Learning about Fiscal Policy and the Effects of Policy Uncertainty

Josef Hollmayr*and Christian Matthes**

Abstract

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 are uncertain about the conduct of fiscal policy and act as econometricians by estimating fiscal policy rules that might change over time.

We find that assuming that agents are not instantaneously aware of the new fiscal policy regime in place leads to substantially more volatility in the short run and persistent differences in average outcomes. We highlight issues that can arise when a policymaker wants to announce a policy change.

From a methodological perspective, we introduce a novel way to model learning in the face of discrete policy changes.

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, Joe Johnson, Michael Krause, Geert Langenus, Thomas Lubik, Stéphane Moyen, Albert Marcet, Taisuke Nakata, Johannes Pfeifer, Ctirad Slavik, Nikolai Stähler and Alexander Wolman, seminar participants at the University of Hamburg and the Bundesbank as well as workshop participants in Mannheim, Bucharest, Gothenburg (EEA) and Minneapolis (Midwest Macro) for helpful comments. This project started while Matthes was visiting the Bundesbank, whose hospitality is gratefully acknowledged. The Online Supplementary Material is available at `https://sites.google.com/site/christianmatthes/`.

1 Introduction

Motivated by the 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). In this paper we ask how this uncertainty arises, what the exact transmission mechanism is and how this uncertainty affects equilibrium outcomes. We propose one model of fiscal policy uncertainty: an RBC model with distortionary taxation and government debt, in which agents are uncertain about the conduct of fiscal policy and act as econometricians to update their beliefs about fiscal policy every period.[1] 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). 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 uncertainty.[2]

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 differences 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)

[1] We use an RBC model that is relatively simple compared to many DSGE models in use today. Nonetheless, DSGE models similar to ours are being used to quantitatively evaluate fiscal policies: See for example Leeper *et al.* (2010).

[2] They 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.".

estimated. Thus, even though we only consider changing a small subset of the fiscal policy coefficients, this uncertainty creeps into other fiscal variables.[3]

There is substantial evidence that fiscal policy rules have changed over time: Davig and Leeper (2007) estimate policy rules for taxes in the US and find substantial time variation. Bianchi and Ilut (2015) estimate a DSGE model that allows for changes in both monetary and fiscal policy rules and again find strong evidence in favor of changes in fiscal policy rules. Given this evidence for changes in fiscal policy over time, we find it natural to study the role of learning about these changes. The aforementioned papers also find evidence in favor of changes in the volatility of policy errors. While we abstract from that possibility in our benchmark, we also study a version of the model where agents consider changes both in policy rule coefficients and the policy error variances.

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 Eusepi and Preston (2011) and Eusepi and Preston (2012) 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.[4] 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.* (2013) 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. Gasteiger and Zhang (2014) introduce distortionary taxation in an RBC model with learning, but in their model agents know the path of fiscal policy instruments and instead have to learn about the dynamics of prices in the economy.

What sets our model apart is the way agents form their beliefs about the stance of fiscal policy. We want the agents in our model to depart from rational expectations as little as possible while simultaneously making the assumption of learning and imperfect information tractable. In contrast to the previously mentioned papers, our agents know the structure of the economy and the behavior of all agents except for the fiscal

[3]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.

[4]We also abstract from the zero lower bound on nominal interest rates. Mertens and Ravn (2014) study the set of equilibria in such a setting under adaptive expectations. Adaptive expectations are also used in a nonlinear model of fiscal and monetary policy interaction by Benhabib *et al.* (2014).

authority, whose behavior is only known up to a finite dimensional vector of policy rule parameters. Households and firms become immediately aware that policy has changed in the period in which the policy change occurs. We think it is reasonable to assume that agents realize when large policy changes like the ones considered here start, since they are announced by policymakers and much discussed in the media. The exact magnitudes of the policy change are often less clear because some of the policy changes are spread out over time and subject to budget approval or other political roadblocks (the government shutdown comes to mind). While our model is too simple to capture all this detail of the political process, we think our setup does capture the fact that agents become aware of policy changes when they happen and then continually learn about how policy has changed. We follow one common approach in the study of policy changes in rational expectations models (see, for example, Uhlig (2010)) and do not allow for anticipation effects in the periods before the actual policy change. We study scenarios in which the government reacts quickly (i.e. within one quarter) to a substantially negative, yet unanticipated, productivity shock, so anticipation effects should not be substantial.

Papers such as Eusepi and Preston (2011) and Eusepi and Preston (2012) instead endow the agents with substantially less knowledge of the economy - their private agents have to learn about all equilibrium relationships, while our agents are only uncertain about the policy rules. 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) and Bianchi and Ilut (2015) 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. In our model, agents have to form beliefs about the policy rule coefficients in the current period.

Our approach to learning follows the approach laid out in Cogley *et al.* (2015), who study a model of monetary policy. Firms and households in our model estimate the coefficients of the policy rules and incorporate both these beliefs and all cross-equation restrictions coming from knowledge of the structure of the economy when making their decisions. The methodological contribution of our paper is how knowledge of the timing of the policy change is incorporated by agents into their estimation problem: They update their beliefs using the Kalman filter with a time-varying covariance matrix for the parameters. Relative to the estimation algorithm used by the agents in Cogley *et al.* (2015), the learning algorithm in our model is substantially faster, opening up the possibility of estimating models with this kind of policy uncertainty in future work.[5]

[5] This comes at the cost of restricting the prior distributions for the parameters to be Gaussian.

