Kriwoluzky (2013).

2.4 Market Clearing

Demand on the part of the government and households must fully absorb the output of the competitive firm:

$$Y_t = C_t + I_t + G_t$$

The bond market in our model is simple and market clearing in this market implies that all bonds issued by the government are bought by the households in the economy.

3 Calibration

The model is calibrated to the U.S. economy at a quarterly frequency. All parameters of the model are chosen to be consistent with other dynamic stochastic general equilibrium models in the literature. Therefore, the discount factor, the parameter which indicates the impatience of households, β , is set to 0.99. This value yields a steady state real interest rate of 3.6 percent in annual terms. The capital share in the Cobb-Douglas function α is one-third ⁶ and the depreciation rate of capital is set at 0.025, which is equivalent to a total annual depreciation of 10 percent. These values are in line with accounting standards. The CES parameters σ and ϕ govern the utility function, which takes as its input consumption and labor. Both parameters are fixed at 2.

Lastly, all coefficients in the fiscal rules come from the estimation of the DSGE model in Leeper *et al.* (2010). Although their model includes more frictions such as consumption habits and a capital utilization rate, we think that it is reasonable to adopt their estimation results for these parameters.

To obtain the same steady state values as Leeper *et al.* (2010) for tax rates, government spending over GDP, and debt capital over GDP, we set the respective constants accordingly. The steady state values for the consumption tax, the capital tax, and the labor tax are therefore 0.0287, 0.2452, and 0.1886, respectively. The ratio for the shares of government spending and capital to GDP are 0.09 and 7.10. The volatilities of all shock processes are also taken from the estimation in Leeper *et al.* (2010). We discuss the parameters governing initial beliefs and learning when we present the learning algorithm in the next section.

All parameters and steady state values are shown in tables 1 and 2, respectively.

⁶This value is within the band that is implied by the prior mean by Smets and Wouters (2007)(0.3) and the calibrated parameter by Bernanke *et al.* (1999) (0.35)

4 A Change in Fiscal Policy

We want to ask how beliefs and economic outcomes evolve during a recession when fiscal policy acts to counteract the recession. This section lays out the main policy experiment we consider. As initial values for the policy rule coefficients we use the estimates from Leeper *et al.* (2010), which we reproduce in table 1. The analysis is carried out via a Monte Carlo simulation - 1000 simulations of 100 periods each. In period 9, a negative technology shock hits that puts the technology level 5 percent below its steady state level. In the next period, the fiscal policy authority changes the process for government spending. We consider a permanent policy change in which only the intercept in the policy rule changes to reflect an average increase of government spending across the board. All other coefficients of the fiscal policy rules remain fixed at the original levels (including the intercepts in the respective policy rules)⁷.

We pick the size of the change in G_c using the following thought experiment: Given the original steady state values for debt and GDP, by how much would we have to change G_c to increase the steady state level of government spending by 1 percent of GDP? The '1 percent of GDP' number is in line with the maximum increase in G_t used by Cogan *et al.* (2010), who calibrate their G_t sequence to the ARRA spending program. To illustrate our choice of the change in G_c , it is useful to look at equation (7) in levels at the original steady state:

$$G = \exp(G_c)Y^{-\rho_{g,y}}B^{-\rho_{g,b}} \quad (12)$$

Uppercase letters without a subscript denote the original steady state in this case. We solve for the new value of the intercept in the log version of the government spending rule G_c^* using the following equation:

$$G + 0.01Y = \exp(G_c^*)Y^{-\rho_{g,y}}B^{-\rho_{g,b}} \quad (13)$$

This is a back-of-the-envelope calculation since it does not take into account that a change in G_c will affect the steady state values of GDP and debt, and thus it will not lead to an increase of 1 percent of GDP. In our benchmark case the actual increase in G due to this policy change is 0.81 percent of original GDP, so the back-of-the-envelope calculation is not far off. We use this calculation because it is a calculation a

⁷This implies that we do not change how the government raises revenues - the way government spending is paid for is still encoded in the policy rule coefficients we have borrowed from Leeper *et al.* (2010). The endogenous variables in our model will adjust to make sure that those policy rules imply that the increase in government spending is paid for.

government might carry out without knowledge of the entire model as long as precise estimates of the original steady state values are available.

5 Learning about Fiscal Policy

The agents in our model act as Bayesian econometricians. They observe all relevant economic outcomes and use those observations to estimate the coefficients of the policy rules (7)-(11). Firms and households know all other aspects of the model.

We first describe how agents update their estimates of fiscal policy coefficients, then go on to derive the beliefs about the equilibrium dynamics induced by those estimates and finally derive expressions for the equilibrium dynamics in our model.

All private agents share the same beliefs and carry out inference by using the Kalman filter⁸, which means that they recursively apply Bayes' law. If we denote by Ω_t the vector of coefficients of all fiscal policy rules (which is exactly what the agents have to estimate) and by τ_t the vector of fiscal policy instruments at time t (i.e., the left-hand side variables of equations (7)-(11)), then the observation equation for the state space system used by the Kalman filter is given by:

$$\tau_t = X_{t-1}\Omega_t + \eta_t \quad (14)$$

where η_t collects the iid disturbances in the fiscal policy rules. X_{t-1} collects the right-hand side variables in the fiscal policy rules. In a previous section we have laid out how policy actually changes. Now we have to specify the perceived law of motion for Ω_t - how do firms and households in the economy think policy rule coefficients change over time? While we move away from the assumption of full-information rational expectations, the agents' views on policy changes are very close to the actual law of motion of the policy rule coefficients (i.e. the actual policy change we consider). In particular, our agents know at what time the policy rule coefficients change - they just do not know *what coefficients change* and the *magnitude of the change*. To be clear, agents also update their beliefs about fiscal policy in the periods in which the policy does not change. The following law of motion for the coefficients encodes these assumptions:

$$\Omega_t = \Omega_{t-1} + \mathbf{1}_t\nu_t \quad (15)$$

⁸For a comparison of learning when using the Kalman filter versus learning when using the common recursive least squares approach, see Sargent and Williams (2005).

$\mathbf{1}_t$ is an indicator function that equals 1 in the period in which fiscal policy changes⁹ and ν_t is a Gaussian vector with mean 0 for each element. This law of motion is inspired by the literature on time-varying coefficient models in empirical macroeconomics (such as Cogley and Sargent (2005) or Primiceri (2005))¹⁰. The perceived law of motion for the coefficients makes agents realize that fiscal policy changes infrequently. A similar modeling device has been introduced in time-varying parameter VAR models by Koop *et al.* (2009), who replace $\mathbf{1}_t$ with a random variable that can take on only the values 0 or 1. In the literature on learning in macroeconomic models, Marcet and Nicolini (2003) propose a learning mechanism in a similar spirit: Agents place greater weight on recent data if they suspect that there has been a structural change (i.e., whenever the estimated coefficients fit the data poorly). Introducing $\mathbf{1}_t$ into the agents' learning algorithm helps us to match the pattern of uncertainty displayed in figure 1.

If we were to set the variance of ν_t to a conformable matrix of zeros, then the private agents in our model would believe that fiscal policy rule coefficients do not change and they would estimate unknown constant coefficients. A non-zero covariance matrix for ν_t implies the belief that fiscal policy rule coefficients change when the actual policy change happens. This begs the question of how we calibrate the covariance matrix for ν_t , Σ_ν . We set this matrix to a scaling factor s times a diagonal matrix with the i th element on the diagonal being equal to the square of the i th element of Ω_0 . Ω_0 is the initial estimate of the policy rule coefficients, which we set to the true pre-policy-change values. This assumption makes any calibration for s easily interpretable - if $s = 1$, then a 1-standard-deviation shock can double the parameter, for example. We choose different values for s that endow the agents with different views on how likely or unlikely the actual policy change is - we calibrate s such that the policy changes we consider in our subsequent simulations represent either a 1, 2, or 3-standard-deviation shock according to Σ_ν . In order to be able to use the Kalman filter for the agents' inference problem, we have to assume that agents know the variance of the shocks in the policy rules.

Next, we move on to describe how the private agents in our model view the world - what is their perceived law of motion?

Given beliefs for Ω_t , agents in our model will adhere to the anticipated utility theory of decision-making (Kreps (1998)): they will act as if Ω_t is going to be fixed at the cur-

⁹We thus implicitly assume that the government can credibly announce that there is a change in fiscal policy, but it cannot credibly communicate in what way fiscal policy changes.

¹⁰An assumption of this kind (with $\mathbf{1}_t = 1\forall t$) has been applied in the learning literature by Sargent *et al.* (2006), for example.

rently estimated level forever more ¹¹. This is a common assumption in the literature on learning, see for example Milani (2007) or Sargent *et al.* (2006). Cogley *et al.* (2007) show that in a model of monetary policy the differences between anticipated-utility decision making and fully Bayesian learning are not large. They succinctly summarize the relationship between uncertainty and anticipated-utility decision making: "Although an anticipated-utility decision maker learns and takes account of model uncertainty, he does not design his decisions intentionally to refine future estimates."