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.
Full information analyses of the American Recovery and Reinvestment Act (ARRA) include Cogan *et al.* (2010), Uhlig (2010) and Drautzburg and Uhlig (2011). As will become clear in the discussion of our learning algorithm, the assumption that we study a permanent change is not overly strong since for any $j > 0$, the equilibrium outcomes in our model will be the same for a model with a permanent change and a model where the policy change only lasts j periods. Uhlig (2010) uses a calibration for an ARRA-type scenario that "reflects a skepticism that stimulus spending will truly return back to normal as quickly as envisioned by the ARRA." He goes on to state that "It may be appropriate to build this skepticism into a rational expectations model such as this one...". We agree with Uhlig on the model building, but refrain from using full information rational expectations.

Another strand of the literature that studies uncertainty[6] (or risk) about future fiscal policy is represented by Born and Pfeifer (2014) 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), Bianchi and Ilut (2015) 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 (2014) 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), Bianchi and Ilut (2015) Born and Pfeifer (2014) 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 - uncertainty about fiscal policy does not change the private agents' behavior in each period once they have formed their beliefs.

[6] When 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). Knightian uncertainty makes agents focus on the worst case scenario, while our approach makes agents focus on the most likely scenario for fiscal policy. An analysis of optimal fiscal policy when agents are learning is provided by Caprioli (2015). Both papers use a smaller set of fiscal policy instruments than we do.

2 Model

We use a real model of a closed economy without habits and other frictions. The model is based on Leeper *et al.* (2010). The only deviation from the standard RBC model (King *et al.* (1988)) besides the learning is the rich fiscal sector with distortionary taxation, government spending, and transfers. This section describes how private agents act once they have formed beliefs and how the government acts in our economy. In later sections, we describe the belief formation and then derive equilibrium dynamics. First-order conditions and the complete log-linearized model can be found in the Online Supplementary Material.

2.1 Households

Households each period maximize their expected utility conditional on their beliefs about fiscal policy rule parameters. [7] The instantaneous utility function of the representative household takes the following form:

$$U_t = \frac{C_t^{1-\sigma}}{1-\sigma} - \frac{L_t^{1+\phi}}{1+\phi} \tag{1}$$

with consumption C_t and labor L_t. In 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 in each period t they maximize the infinite sum of discounted utility

$$E_t^* \sum_{j=0}^{\infty} \beta^{t+j} U_{t+j} \tag{2}$$

The E_t^* operator is the expectations operator under the perceived probability measure at time t - expectations each period are formed conditional on beliefs about fiscal policy rules.

The utility maximization each period is subject to the sequence of the following

[7] We thus 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.

constraints which have to hold for all periods:

$$C_t(1+\tau_t^C) + B_t + I_t = W_t L_t(1-\tau_t^L) + (1-\tau_t^K)R_t^K K_{t-1} + R_{t-1}B_{t-1} + Z_t \tag{3}$$

$$K_t = (1-\delta)K_{t-1} + I_t \tag{4}$$

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:

$$Y_t = \exp(A_t)K_{t-1}^{\alpha}L_t^{1-\alpha} \tag{5}$$

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):

$$A_t = \rho_a A_{t-1} + \epsilon_t^A \tag{6}$$

2.3 Government

The government in our model makes decisions according to log-linear decision rules and respects the government budget constraint. This government budget constraint is given by:

$$B_t = B_{t-1}R_{t-1} - R_t^K K_t \tau_t^K - W_t L_t \tau_t^L - C_t \tau_t^C + G_t + Z_t \tag{7}$$

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 strong.[8] Each policy rule features an iid Normal policy error denoted by ϵ_t^i for each policy instrument $i \in \{G, Z, \tau^C \tau^L, \tau^K\}$. These error terms cannot be observed by firms and households in our model. The presence of these policy errors makes the learning problem of agents in our models interesting: Firms and households have to estimate if changes in a policy instrument are driven by persistent changes in policy rule parameters or exogenous iid shocks in each policy rule.
Government expenditure G_t is determined by the following rule:

$$\log(G_t) = G_c - \rho_{g,y} \log(Y_{t-1}) - \rho_{g,b} \log(B_{t-1}) + \epsilon_t^G \tag{8}$$

Transfers Z_t are given by:

$$\log(Z_t) = Z_c - \rho_{z,y} \log(Y_{t-1}) - \rho_{z,b} \log(B_{t-1}) + \epsilon_t^Z, \tag{9}$$

Consumption taxes follow an iid process to match their properties in US data (see Leeper *et al.* (2010)):

$$\log(\tau_t^C) = \tau_c^c + \epsilon_t^C \tag{10}$$

Labor and capital taxes are set as a log-linear function of log output and log debt in the previous period:

$$\log(\tau_t^L) = \tau_c^l + \rho_{L,y} \log(Y_{t-1}) + \rho_{L,b} \log(B_{t-1}) + \epsilon_t^L \tag{11}$$

$$\log(\tau_t^K) = \tau_c^k + \rho_{K,y} \log(Y_{t-1}) + \rho_{K,b} \log(B_{t-1}) + \epsilon_t^K \tag{12}$$