A change in beliefs about fiscal policy will also induce a change in the beliefs about the steady state of the economy (see the description of the perceived steady state in the Appendix for details). If we denote the vector of all variables (plus a constant intercept) in the model economy by \mathbb{Y}_t , then we can stack the log-linearized equilibrium conditions (approximated around the perceived steady state) and the estimated fiscal policy rules to get the log-linearized perceived law of motion in the economy¹²:

$$A(\Omega_{t-1})\mathbb{Y}_t = B(\Omega_{t-1})E_t^*\mathbb{Y}_{t+1} + C(\Omega_{t-1})\mathbb{Y}_{t-1} + D\varepsilon_t^* \quad (16)$$

The asterisked expectations operator denotes expectations conditional on private sector beliefs about the economy. The asterisked vector of shocks ε_t^* includes the perceived fiscal policy shocks as well as the technology shock that agents can observe perfectly. This expectational difference equation can be solved using standard algorithms to yield the perceived law of motion for the economy at time t :

$$\mathbb{Y}_t = S(\Omega_{t-1})\mathbb{Y}_{t-1} + G(\Omega_{t-1})\varepsilon_t^* \quad (17)$$

$S(\Omega_{t-1})$ solves the following matrix quadratic equation¹³:

$$S(\Omega_{t-1}) = (A(\Omega_{t-1}) - B(\Omega_{t-1})S(\Omega_{t-1}))^{-1}C(\Omega_{t-1}) \quad (18)$$

and $G(\Omega_{t-1})$ is given by

$$G(\Omega_{t-1}) = (A(\Omega_{t-1}) - B(\Omega_{t-1})S(\Omega_{t-1}))^{-1}D \quad (19)$$

The beliefs in those equations are dated $t-1$ because of our timing assumption: Agents enter the current period (and make decisions in that period) with beliefs updated at

¹¹We use the posterior mean produced by the Kalman filter as a point estimate that the agents in the model condition on when forming expectations.

¹²This derivation follows Cogley *et al.* (2011). We also borrow their use of a projection facility: If no stable perceived law of motion exists, agents use the previous period's estimates.

¹³The perceived law of motion can be derived by assuming a VAR perceived law of motion of order 1 and then using the method of undetermined coefficients.

the end of the previous period. This makes the solution method recursive, otherwise we would have to jointly solve for outcomes and beliefs in every period.

Having described how agents update their estimates and their views on the dynamics of the variables in the model, we are now in a position to derive the equilibrium dynamics - the actual law of motion of the economy. This actual law of motion can be derived as follows: we modify $C(\Omega_{t-1})$ to now include the *true* policy coefficients. We call this matrix $C^{true}(\Omega_{t-1})$. Then the actual law of motion solves:

$$A(\Omega_{t-1})\mathbb{Y}_t = B(\Omega_{t-1})E_t^*\mathbb{Y}_{t+1} + C^{true}(\Omega_{t-1})\mathbb{Y}_{t-1} + D\varepsilon_t \quad (20)$$

where we now use the actual shock vector ε_t . Using the perceived law of motion to solve out for the expectations gives

$$\mathbb{Y}_t = H(\Omega_{t-1})\mathbb{Y}_{t-1} + G(\Omega_{t-1})\varepsilon_t \quad (21)$$

As can be seen from this derivation, actual economic outcomes will depend on both perceived and actual policy rule coefficients. H is given by:

$$H(\Omega_{t-1}) = S(\Omega_{t-1}) + (A(\Omega_{t-1}) - B(\Omega_{t-1})S(\Omega_{t-1}))^{-1}(C^{true}(\Omega_{t-1}) - C(\Omega_{t-1})) \quad (22)$$

We calibrate the initial covariance matrix of the estimators so that the initial standard deviation for each parameter is equal to 10 percent of its original value (which is also the true pre-policy-change value). We want agents to be reasonably confident about the pre-policy-change fiscal policy rules (so that before the policy change our agents behave very similarly to agents who know the fiscal policy rules perfectly). Since the policy change in our simulations only happens in period 10 and the agents update their estimates as well as the associated covariance matrix in the first 9 periods of the simulations, the exact calibration of the initial covariance matrix is not critical.

6 Results

6.1 A Roadmap

We will first present results for the full-information rational expectations benchmark¹⁴. We will then show how learning affects equilibrium outcomes by first discussing results in our benchmark specification, in which agents think that the true policy change is a 2-standard-deviation shock. We then go on to show how our different beliefs about the possible size of the policy change affect outcomes. After that we ask if learning would have any effects if there were no actual policy change.

Next, we ask how different information structures affect our results: Does it matter if agents know that only one specific coefficient changes or if agents think that other variables could affect fiscal policy?

We also assess the robustness of our result with respect to the specification of preferences: As we will see below, the behavior of labor supply seems to play an important role in the dynamics of our model. We thus check to see if our results hold under two preference specifications that imply very different behavior of labor supply: the preferences of King *et al.* (1988) and Greenwood *et al.* (1988), respectively. Finally, we show that our findings are robust to the choice of policy instrument that is changed: We consider a decrease in the intercept of the policy rule for the capital tax rate.

6.2 Rational Expectations

Figure 2 plots the median of the logarithm of the outcomes for our experiment under full-information rational expectations¹⁵. We see that there are very persistent effects on output, but ultimately output returns to a level very close to the initial steady state. The steady state of other variables is very much affected by the policy change though: Debt and the capital tax rate are permanently higher, leading to a permanently lower capital stock. The long-run level of the labor tax, on the other hand, remains basically unchanged, stemming from the parameter values of the policy rule for that instrument. Consumption shows very persistent effects and converges toward a lower

¹⁴Full-information rational expectations might be a misnomer since the agents in this economy do not anticipate the policy change - a common assumption when analyzing structural change in rational expectations models. When the change in fiscal policy happens, the agents are fully aware of the new policy, though.

¹⁵Mean outcomes are very similar.

steady state. Households raise their labor supply to partially offset the drop in capital. Overall, the effects of the policy change are a short-term small increase in output relative to a scenario in which the policy rule does not change (shown in figure 13 in the Appendix), coming at the cost of changes in the long-run behavior of the economy. As mentioned above, we will later check how robust our outcomes are to different preference specifications that lead to different behavior of the labor supply.

6.3 Benchmark Results

Now we turn to the economy under learning. First, we ask to what extent outcomes are different under learning relative to rational expectations when agents' beliefs about time variation are calibrated in such a way that the actual policy represents a 2-standard-deviation shock under the beliefs of the agents in the economy. Figure 3 shows a summary of the outcomes in that environment. The bottom panel shows the distribution of point estimates (median as well as 5th and 95th percentile bands) across simulations for the parameters in the government spending policy rule¹⁶. Agents quickly pick up on the change in G_c . Before the policy change, the uncertainty surrounding policy rule parameters is very small. There is a substantial increase in that uncertainty, as measured by the difference of the percentile bands, as policy changes. The uncertainty decreases again after the policy change for G_c . These patterns are consistent with the uncertainty index constructed by Baker *et al.* (2012)¹⁷. The uncertainty surrounding the response coefficients grows over time, but is very small in magnitude. There is also a slight bias in the estimation of these coefficients, but by inspecting the y-axis of these graphs one can see that the bias is small, too¹⁸. Thus, agents in this setup learn fast and the largest uncertainty in quantitative terms (that around G_c) disappears reasonably quickly. Does learning have any effect on outcomes then?

The top panel shows how average outcomes change relative to full-information rational expectations¹⁹: We plot the cumulated difference between median outcomes under

¹⁶Agents estimate the coefficients in all policy rules, but since the policy change occurs in the government spending policy rule we focus on those parameters.

¹⁷If we were to set $\mathbf{1}_t = 1 \forall t$ we would not get this strong reduction in uncertainty.

¹⁸The uncertainty in these response coefficients does not make a substantial difference for our results. This will become clear in the robustness check below in which agents only have to estimate G_c . The qualitative results in this case are the same as in our benchmark case.

¹⁹Note that the results under learning up to any period t are the same under our assumption of a permanent change in fiscal policy as they would be under the assumption of a temporary change that ends in period $t + 1$. This is not true under full-information rational expectations.

learning and under rational expectations relative to the original steady state. We thus plot

$$Diff_j^W = \sum_{t=1}^j \frac{(W_t^{learning} - W_t^{RE})}{W} \quad (23)$$

where W_t is the median of the variable of interest in *levels*, W is the associated original steady state, and the superscripts denote outcomes under learning and rational expectations²⁰. We see that before the negative technology shock and the associated policy change the cumulative differences are basically zero - there is no difference in average outcomes between learning and the full-information case. After the technology shock and the fiscal policy change in period 10 differences emerge - for a while consumption is higher under learning and hours worked lower. In those periods the agents in the learning model are actually better off on average. After a few periods the cumulative difference in consumption decreases again and ultimately becomes negative. The cumulative difference for GDP stays negative throughout. These effects are quantitatively significant: 40 periods (10 years) after the policy change the cumulative loss in GDP is 2 percent of the original steady state. The cumulative difference in the capital stock is persistently negative, which explains the differences in GDP given that the cumulative difference in hours is small. When it comes to fiscal policy instruments, we see that the cumulative difference in capital tax rates is basically zero, but that there are huge differences when it comes to debt. To summarize, not taking into account learning can have sizable effects on average outcomes in the economy. This is only one side of the coin though - the middle panel of figure 3 shows the standard deviation of (the log of) each variable relative to the volatility across the simulations under rational expectations. Consumption is substantially more volatile under learning at the time of the policy change (a 20 percent increase). Volatility also increases for GDP (around 2 percent) and other variables. These increases in volatility are smaller than those for GDP, but they are still significant. The changes in standard deviations are short-lived though, which is consistent with our observations that the estimated coefficients converge quickly. Why then are average outcomes affected so much? The sudden large fall in average investment under learning has very persistent effects via the capital stock. Thus, even though agents pick up quickly on changes, the short period of 'confusion' has persistent effects. This in turn stems from the underestimation of the persistence of the increase in government spending by agents - it takes them a few periods to fully grasp that the increase in government spending comes from an increase in G_c rather than a sequence of large shocks. The belief that part of the changes in government