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.
We include a consumption tax in our model to make the agents' estimation problem more realistic. As will become clear in the section that describes the learning algorithm, private agents entertain a view of the world where changes in parameters could be correlated across policy rules. Thus reducing the number of tax instruments and number of policy rule parameters might have an influence on the private agents' estimates even if there is no direct feedback from endogenous variables to the consumption

[8] For a discussion of the link between simple fiscal policy rules like the ones employed here and optimal fiscal policy, see Kliem and Kriwoluzky (2014).

tax rate.
In our policy experiments we will consider one-time changes in some of the policy rule coefficients.

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 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 the Online Supplementary Material. The analysis is carried out via a Monte Carlo simulation - 1000 simulations of 100 periods each. We start off the simulations at the steady state associated with the original policy rule parameters. In period 9, a negative technology shock hits in all simulations 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 once and for all. We consider a permanent policy change in which only the intercept in the policy rule G_c changes from the original value G_c^{old} to G_c^* 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 policy rules)[9].
We pick the size of the change in G_c using the following thought experiment: Given the

[9]This 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. Under full information rational expectations, agents would immediately realize that this policy change implies a new steady state of the economy. Under learning agents only realize over time how much the steady state has changed.

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 (8) in levels at the original steady state:

$$G = \exp(G_c^{old}) Y^{-\rho_{g,y}} B^{-\rho_{g,b}} \tag{13}$$

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:

$$G + 0.01Y = \exp(G_c^*) Y^{-\rho_{g,y}} B^{-\rho_{g,b}} \tag{14}$$

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

We choose to model our policy experiment as a change in the policy rule rather than a pre-determined sequence of exogenous shocks, but since we could pick a sequence of shocks to deliver the same fiscal policy instruments these two modeling approaches deliver the same outcome under learning, as we discuss further in section 4. This is not the case for full-information rational expectations, where there are cross-equation restrictions between the private sector's decision rules and the actual policy rules.

4 Learning about Fiscal Policy

The agents in our model act as Bayesian econometricians. All private agents share the same beliefs. They observe all relevant economic outcomes and use those observations to estimate the coefficients of the policy rules (8)-(12). 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.

Our agents are endowed with initial beliefs that are centered around the true values of the parameters before any parameter changes that we consider in the various policy experiments later. Given those initial beliefs and the observed policy instruments and right-hand side variables for each policy rule, agents can update their beliefs every period according to Bayes' rule.

We will now describe an environment where using Bayes' rule amounts to using the Kalman filter, which makes our assumption of agents acting as Bayesian econometricians operational. To do so, we first stack the policy instruments at time t:

$$\tau_t = \begin{pmatrix} \log(G_t) \\ \log(Z_t) \\ \log(\tau_t^C) \\ \log(\tau_t^L) \\ \log(\tau_t^k) \end{pmatrix} \tag{15}$$

Next, we denote by Ω_t the 13-dimensional vector of coefficients of all fiscal policy rules (8)-(12):

$$\Omega_t = \begin{pmatrix} G_{c,t} \\ \rho_{g,y,t} \\ \rho_{g,b,t} \\ Z_{c,t} \\ \rho_{z,y,t} \\ \rho_{z,b,t} \\ \tau^c_{c,t} \\ \tau^l_{c,t} \\ \rho_{L,y,t} \\ \rho_{L,b,t} \\ \tau^k_{c,t} \\ \rho_{K,y,t} \\ \rho_{K,b,t} \end{pmatrix}$$

t subscripts attached to coefficients identify estimates of that coefficient at time t. We assume for simplicity that agents know the volatility of the errors in the policy rules, which means that agents can in fact use the Kalman filter for inferring the values of all other policy rule parameters. Furthermore, we denote by η_t the vector of iid Gaussian disturbances in the fiscal policy rules (with covariance matrix Σ_η) and by X_t a matrix whose entries are either 0 or one of the right hand-side variables in the policy rules (a constant, $\log(Y_t)$ or $\log(B_t)$). The full X_t matrix is given in the Online Supplementary Material. Then the observation equation for the Kalman filter that links the unobserved parameters that agents need to estimate with the observed variables can be written as:

$$\tau_t = X_{t-1}\Omega_t + \eta_t \tag{16}$$

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? This perceived law of motion will serve as the state equation in the Kalman filter used by the agents. We want our agents to act as sophisticated applied economists. We therefore endow our agents with a perceived law of motion of the parameters that has been the benchmark in the literature on time-varying coefficient models in empirical macroeconomics (such as Cogley and Sargent (2005) or Primiceri (2005)) [10]. Furthermore, the perceived law of motion our agents are endowed with nests the changes in policy rule parameters that we will study later. Thus, our agents use a correctly specified model. Endowing agents with a view of the world that does not nest the true data-generating process would preclude them from ultimately learning the true parameter values and knowing the true equilibrium dynamics. While this might certainly be interesting, we do not want to force our agents to not be able to learn the truth asymptotically. Our agents know at what time the policy rule coefficients change - they just do not know *which 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:

$$\Omega_t = \Omega_{t-1} + \mathbf{1}_t \nu_t \tag{17}$$

$\mathbf{1}_t$ is an indicator function that equals 1 in the period in which fiscal policy changes[11] and ν_t is a Gaussian vector with mean 0 for each element.
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

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

[11] We 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.