²⁰In this calculation the outcomes under rational expectations and learning are calculated using the same shock sequences.

spending are temporary leads agents to believe that permanent increases in debt and capital taxes are not as large as they actually are, which substantially affects their investment decisions. Further evidence for this can be gathered by looking at figure 12. The figure plots the actual median path of the capital rate in levels under learning (this path is very similar under learning and rational expectations), the steady state capital tax rate associated with the original policy, the steady state capital tax rate associated with the new policy rule and the median perceived steady state across simulations. As the policy change happens, the rational expectations agents immediately realize that the new steady state of capital taxes is the green line, whereas agents under learning think the steady state is given by the perceived steady state. Thus, relative to steady state rational expectations agents find it more profitable to invest even at the time of the crisis because they know that the capital tax will be higher on average than the learning agents think. In more technical terms, the log-linearized equilibrium conditions we use will give investment as a negative function of (among other things) $\log(\tau_t^K) - \log(\tau^K)$, which will be larger in absolute value for the rational expectations agents because they know that the steady state is larger. This is only a partial explanation because the coefficients multiplying the log difference term are also a function of the (perceived or actual) steady state. Nonetheless, the dynamics of the perceived steady state of capital taxes seem to be one factor contributing to the difference in investment. This also sheds light on an interesting feature of our model: The agents are very much certain about the coefficients of the capital tax policy rule (they estimate them, but the associated estimates do not move significantly), but they are still very uncertain about the steady state value of that policy instrument. This is due to their uncertainty about the steady state of debt and GDP owing to the uncertainty surrounding government spending. GDP and debt enter the right-hand side of the capital tax policy rule and thus influence the steady state of the capital tax rate. In at least one direction we are underestimating the average effects of learning: If the policy were autocorrelated, it would take the agents longer to figure out that a change in G_c drives the policy change, rather than a sequence of shocks.

6.4 The Effect of Agents' Beliefs

Next we ask to what extent outcomes under learning would be different if agents either think that the same policy change is more likely than before (it represents a 1-standard-deviation shock) or less likely (it represents a 3-standard-deviation shock). The shape of the plotted objects remains the same as before. However, the magnitudes do change

substantially and there is a clear pattern: The less likely agents find a large change in policy, the bigger the differences in average outcomes between learning and rational expectations - it takes agents longer to learn. This longer transition has the effect of substantially decreasing volatility. Thus it is not clear if a policymaker contemplating a policy change would want agents to be uncertain about policy and consider large changes, or if that policymaker would want agents to believe that there will be only small policy changes. Ultimately this will depend on the preferences and the decision horizon of the policymaker.

6.5 Learning When There is no Policy Change

An important question is what drives the differences between learning and rational expectations: Is it the change in policy or would learning also lead to different outcomes when there is no policy change? The pre-policy-change part of the results above strongly indicates that if agents did not contemplate a policy change (i.e., $\mathbf{1}_t = 0 \forall t$), then there would be no noticeable difference between learning and rational expectations. But what would happen if the agents did contemplate a policy change just as above, but there was none? Figure 6 tackles that question. Comparing this figure with figure 3, we see that the mere suspicion of a policy change on the part of the agents already leads to substantial increases in volatility (which are smaller than in the case with changes to fiscal policy, though), but average effects are substantially smaller.

6.6 Information Structure

Does it matter whether agents know exactly what parameter in the fiscal policy rule changes or what variables enter into the fiscal policy rules? We turn to these questions next. Both of these experiments use the benchmark calibration for the agents' beliefs. First, we endow agents with the knowledge that only G_c changes. The results of this exercise are given in figure 7. In this case volatilities are dampened relative to our benchmark case depicted in figure 3, but average outcomes behave very similarly.

Next we ask what would happen if the agents thought that another variable (in our case consumption) would enter the right-hand side of the policy rule for government spending. We initialize the beliefs about the coefficient on consumption at zero. Figure 8 shows the relevant outcomes. The parameter estimates for the other coefficients are very similar to our benchmark case (the estimate for the coefficient on consumption

stays centered on zero throughout). Average outcomes and volatilities are very similar to the benchmark case as well - it seems that agents entertaining more general models (within certain bounds) does not substantially change our conclusions.

6.7 Preferences

Do our results hold when agents have different preferences? To address this issue with a particular focus on the behavior of labor supply, we redo our benchmark analysis for two classes of preferences that imply very different wealth effects on labor supply: the preferences of Greenwood *et al.* (1988) and those of King *et al.* (1988). The equations for both cases are laid out in the Appendix. Figures 9 and 10 show the results for these two cases. While the dynamics differ from our benchmark case for both preferences, the big picture remains the same: We see substantial differences in average outcomes and increases in volatility relative to rational expectations.

6.8 Capital Tax Change

After a negative shock hits the economy, government spending is not the only instrument the fiscal sector can change to boost the economy. In figure 11 we study a capital tax decrease equivalent to 1 percent of GDP. This is calculated along the lines of Leeper *et al.* (2010) and our own calculations for the government spending case, so that the decrease of total capital tax revenues approximately equals one percent of overall pre-policy-change steady state GDP. Qualitatively the results are the same as under the scenario of an increase of government spending. Cumulated GDP is lower by about 5 percent after the end of our simulation horizon while cumulated debt is around 15 percent higher in the case of learning compared to the rational expectations outcome. Investment and therefore also capital are decreasing constantly throughout. Volatility increases are quite small for all variables.

7 Conclusion

Our experiments point to the conclusion that we should be cautious when evaluating fiscal policy proposals solely on the basis of a full-information analysis. We have

endowed agents with substantial knowledge of the structure of the economy and the timing of the policy change, thus focusing the uncertainty agents face on a very specific aspect - the post-policy-change values of the policy rule coefficients. Yet we still find meaningful differences between a rational expectations model and our learning model. The views that agents hold about the magnitude of possible policy changes has a significant impact on outcomes, pointing toward a possible role for communicating policy changes. However, a policymaker would have to be sure of the effects of their communication on the public's views to avoid undesired outcomes - if that communication only increases the probability that private agents assign to large policy changes then communication would lead to substantially more volatility after the policy change.

References

- Baker, S. R., Bloom, N., and Davis, S. J. (2012). Has economic policy uncertainty hampered the recovery? Technical report.
- Bernanke, B., Gertler, M., and Gilchrist, S. (1999). *The financial accelerator in a quantitative business cycle framework*, volume 1, chapter 21, pages 1341–1393. Handbook of Macroeconomics.
- Born, B. and Pfeifer, J. (2013). Policy risk and the business cycle. CESifo Working Paper Series 4336, CESifo Group Munich.
- Caprioli, F. (2010). Optimal fiscal policy when agents are learning. Working paper.
- Cogan, J. F., Cwik, T., Taylor, J. B., and Wieland, V. (2010). New Keynesian versus old Keynesian government spending multipliers. *Journal of Economic Dynamics and Control*, **34**, 281–295.
- Cogley, T. and Sargent, T. J. (2005). Drift and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics*, **8**(2), 262–302.
- Cogley, T., Colacito, R., and Sargent, T. J. (2007). Benefits from U.S. monetary policy experimentation in the days of Samuleson and Solow and Lucas. *Journal of Money, Credit & Banking*, **39**, 67–99.
- Cogley, T., Matthes, C., and Sbordone, A. M. (2011). Optimal disinflation under learning. Technical report.
- Davig, T., Leeper, E. M., and Walker, T. B. (2010). Unfunded liabilities and uncertain fiscal financing. *Journal of Monetary Economics*, **57**(5), 600–619.
- Eusepi, S. and Preston, B. (2011). Learning the fiscal theory of the price level: Some consequences of debt-management policy. *Journal of the Japanese and International Economies*, forthcoming.
- Fernandez-Villaverde, J., Guerron-Quintana, P., Kuester, K., and Rubio-Ramirez, J. (2011). Fiscal volatility shocks and economic activity. PIER Working Paper Archive 11-022, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania.
- Giannitsarou, C. (2006). Supply-side reforms and learning dynamics. *Journal of Monetary Economics*, **53**(2), 291–309.
- Greenwood, J., Hercowitz, Z., and Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. *American Economic Review*, **78**(3), 402–17.

- Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, **99**(4), 1097–1118.
- Karantounias, A. G. (2013). Managing pessimistic expectations and fiscal policy. *Theoretical Economics*, **8**, 193–231.
- King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles : I. the basic neoclassical model. *Journal of Monetary Economics*, **21**(2-3), 195–232.
- Kliem, M. and Kriwoluzky, A. (2013). Toward a Taylor rule for fiscal policy. Working paper, Deutsche Bundesbank and Rheinische Friedrich-Wilhelms Universität Bonn.
- Koop, G., Leon-Gonzalez, R., and Strachan, R. (2009). On the evolution of monetary policy. *Journal of Economic Dynamics and Control*, **33**, 997–1017.
- Kreps, D. (1998). *Anticipated Utility and Dynamic Choice*, pages 242–274. Frontiers of Research in Economic Theory. Cambridge University Press.
- Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, **50**(6), 1345–70.
- Leeper, E. M., Plante, M., and Traum, N. (2010). Dynamics of fiscal financing in the United States. *Journal of Econometrics*, **156**, 304–321.
- Marcet, A. and Nicolini, J.-P. (2003). Recurrent hyperinflations and learning. *American Economic Review*, **93**, 1476–1498.
- Milani, F. (2007). Expectations, learning and macroeconomic persistence. *Journal of Monetary Economics*, **54**(7), 2065–2082.
- Mitra, K., Evans, G. W., and Honkapohja, S. (2012). Fiscal policy and learning. Research Discussion Papers 5/2012, Bank of Finland.
- Primiceri, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, **72**(3), 821–852.
- Sargent, T., Williams, N., and Zha, T. (2006). Shocks and government beliefs: The rise and fall of American inflation. *American Economic Review*, **96**(4), 1193–1224.
- Sargent, T. J. and Williams, N. (2005). Impacts of priors on convergence and escapes from Nash inflation. *Review of Economic Dynamics*, **8**, 360–391.
- Sims, C. A. (2001). Solving linear rational expectations models. *Computational Economics*, **20**, 1–20.
- Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, **97**(3), 586–606.

Appendix

A First-Order Conditions

Households:

$$\begin{aligned}\frac{C_t^{-\sigma}}{1 + \tau_t^C} &= E_t \frac{\beta R_t C_{t+1}^{-\sigma}}{1 + \tau_{t+1}^C} \\ L_t^{1+\phi}(1 + \tau_t^C) &= C_t^{-\sigma}(1 - \tau_t^L)(1 - \alpha)Y_t \\ 1 &= \beta E_t \frac{C_{t+1}^{-\sigma}(1 + \tau_{t+1}^C)}{C_t^{-\sigma}(1 + \tau_t^C)} \left((1 - \tau_{t+1}^K) \frac{\alpha Y_{t+1}}{K_t} + (1 - \delta) \right)\end{aligned}$$

Firms:

$$\begin{aligned}W_t &= \frac{(1 - \alpha)Y_t}{L_t} \\ R_t^K &= \frac{\alpha Y_t}{K_{t-1}}\end{aligned}$$

B Log-Linearized Model

Households:

$$\begin{aligned}(1 + \phi)\log(L_t) + \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) &= Const^L + \log(Y_t) - \left(\frac{\tau_c^L}{1 - \tau_c^L} \right) \log(\tau_t^L) - \sigma \log(C_t) \\ \log(C_t) &= Const^C - \frac{1}{\sigma} \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) + \frac{1}{\sigma} \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_{t+1}^C) + \log(C_{t+1}) - \frac{1}{\sigma} \log(R_t) \\ \log(K_t) &= Const^{LoM} + (1 - \delta)\log(K_{t-1}) + \delta \log(I_t) \\ \log(Y_t) &= Const^Y + \log(A_t) + \alpha \log(K_{t-1}) + (1 - \alpha)\log(L_t) \\ \sigma E_t \log(C_{t+1}) &= Const^K + \sigma \log(C_t) - \frac{\tau_c^C}{(1 + \tau_c^C)} E_t \log(\tau_{t+1}^C) + \frac{\tau_c^C}{(1 + \tau_c^C)} \log(\tau_t^C) \\ &+ \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(Y_{t+1}) - \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(K_t) - \beta \tau_c^K \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(\tau_{t+1}^K)\end{aligned}$$

Firms:

$$\begin{aligned}\log(Y_t) &= \text{Const}^{Agg} + \frac{C_{ss}}{Y_{ss}} \log(C_t) + \frac{I_{ss}}{Y_{ss}} \log(I_t) + \frac{G_{ss}}{Y_{ss}} \log(G_t) \\ \log(A_t) &= \text{Const}^A + \rho_a \log(A_{t-1}) + \epsilon_{tA}\end{aligned}$$

Policy Rules:

$$\begin{aligned}\log(B_t) &+ \tau_c^K \alpha \frac{Y_{ss}}{B_{ss}} (\log(\tau_t^K) + \log(Y_t)) + \tau_c^L (1 - \alpha) \frac{Y_{ss}}{B_{ss}} (\log(\tau_t^L) + \log(Y_t)) + \tau_c^C \frac{C_{ss}}{B_{ss}} (\log(\tau_t^C) + \log(C_t)) \\ &= \text{Const}^B + \frac{1}{\beta} \log(R_{t-1}) + \frac{1}{\beta} \log(B_{t-1}) + \frac{G_{ss}}{B_{ss}} \log(G_t) + \frac{Z_{ss}}{B_{ss}} \log(Z_t) \\ \log(G_t) &= G_c - \rho_{g,y} \log(Y_{t-1}) - \rho_{g,b} \log(B_{t-1}) + \epsilon_t^G \\ \log(Z_t) &= Z_c - \rho_{z,y} \log(Y_{t-1}) - \rho_{z,b} \log(B_{t-1}) + \epsilon_t^Z \\ \log(\tau_t^C) &= \tau_c^c + \epsilon_t^C \\ \log(\tau_t^L) &= \tau_c^l + \rho_{L,y} \log(Y_{t-1}) + \rho_{L,b} \log(B_{t-1}) + \epsilon_t^L \\ \log(\tau_t^K) &= \tau_c^k + \rho_{K,y} \log(Y_{t-1}) + \rho_{K,b} \log(B_{t-1}) + \epsilon_t^K\end{aligned}$$

with the constants given by:

Constant	Expression
G_c	$\log(G_c) + \rho_{g,y} \log(Y_{ss}) + \rho_{g,b} \log(B_{ss})$
Z_c	$\log(Z_c) + \rho_{z,y} \log(Y_{ss}) + \rho_{z,b} \log(B_{ss})$
τ_c^l	$\log(\tau_c^L) - \rho_{L,y} \log(Y_{ss}) - \rho_{L,b} \log(B_{ss})$
τ_c^k	$\log(\tau_c^K) - \rho_{K,y} \log(Y_{ss}) - \rho_{K,b} \log(B_{ss})$
τ_c^c	$\log(\tau_c^C)$
Const^B	$\log(B_{ss}) (1 - \frac{1}{\beta}) + \tau_c^K \alpha \frac{Y_{ss}}{B_{ss}} (\log(\tau_c^K) + \log(Y_{ss})) + \tau_c^L (1 - \alpha) \frac{Y_{ss}}{B_{ss}} (\log(Y_{ss}) + \log(\tau_c^L))$ $+ \tau_c^C \frac{C_{ss}}{B_{ss}} (\log(\tau_c^C) + \log(C_{ss})) - \frac{1}{\beta} \log(R_{ss}) - \frac{G_{ss}}{B_{ss}} \log(G_{ss}) - \frac{Z_{ss}}{B_{ss}} \log(Z_{ss})$
Const^{LoM}	$\delta (\log(K_{ss}) - \log(I_{ss}))$
Const^L	$(1 + \phi) \log(L_{ss}) + \frac{\tau_c^C}{1 + \tau_c^C} \log(\tau_c^C) - \log(Y_{ss}) + \frac{\tau_c^L}{1 + \tau_c^L} \log(\tau_c^L)$
Const^C	$\frac{1}{\sigma} \log(R_{ss})$
Const^Y	$\log(Y_{ss}) - \log(A_{ss}) - \alpha \log(K_{ss}) - (1 - \alpha) \log(L_{ss})$
Const^A	$\log(A_{ss})$
Const^{Agg}	$\log(Y_{ss}) - \frac{C_{ss}}{Y_{ss}} \log(C_{ss}) - \frac{G_{ss}}{Y_{ss}} \log(G_{ss}) - \frac{I_{ss}}{Y_{ss}} \log(I_{ss})$
Const^K	$-\beta (1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(Y_{ss}) + \beta (1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(K_{ss}) + \beta (\tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(\tau_c^K)$

C Parameters

Calibrated Parameters

Description	Parameter	Value
Impatience	β	0.99
Capital share	α	0.33
Depreciation rate	δ	0.025
CES utility consumption	σ	2
CES utility labor	ϕ	2
Coeff. on Y in gov. exp. rule	$\rho_{g,y}$	0.034
Coeff. on B in gov. exp. rule	$\rho_{g,b}$	0.23
Coeff. on Y in transfer rule	$\rho_{z,y}$	0.13
Coeff. on B in transfer rule	$\rho_{z,b}$	0.5
Coeff. on Y labor tax rule	$\rho_{L,y}$	0.36
Coeff. on B labor tax rule	$\rho_{L,b}$	0.049
Coeff. on Y capital tax rule	$\rho_{K,y}$	1.7
Coeff. on B capital tax rule	$\rho_{K,b}$	0.39
AR parameter technology	ρ_a	0.9
Std. deviation technology	σ_a	0.0062
Std. deviation gov. spending	σ_g	0.031
Std. deviation transfers	σ_z	0.034
Std. deviation cons.tax	σ_c	0.04
Std. deviation labor tax	σ_l	0.03
Std. deviation capital tax	σ_k	0.044

Table 1: Calibrated parameters of the model

Initial Steady State Values of the Actual Law of Motion

Description	Parameter	Value
Output	Y_{ss}	2.0601
Consumption	C_{ss}	1.5010
Cons. tax rate	τ_c^C	0.0287
Capital tax rate	τ_c^K	0.2452
Labor	L_{ss}	0.7847
Investment	I_{ss}	0.3655
Capital	K_{ss}	14.6195
Debt	B_{ss}	0.5623
Labor tax rate	τ_c^L	0.1886
Government spending	G_c	0.1936
Transfers	Z_c	0.2709
Technology	A_{ss}	1
Interest rate	R_{ss}	1.01

Table 2: Calibrated parameters of the model

Perceived Steady States

The perceived steady states in the updating algorithm of the agents are given by the following twelve equations:

$$\begin{aligned}
R &= \frac{1}{\beta} \\
\frac{\alpha Y}{K} &= \frac{\frac{1}{\beta} - (1 - \delta)}{1 - \tau^K} \\
L^{1+\phi}(1 + \tau^C) &= C^{-\sigma}(1 - \tau^L)(1 - \alpha)Y \\
Y &= AK^\alpha L^{1-\alpha} \\
Y &= C + I + G \\
I &= \delta K \\
B &= B \frac{1}{\beta} - \tau^K \alpha Y - \tau^L (1 - \alpha)Y - \tau^C C + G + Z \\
Const^G &= \log(G) + \rho_{g,y} \log(Y) + \rho_{g,b} \log(B) \\
Const^Z &= \log(Z) + \rho_{z,y} \log(Y) + \rho_{z,b} \log(B) \\
Const^{\tau^L} &= \log(\tau_c^L) - \rho_{L,y} \log(Y) - \rho_{L,b} \log(B) \\
Const^{\tau^K} &= \log(\tau_c^K) - \rho_{K,y} \log(Y) - \rho_{K,b} \log(B) \\
Const^{\tau^C} &= \log(\tau_c^C)
\end{aligned}$$

for the twelve variables: $Y, K, L, C, G, Z, \tau^L, \tau^K, \tau^C, B, I, R$, which are solved numerically.