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.

4.1 Subjective Uncertainty and the Perceived Model of Fiscal Policy

If we were to set the covariance matrix 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. We will measure uncertainty by the dispersion of beliefs (i.e. estimates of policy coefficients) *across simulations* after the policy change. This dispersion is a function of the covariance matrix of ν_t - the larger this matrix is, the larger will be the impact of observed data on beliefs directly after the policy change. We can not directly build a comparable measure from data since we compare beliefs across simulations. We can, however, look at other measures of policy uncertainty to at least informally guide our choice of increasing the subjective uncertainty at the time of the policy change. Baker *et al.* (2012) construct various measures of policy uncertainty. Figure 1 plots one of their indices of fiscal uncertainty based on the tax code expiration data. The index thus gives an objective measure of how much uncertainty there is about whether fiscal policies in place are temporary or permanent. This is *exactly* the same question the agents in our model try to answer. The objective measure of 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. Our measure of subjective uncertainty will show the same pattern. [12] Very broadly speaking, we think of the objective measure as a signal of uncertainty that can inform private agents' priors as they have access to this information via news outlets. This is why we chose to adapt a similar pattern for the subjective uncertainty we endow our agents with.

[12] The 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. We do not focus on this measure since our model features a representative agent and thus has no role for forecast disagreement.

We set the covariance matrix for ν_t, Σ_ν, to a scaling factor s times a diagonal matrix with the ith element on the diagonal being equal to the square of the ith 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. We provide the relevant Kalman filter equations in the Online Supplementary Material.

We could have alternatively assumed that our agents use recursive least squares algorithms to update their beliefs. For a comparison of learning when using the Kalman filter versus learning when using the common recursive least squares approach, see Sargent and Williams (2005). From a frequentist perspective, the Kalman filter delivers the optimal estimator in scenarios with time varying system matrices such as those studied here - see Hamilton (1994). It seems reasonable to assume that our agents use an optimal estimator as a benchmark.

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 currently estimated level forever more [13]. This is a common assumption in the literature on learning, see for example Milani (2007) or Sargent *et al.* (2006).[14]

The agents' prior about the initial policy rule parameters is assumed to be Gaussian since we want to use the Kalman filter. The prior mean is set to the pre-policy change true parameters. We calibrate the initial covariance matrix of the estimators to be a diagonal matrix (with all off-diagonal elements set to zero) so that the initial standard deviation for each parameter is equal to 10 percent of the prior mean. 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

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

[14] 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."

matrix in the first 9 periods of the simulations, the exact calibration of the initial covariance matrix is not critical. The restriction that the prior is Gaussian is necessary to enable us to use the Kalman filter - the other restrictions are made for convenience. 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 Online Supplementary Material for details).

4.2 Solving for Equilibrium Dynamics

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 equations that can be solved for the perceived law of motion in the economy[15]:

$$A(\Omega_{t|t-1})\mathbb{Y}_t = B(\Omega_{t|t-1})E_t^*\mathbb{Y}_{t+1} + C(\Omega_{t|t-1})\mathbb{Y}_{t-1} + D\varepsilon_t^* \tag{18}$$

We follow Cogley *et al.* (2015) by log-linearizing each period around the perceived steady state. Cogley *et al.* (2015) interpret this as a behavioral assumption: Private agents in the model need to solve their perceived model. To do so, they log-linearize around their best estimate of the steady state, the perceived steady state. Alternatively, one can view this as a computational procedure to better approximate the behavior of the true non-linear model when it is far away from the full-information steady state.

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.

The coefficient matrices depend on $\Omega_{t|t-1} = E_{t-1}(\Omega_t)$, which is the one period ahead forecast of the policy rule coefficients that is produced by the Kalman filter. Because of the random walk nature of the law of motion for the policy rule coefficients in the state space system used by the Kalman filter, this estimate is not updated as long as no new data arrives: $\Omega_{t|t-1} = E_{t-1}(\Omega_t) = E_{t-1}(\Omega_{t-1})$.

The estimated policy coefficients enter directly as entries in $C(\Omega_{t|t-1})$ and indirectly in $A(\Omega_{t|t-1})$ and $B(\Omega_{t|t-1})$ because they influence the perceived steady state around which we log-linearize each period, thus influencing the coefficients in the log-linearized

[15] This derivation follows Cogley *et al.* (2015). We also borrow their use of a projection facility: If no stable perceived law of motion exists, agents use the previous period's estimates.

equilibrium conditions.

The perceived policy shocks can be derived by replacing the true policy coefficients in equations (8)-(12) with their estimated counterparts that are elements of $\Omega_{t|t-1}$.