D Figures

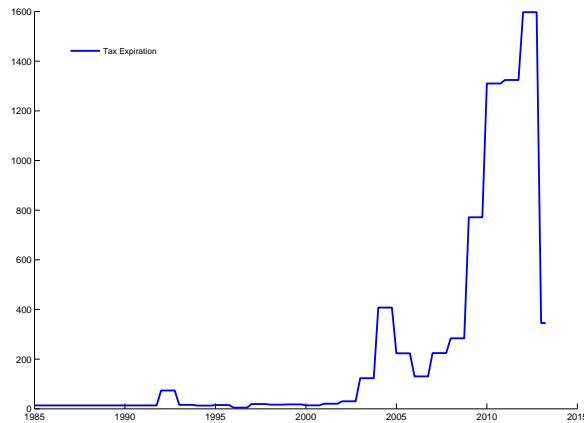


Figure 1: Fiscal uncertainty index by Baker *et al.* (2012)

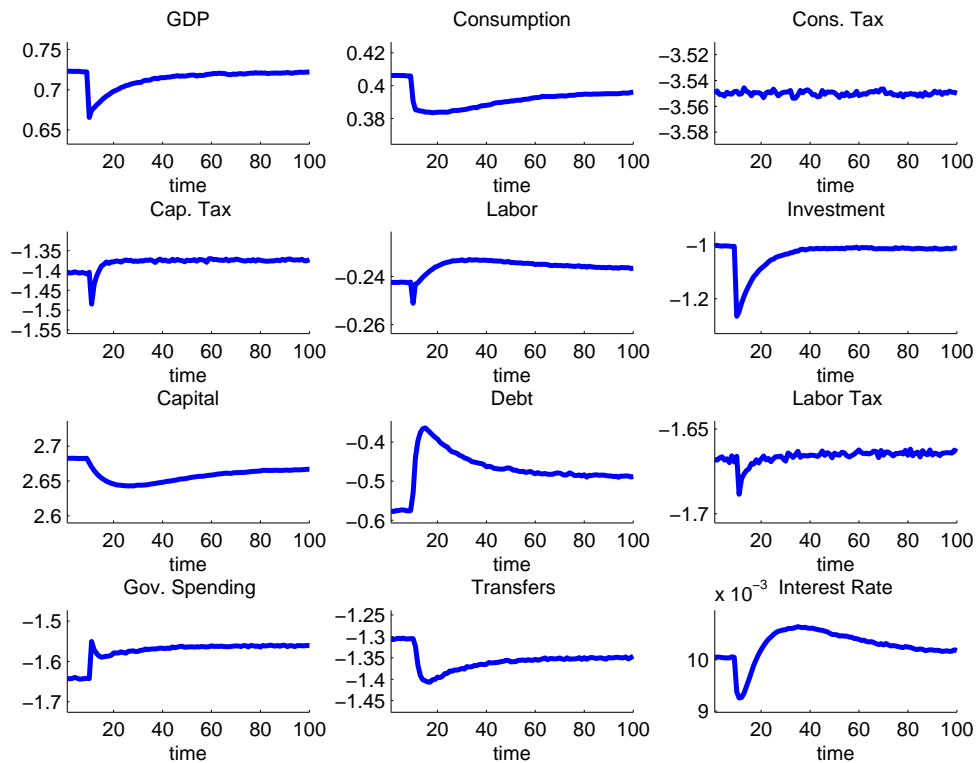


Figure 2: Log outcomes under rational expectations

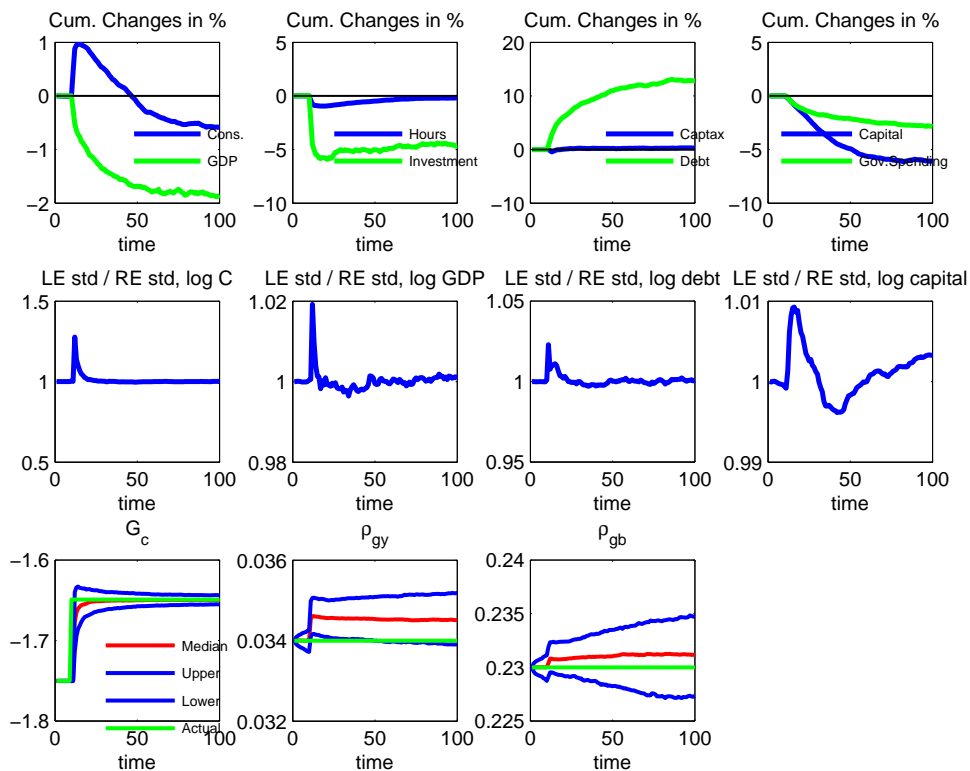


Figure 3: Summary of outcomes under learning

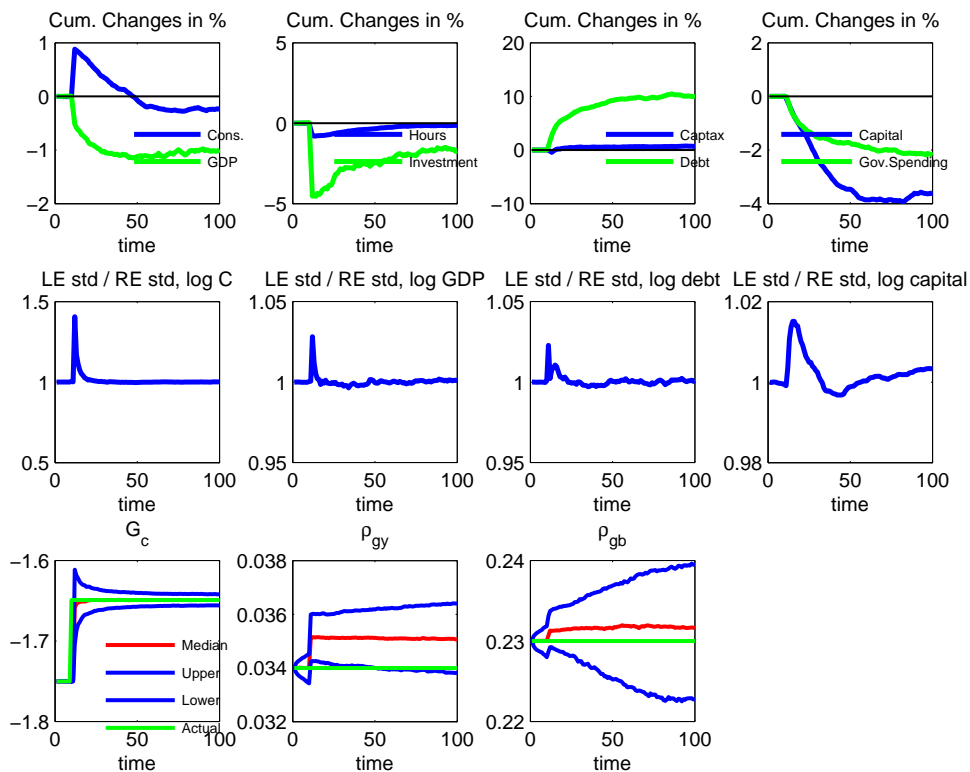


Figure 4: Summary of outcomes under learning, 1-standard-deviation case

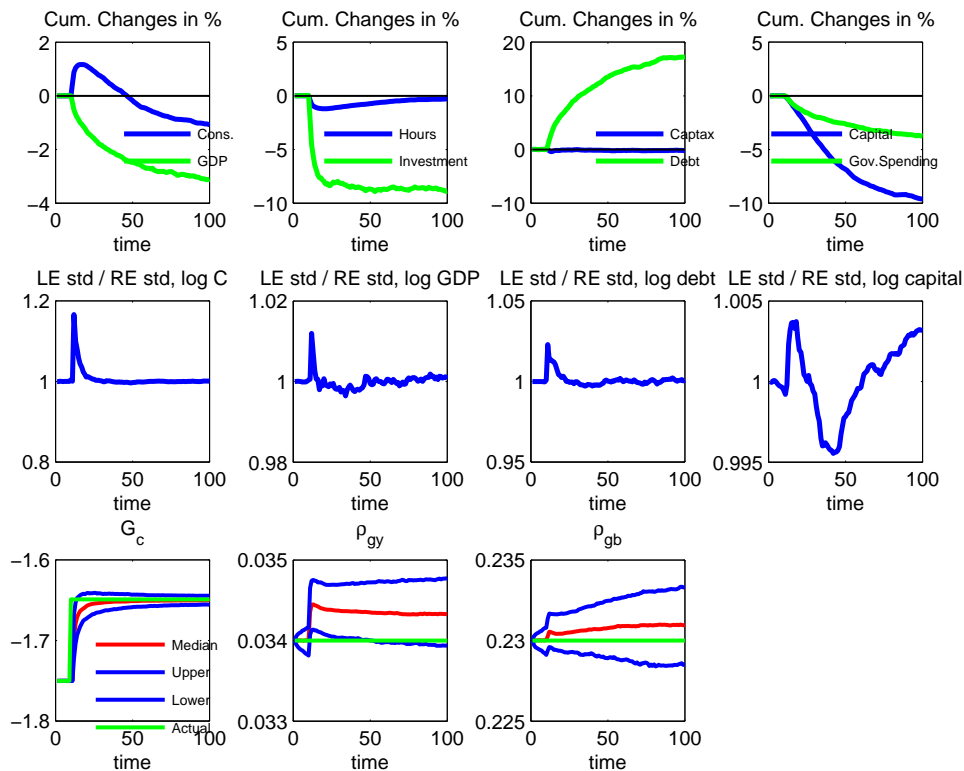


Figure 5: Summary of outcomes under learning, 3-standard-deviations case

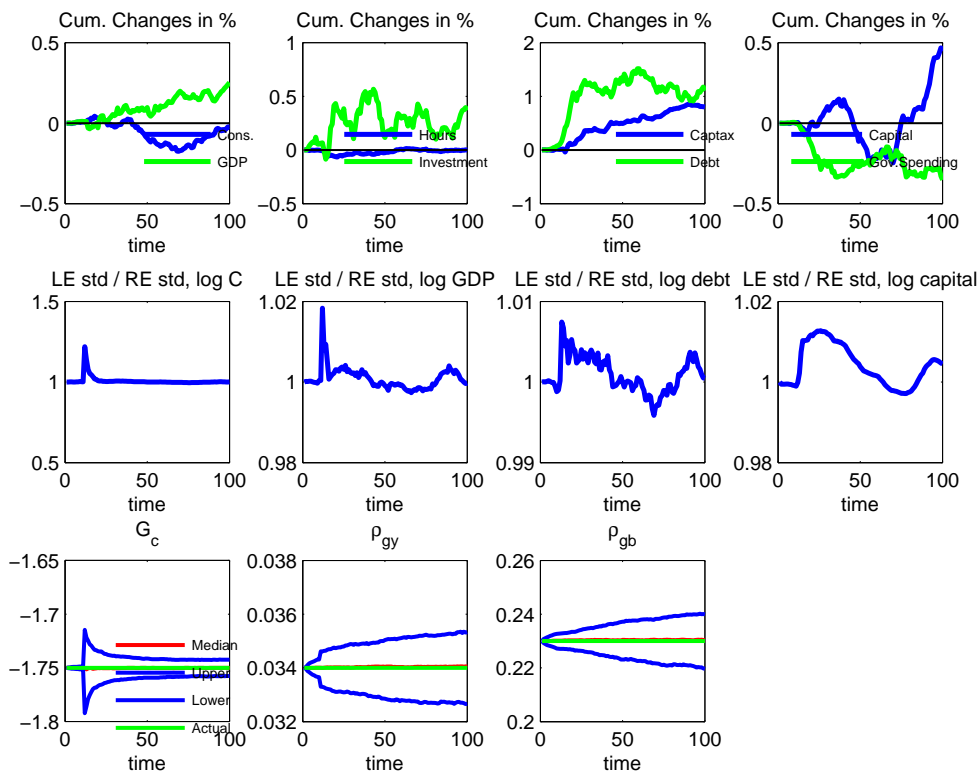


Figure 6: Summary of outcomes under learning when there is no fiscal policy change

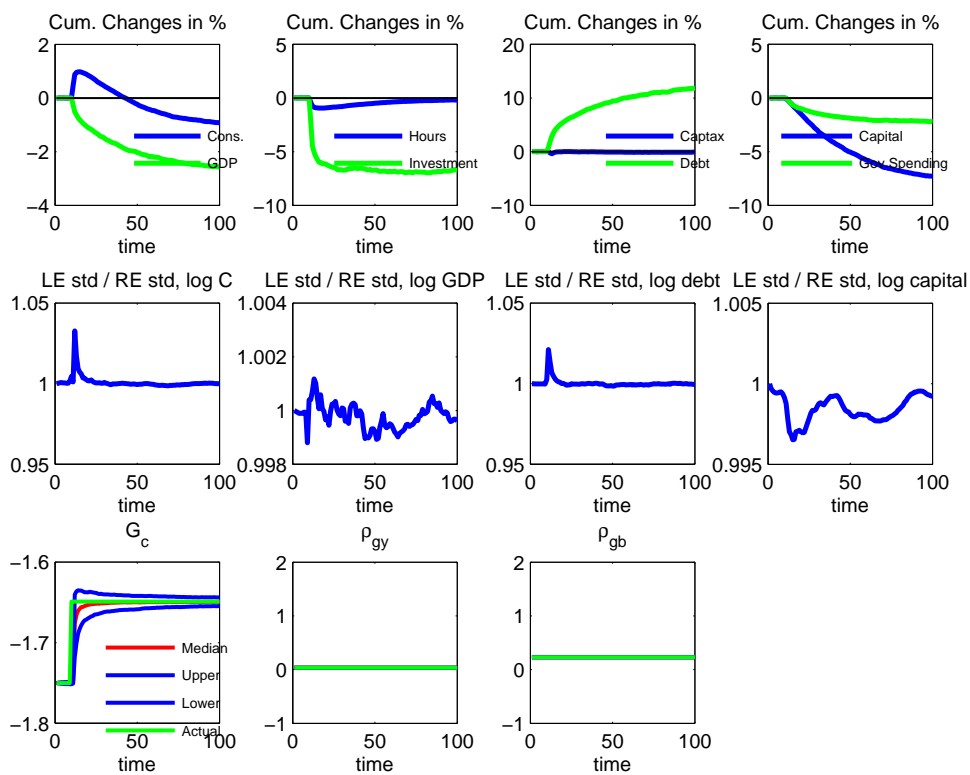


Figure 7: Summary of outcomes under learning when agents only need to learn about G_c

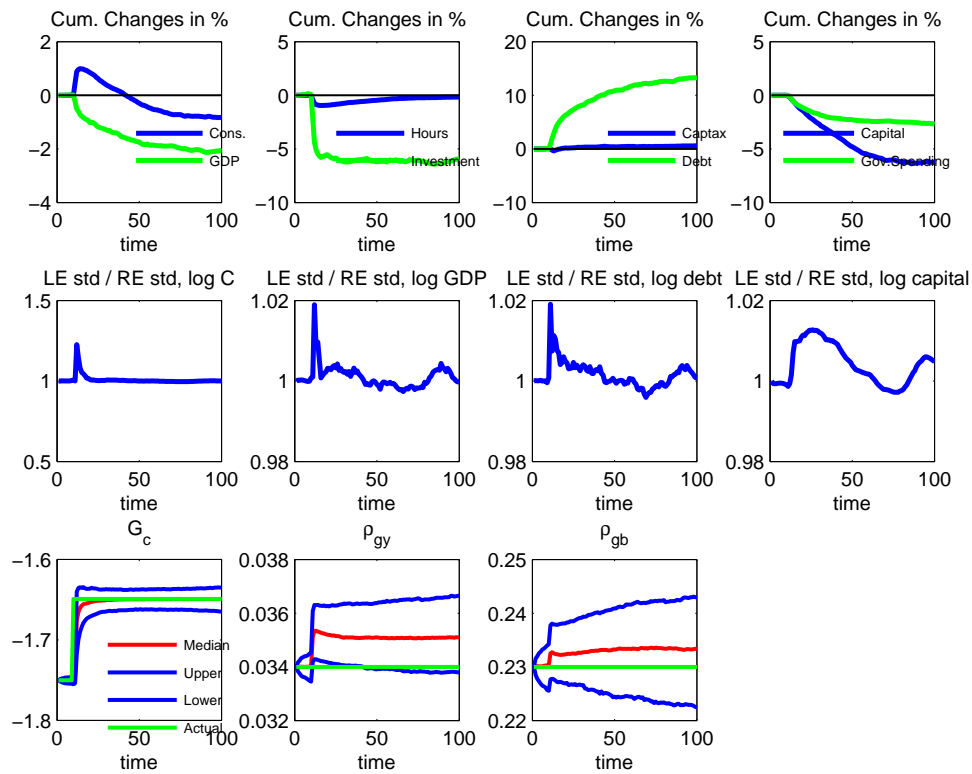


Figure 8: Summary of outcomes under learning when agents think that consumption enters the policy rule for government spending