This expectational difference equation can be solved using standard algorithms to yield the perceived law of motion for the economy at time t:

$$\mathbb{Y}_t = S(\Omega_{t|t-1})\mathbb{Y}_{t-1} + G(\Omega_{t|t-1})\varepsilon_t^* \tag{19}$$

$S(\Omega_{t-1})$ solves the following matrix quadratic equation[16]:

$$S(\Omega_{t|t-1}) = (A(\Omega_{t|t-1}) - B(\Omega_{t-1})S(\Omega_{t|t-1}))^{-1}C(\Omega_{t|t-1}) \tag{20}$$

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

$$G(\Omega_{t|t-1}) = (A(\Omega_{t|t-1}) - B(\Omega_{t|t-1})S(\Omega_{t|t-1}))^{-1}D \tag{21}$$

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 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 replace the estimated policy coefficients in $C(\Omega_{t|t-1})$ with the *true* policy coefficients. We call this matrix $C^{true}(\Omega_{t|t-1})$. It is still a function of $\Omega_{t|t-1}$ because we log-linearize around the perceived steady state, which is a function of $\Omega_{t|t-1}$. This can have an effect on the entries in rows of $C^{true}(\Omega_{t|t-1})$ that are not associated with the policy rules. Then the actual law of motion solves:

$$A(\Omega_{t|t-1})\mathbb{Y}_t = B(\Omega_{t|t-1})E_t^*\mathbb{Y}_{t+1} + C^{true}(\Omega_{t|t-1})\mathbb{Y}_{t-1} + D\varepsilon_t \tag{22}$$

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

$$\mathbb{Y}_t = H(\Omega_{t|t-1})\mathbb{Y}_{t-1} + G(\Omega_{t|t-1})\varepsilon_t \tag{23}$$

[16] The 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.

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

$$H(\Omega_{t|t-1}) = S(\Omega_{t|t-1}) + (A(\Omega_{t|t-1}) - B(\Omega_{t|t-1})S(\Omega_{t|t-1}))^{-1}(C^{true}(\Omega_{t|t-1}) - C(\Omega_{t|t-1})) \tag{24}$$

Equation (23) gives some insights into the relative roles played by the perceived law of motion and the actual policy rules: actual policy rules only matter insofar as they deliver a given path of the policy instruments τ_t. Any two actual policy rules that give the same sample of τ_t from $t = 0$ to some time period T will result in the same equilibrium dynamics because, given that these two rules produce the same path for τ_t, they also induce the same beliefs for agents in this sample. Then the only difference in the actual law of motion for the two policy rules are the equations for the policy instruments themselves. While those equations are different, they yield the same sample paths for τ_t by our assumption. Thus equilibrium outcomes for this sample from $t = 0$ to T are the same.[17] This also implies that it does not matter for the sample ending in period T whether or not the true policy rule will change again in any period $T + s$, $s > 0$. A policy experiment that features only a temporary change in fiscal policy will thus yield the same equilibrium outcomes as the model with a permanent change as long as the temporary change is in place. A model with a temporary policy change would feature a second learning transition back to the original policy values. Depending on when the reversion back to the original rule happens and if we assume $\mathbf{1}_t = 1$ for the period when the reversion begins, this transition back to the original policy rule parameters could be faster or slower than the first learning transition that we focus on in this paper. For the sake of brevity, we focus on the case with only one learning transition in this paper.

The shape of the perceived policy rules could matter substantially for equilibrium outcomes. To check this, we below also solve a version of the model where agents are uncertain about changes in the volatility of policy errors as well as policy rule coefficients. Davig and Leeper (2007) identify changes in both coefficients and error volatilities when estimating Markov-switching policy rules for taxes in the US.

To allow agents to entertain uncertainty about the volatility of the errors in the policy rule (which is how Born and Pfeifer (2014) and Fernandez-Villaverde *et al.* (2011) model fiscal policy uncertainty), we augment our framework by letting agents consider a set of J possible covariance matrices $\Sigma_{\eta,j}$, $j = 1, ..., J$ for the policy errors. Note that

[17] In particular, this means that, conditional on a given path of realized policy instruments, it does not matter whether these paths of the policy instruments were generated by a policy rule with changing coefficients or stochastic volatility.

$\Sigma_{\eta,j}$ can not be directly calculated using the Kalman Filter output because it can not be written as an unobserved state (it multiplies a vector of unobserved shocks).

We think this extension could be of interest more generally in learning models where some parameters can be estimated directly via the Kalman filter, while others cannot. It is similar in spirit to Cho and Kasa (2012), where agents also entertain multiple models while learning, but our approach exploits that some parameters can be easily estimated using the Kalman filter. The agents consider J laws of motion for the policy instruments:

$$\tau_t = X_{t-1}\Omega_{t,j} + \eta_{t,j} \tag{25}$$

where $E\eta_{t,j}\eta'_{t,j} = \Sigma_{\eta,j}$. The agents solve J filtering problems in parallel every period and thus have access to J estimates of the policy rule coefficients $\Omega_{t|t-1,j}$, $j = 1, ..., J$. We then need to posit a model selection rule $\Omega^*_t = F(\Omega_{t|t-1,1}, ..., \Omega_{t|t-1,J})$. In our application below, we let agents choose the estimate $\Omega_{t|t-1,j}$ associated with the covariance matrix $\Sigma_{\eta,j}$ that yields the highest likelihood values (which can be recursively computed using the Kalman filter). Conditional on an estimate Ω^*_t, we can solve the model as before, replacing $\Omega_{t|t-1}$ in the equations above with Ω^*_t.

5 Calibration

This section describes the calibration of the standard (i.e. not related to learning) parameters of the model. 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 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 [18] and the depreciation rate of capital is set at 0.025, which is equivalent to a total annual depreciation of 10 percent. 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.

We set the constants in the policy rules to obtain the same steady state values as Leeper

[18] This 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)

et al. (2010) for tax rates, government spending over GDP, and debt capital over GDP. The steady state values for the consumption tax, the capital tax, and the labor tax are therefore 0.03, 0.25, and 0.19, 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).
All parameter values and steady state values are reported in the Online Supplementary Material.

6 Results

6.1 A Roadmap

We will first present results for the full-information rational expectations case. Full-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. A different interpretation of this rational expectations case is one where agents are learning, but there is a *credible* announcement of policy change so that beliefs directly collapse to the correct post-policy change values. The alternative learning specifications studied below feature agents that realize that there was a policy change, but they do not view announcements of new policy rule parameters by the government as credible. Instead they use data and Bayes' theorem to update their beliefs.
We will show how learning affects equilibrium outcomes 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. In addition, we ask if learning would have any effects if there were no actual policy change.
We also explore 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?
Finally we also let agents entrain the possibility of changes in the volatility of shock in the policy rules.
In the Online Supplementary Material, we present two additional robustness checks: We 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. There, we also 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 expectations[19]. 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 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 3), 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 4 shows a summary of the outcomes in that environment. The bottom row 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 rule[20]. Agents quickly pick up

[19] Mean outcomes are very similar.

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

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)[21]. 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, too[22]. Thus, agents in this setup learn fast and the largest uncertainty in quantitative terms (that around G_c) disappears reasonably quickly[23]. Does learning have any effect on outcomes then?

The top row shows how average outcomes change relative to full-information rational expectations[24]: We plot the cumulated difference between median outcomes under learning and under rational expectations relative to the original steady state. Our focus is on cumulative outcomes because they more clearly highlight the differences between full information rational expectations and our learning setup. We thus plot

$$Diff_j^W = \sum_{t=1}^{j} \frac{(W_t^{learning} - W_t^{RE})}{W} \tag{26}$$

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 expectations[25]. 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 quan-

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

[22]The 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.

[23]Our results indicate that convergence to rational expectations does occur. For a comparison of theoretical convergence results under Kalman filter learning and least squares learning see Sargent and Williams (2005).

[24]Note 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.

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

titatively 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 row of figure 4 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 turns 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 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 11. The figure plots the actual median path of the capital tax 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 quite 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 shocks 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 analyze scenarios in which the agents have more or less prior uncertainty than in our benchmark. To do so, we vary the scaling parameter s that multiplies the prior covariance matrix of the parameters. The actual policy change is the same that is studied in the benchmark case, but we now analyze environments in which agents think that this particular policy change is more likely than in the benchmark case (it represents a 1-standard-deviation shock) or less likely (it represents a 3-standard-deviation shock). The results are plotted in figures 5 and 6.

The qualitative patterns that emerge remain 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 7 tackles that question. Comparing this figure with figure 4, 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. We can interpret this situation as policymakers making an announcement but not following through or policymakers making an announcement about a future policy change which makes agents more uncertain about the nature of the policy rules currently in place. A policymaker that makes misleading announcements about policy changes can thus induce substantial volatility in the economy even if the policymaker does not in fact change policy rules.

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 8. In this case volatilities are dampened relative to our benchmark case depicted in figure 4, 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 9 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 Perceived Changes in Volatilities

We now let agents estimate both the coefficients Ω_t and the covariance matrix of the shocks in the policy rules, Σ_η. For simplicity, we focus here on perceived changes in volatility in the policy rules for capital taxes and government spending. The policy experiment is the same as in the benchmark - we only endow agents with a richer model of fiscal policy. We choose evenly spaced 9 point grids for each of those volatilities with a lower bound of half the true value (which we keep unchanged from the benchmark case) and an upper bound of 1.5 times the true value. This gives us 81 different combinations of Σ_η. Figure 10 plots the results. Qualitatively nothing changes from our benchmark when we allow agents to entertain the possibility of changes in volatility.

7 Conclusion

Analyses of large changes in policy will most likely come to wrong conclusions if we as researchers do not take into account how agents form expectations and how credible policy announcements are. Our results represent cautionary tales that back up this claim. 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. While we do not study the optimal choice of policy instruments in our framework and are thus limited in the policy recommendations we can make based on our results, we do find that the views that agents hold about the magnitude of possible policy changes have 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. Studying the optimal policy choice in conjunction with the optimal choice of announcements when agents are learning will require a model of agents that are learning from both data and (possibly incredible) announcements made by policymakers. Our results highlight what can happen when private agents do not put any weight on announcements made by policymakers (our benchmark case) or if the announcements made by policymakers only lead private agents to put more prior weight on larger possible policy changes.