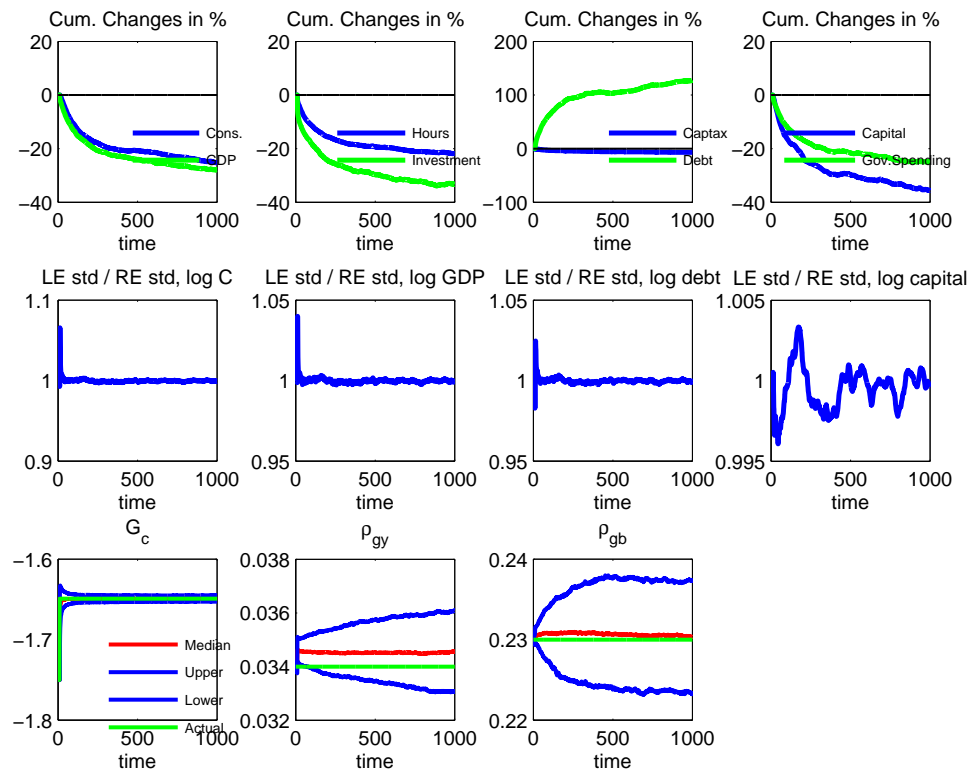


Figure 9: Summary of outcomes under learning when agents have GHH preferences

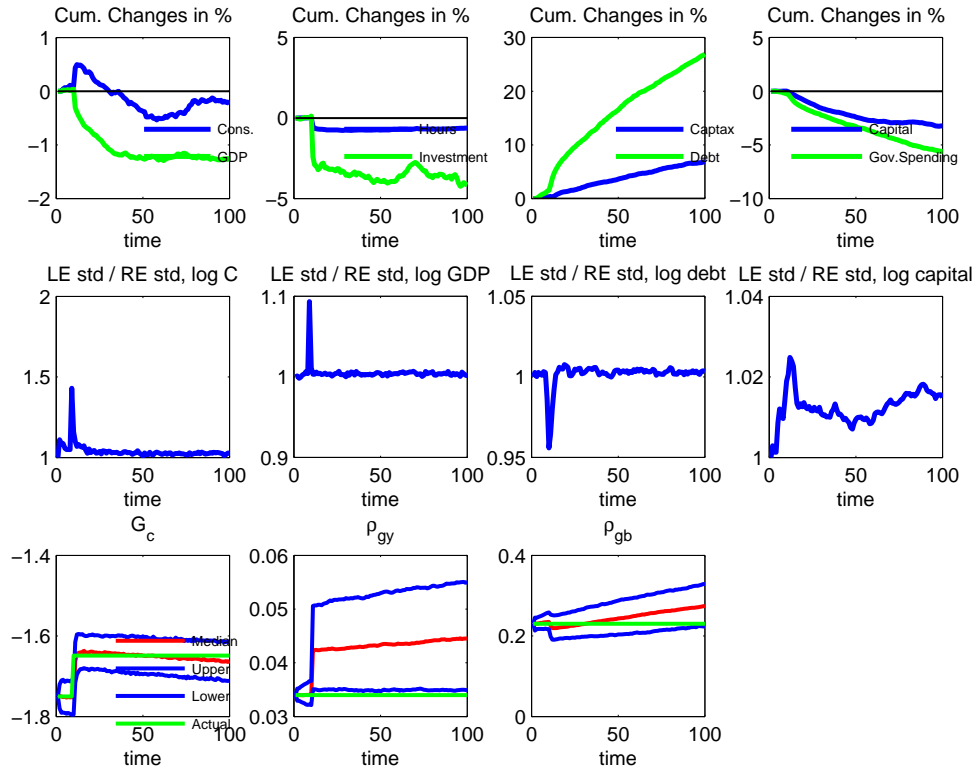


Figure 10: Summary of outcomes under learning when agents have KPR preferences

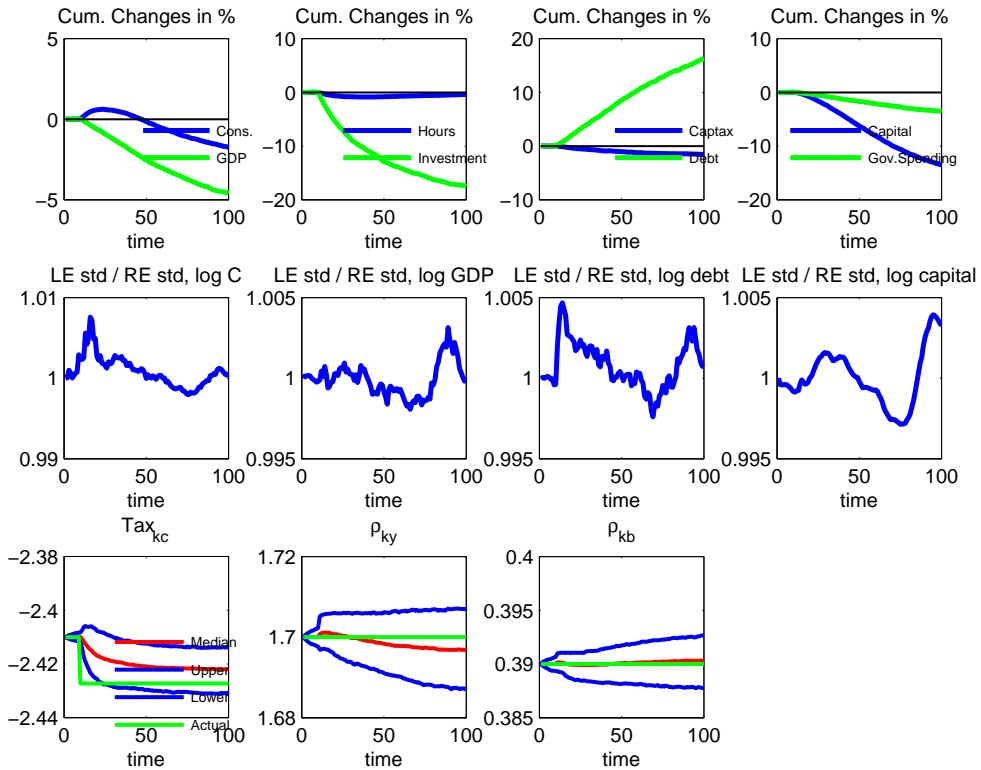


Figure 11: Summary of outcomes under learning when the capital tax policy rule changes

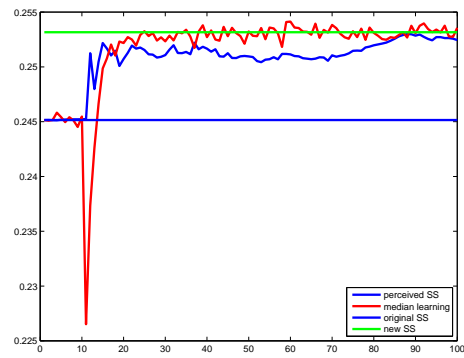


Figure 12: The capital tax rate

E Robustness Checks: Different Utility Function Specifications

A: First-order conditions of households: As robustness checks we consider the following utility function (compare Jaimovich and Rebelo (2009)):

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \frac{(C_t - \psi N_t^\theta X_t)^{1-\sigma} - 1}{1-\sigma} \quad (24)$$

with $X_t = C_t^\gamma X_{t-1}^{1-\gamma}$ which nests both the King *et al.* (1988) Preferences ($\gamma = 1$) and the Greenwood *et al.* (1988) preferences ($\gamma = 0$).

$$\begin{aligned} (C_t - \psi N_t^\theta X_t)^{-\sigma} + \mu_t \gamma C_t^{\gamma-1} X_t^{1-\gamma} &= \lambda_t (1 + \tau_t^c) \\ (C_t - \psi N_t^\theta X_t)^{-\sigma} \psi N_t^\theta + \mu_t &= \beta E_t \left[\mu_{t+1} (1 - \gamma) C_{t+1}^\gamma X_t^{-\gamma} \right] \\ (C_t - \psi N_t^\theta X_t)^{-\sigma} \psi \theta N_t^{\theta-1} X_t &= \lambda_t (1 - \tau_t^l) W_t \\ 1 &= \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \left((1 - \tau_{t+1}^K) R_{t+1}^K + (1 - \delta) \right) \end{aligned}$$

B: First-order conditions in the GHH case:

$$\begin{aligned} \frac{(C_t - \psi N_t^\theta)^{-\sigma}}{(1 + \tau_t^C)} + \beta E_t \frac{R_t (C_{t+1} - \psi N_{t+1}^\theta)^{-\sigma}}{(1 + \tau_{t+1}^C)} \\ \psi \theta N_t^\theta (1 + \tau_t^C) &= (1 - \tau_t^L) (1 - \alpha) Y_t \\ 1 &= \beta E_t \frac{(C_{t+1} - \psi N_{t+1}^\theta)^{-\sigma} (1 + \tau_t^C)}{(C_t - \psi N_t^\theta)^{-\sigma} (1 + \tau_{t+1}^C)} \left((1 - \tau_{t+1}^K) \alpha \frac{Y_{t+1}}{K_t} + (1 - \delta) \right) \end{aligned}$$

C: Log-linearized conditions in the GHH case

$$\begin{aligned}
& \theta \log(L_t) + \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) = \text{Const}^L + \log(Y_t) - \left(\frac{\tau_c^L}{1 - \tau_c^L} \right) \log(\tau_t^L) \\
& - \frac{\sigma C_{ss}}{C_{ss} - \psi N_{ss}^\theta} \log(C_t) + \frac{\sigma \psi \theta N_{ss}^\theta}{C_{ss} - \psi N_{ss}^\theta} \log(N_t) - \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) = \text{Const}^C + R_t - \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_{t+1}^C) \\
& - \frac{\sigma C_{ss}}{C_{ss} - \psi N_{ss}^\theta} \log(C_{t+1}) + \frac{\sigma \psi \theta N_{ss}^\theta}{C_{ss} - \psi N_{ss}^\theta} \log(N_{t+1}) \\
& \frac{\sigma C_{ss}}{C_{ss} - \psi N_{ss}^\theta} E_t \log(C_{t+1}) - \frac{\sigma \psi \theta N_{ss}^\theta}{C_{ss} - \psi N_{ss}^\theta} \log(N_{t+1}) - \frac{\tau_c^C}{(1 + \tau_c^C)} \log(\tau_t^C) - \frac{\sigma C_{ss}}{C_{ss} - \psi N_{ss}^\theta} \log(C_t) \\
& + \frac{\sigma \psi \theta N_{ss}^\theta}{C_{ss} - \psi N_{ss}^\theta} \log(N_t) + \frac{\tau_c^C}{(1 + \tau_c^C)} E_t \log(\tau_{t+1}^C) = \text{Const}^K \\
& + \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(Y_{t+1}) - \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(K_t) \\
& - \beta \tau_c^K \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(\tau_{t+1}^K)
\end{aligned}$$

D: First-order conditions in the KPR case:

$$\begin{aligned}
\frac{(C_t - \psi N_t^\theta C_t)^{-\sigma} (1 - \psi N_t^\theta)}{(1 + \tau_t^C)} &= \beta E_t \frac{R_t (C_{t+1} - \psi N_{t+1}^\theta C_{t+1})^{-\sigma} (1 - \psi N_{t+1}^\theta)}{(1 + \tau_{t+1}^c)} \\
\psi \theta C_t N_t^\theta (1 + \tau_t^C) &= (1 - \psi N_t^\theta) (1 - \tau_t^L) (1 - \alpha) Y_t \\
1 &= \beta E_t \frac{(C_{t+1} - \psi N_{t+1}^\theta C_{t+1})^{-\sigma} (1 - \psi N_{t+1}^\theta) (1 + \tau_t^C)}{(C_t - \psi N_t^\theta C_t)^{-\sigma} (1 - \psi N_t^\theta) (1 + \tau_{t+1}^C)} \\
&\quad \left((1 - \tau_{t+1}^K) \alpha \frac{Y_{t+1}}{K_t} + (1 - \delta) \right)
\end{aligned}$$

E: Log-linearized conditions in the KPR case

$$\begin{aligned}
\theta \log(N_t) &+ \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) + \log(C_t) = \text{Const}^L + \log(Y_t) - \left(\frac{\tau_c^L}{1 - \tau_c^L} \right) \log(\tau_t^L) \\
&- \frac{\psi \theta N_{ss}^\theta}{1 - \psi N_{ss}^\theta} \log(N_t) \\
-\sigma \log(C_t) &- \frac{(1 - \sigma) \psi \theta N_{ss}^\theta}{1 - \psi N_{ss}^\theta} \log(N_t) - \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_t^C) = \text{Const}^C + R_t - \left(\frac{\tau_c^C}{1 + \tau_c^C} \right) \log(\tau_{t+1}^C) \\
&- \sigma \log(C_{t+1}) - \frac{(1 - \sigma) \psi \theta N_{ss}^\theta}{1 - \psi N_{ss}^\theta} \log(N_{t+1}) \\
\sigma E_t \log(C_{t+1}) &+ \frac{(1 - \sigma) \psi \theta N_{ss}^\theta}{1 - \psi N_{ss}^\theta} \log(N_{t+1}) - \frac{\tau_c^C}{(1 + \tau_c^C)} \log(\tau_t^C) - \sigma \log(C_t) - \frac{(1 - \sigma) \psi \theta N_{ss}^\theta}{1 - \psi N_{ss}^\theta} \log(N_t) \\
&+ \frac{\tau_c^C}{(1 + \tau_c^C)} E_t \log(\tau_{t+1}^C) = \text{Const}^K + \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(Y_{t+1}) - \beta(1 - \tau_c^K) \alpha \frac{Y_{ss}}{K_{ss}} \log(K_t) \\
&- \beta \tau_c^K \alpha \frac{Y_{ss}}{K_{ss}} E_t \log(\tau_{t+1}^K)
\end{aligned}$$

F Simulation

The simulation of our learning economy is carried out via the following steps:

1. We endow agents with initial beliefs Ω_0 , which coincide with the true pre-policy-change parameter values.

2. Given the beliefs Ω_{t-1} , the perceived steady states are calculated and then used to log-linearize the equilibrium conditions, which together with the estimated policy rules gives the following expectational difference equation:

$$A(\Omega_{t-1})\mathbb{Y}_t = B(\Omega_{t-1})E_t^*\mathbb{Y}_{t+1} + C(\Omega_{t-1})\mathbb{Y}_{t-1} + D\varepsilon_t^*$$

which yields the perceived law of motion (using the RE solution algorithm Gensys by Sims (2001))

$$\mathbb{Y}_t = S(\Omega_{t-1})\mathbb{Y}_{t-1} + G(\Omega_{t-1})\varepsilon_t^*.$$

3. The actual law of motion takes the perceived steady states but uses the true policy parameters $C^{true}(\Omega_t)$ to arrive at the system:

$$A(\Omega_{t-1})\mathbb{Y}_t = B(\Omega_{t-1})E_t^*\mathbb{Y}_{t+1} + C^{true}(\Omega_{t-1})\mathbb{Y}_t + D\varepsilon_t$$

with the actual shock vector ε_t . To solve out for the expectations we use the perceived law of motion to obtain

$$\mathbb{Y}_t = H(\Omega_{t-1})\mathbb{Y}_{t-1} + G(\Omega_{t-1})\varepsilon_t$$

4. Shocks are realized by drawing from a multivariate Gaussian distribution, which together with the transition matrices produced by step 3 determine the macroeconomic outcomes for period t .

5. Observing these outcomes, beliefs are updated via the Kalman filter, which gives Ω_t .

We simulate the economy for each setting 1000 times with a sample length of $T = 100$.

G Additional Figures

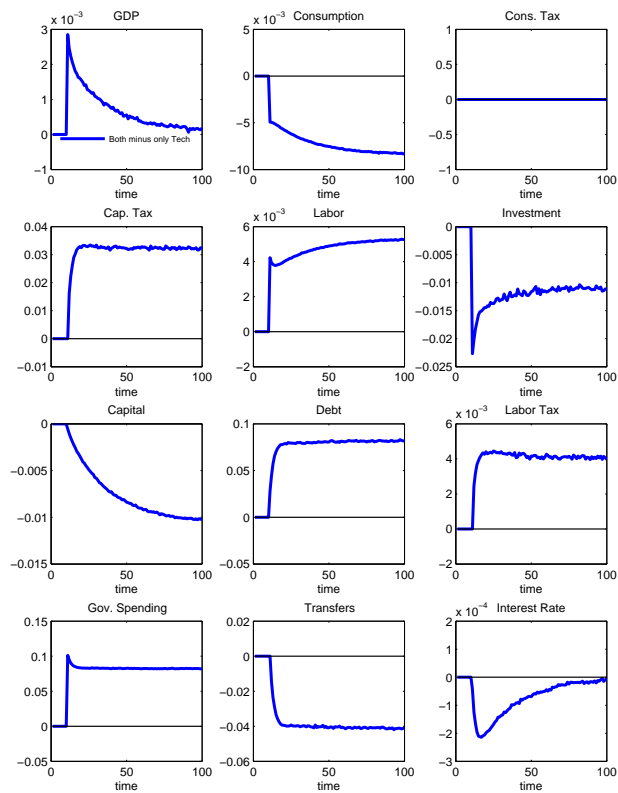


Figure 13: Difference in median (log) outcomes between the RE cases with and without fiscal policy change