References

Baker, S. R., Bloom, N., and Davis, S. J. (2012). Has economic policy uncertainty hampered the recovery? Working papers, Becker Friedman Institute for Research In Economics.

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**(C), 220–238.

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.

Bianchi, F. and Ilut, C. (2015). Monetary/Fiscal Policy Mix and Agents' Beliefs. Working paper.

Born, B. and Pfeifer, J. (2014). Policy risk and the business cycle. *Journal of Monetary Economics*, **68**(C), 68–85.

Caprioli, F. (2015). Optimal fiscal policy under learning. *Journal of Economic Dynamics and Control*, **58**, 101–124.

Cho, I.-K. and Kasa, K. (2012). Model Validation and Learning. Discussion Papers dp12-07, Department of Economics, Simon Fraser University.

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. (2015). Optimized Taylor rules for disiflation when agents are learning. *Journal of Monetary Economics*, **72**, 131–147.

Davig, T. and Leeper, E. M. (2007). Fluctuating Macro Policies and the Fiscal Theory. In *NBER Macroeconomics Annual 2006, Volume 21*, NBER Chapters, pages 247–316. National Bureau of Economic Research, Inc.

Davig, T., Leeper, E. M., and Walker, T. B. (2010). Unfunded liabilities and uncertain fiscal financing. *Journal of Monetary Economics*, **57**(5), 600–619.

Drautzburg, T. and Uhlig, H. (2011). Fiscal Stimulus and Distortionary Taxation. NBER Working Papers 17111, National Bureau of Economic Research, Inc.

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*, **25**(4), 358–379.

Eusepi, S. and Preston, B. (2012). Debt, Policy Uncertainty, And Expectations Stabilization. *Journal of the European Economic Association*, **10**(4), 860–886.

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.

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

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.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.

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. (2014). Toward a Taylor Rule for Fiscal Policy. *Review of Economic Dynamics*, **17**(2), 294–302.

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.

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.

Mertens, K. and Ravn, M. O. (2014). Fiscal Policy in an Expectations Driven Liquidity Trap. *forthcoming, Review of Economic Studies*.

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.

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.

Smets, F. and Wouters, R. (2007). Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. *American Economic Review*, **97**(3), 586–606.

Uhlig, H. (2010). Some Fiscal Calculus. *American Economic Review*, **100**(2), 30–34.

Figures

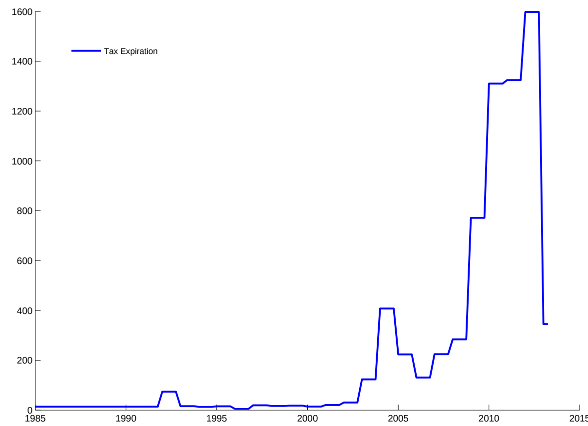

Figure 1: Fiscal uncertainty index by Baker *et al.* (2012)

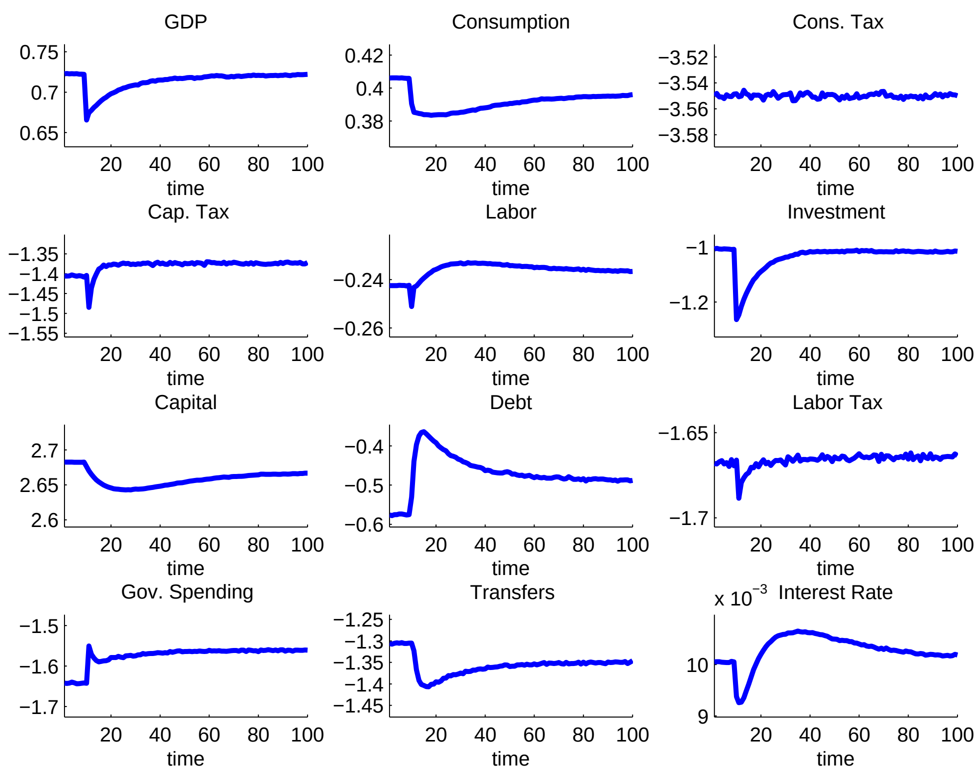

Figure 2: Log outcomes under rational expectations

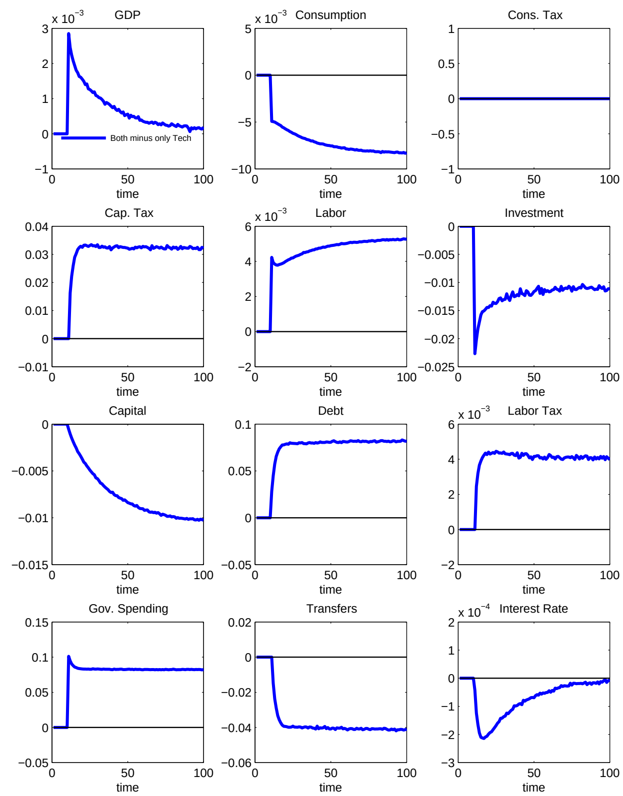

Figure 3: Difference in median (log) outcomes between the RE cases with and without fiscal policy change

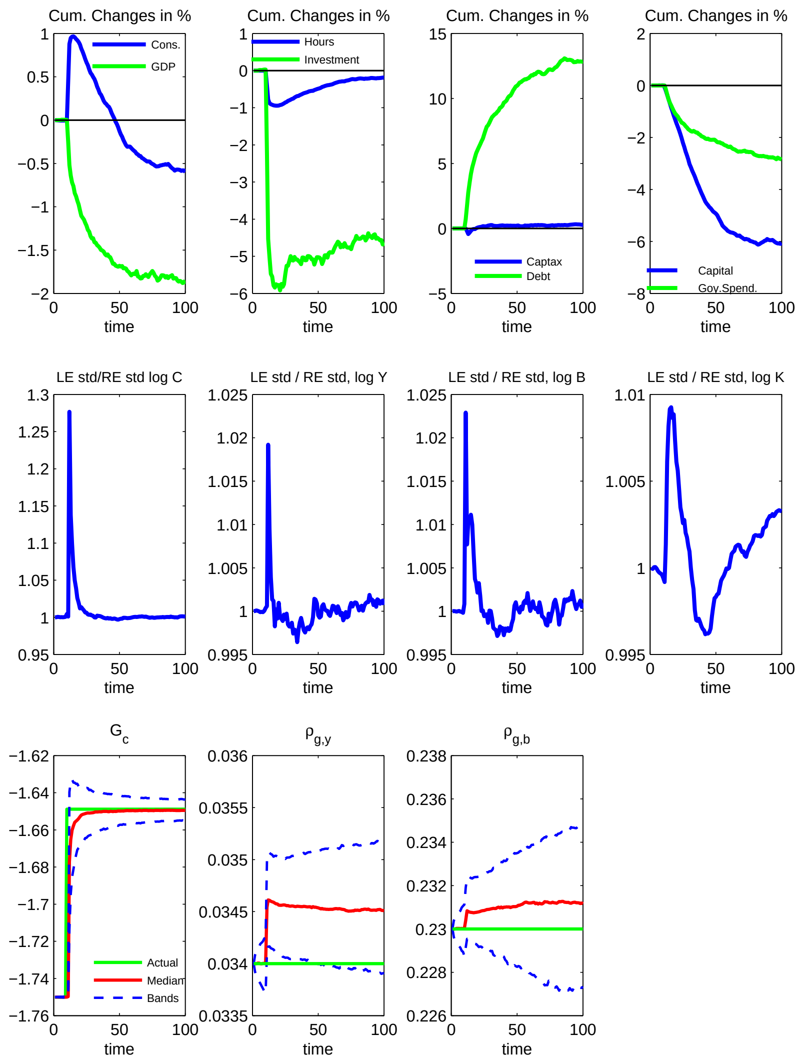

Figure 4: Summary of outcomes under learning

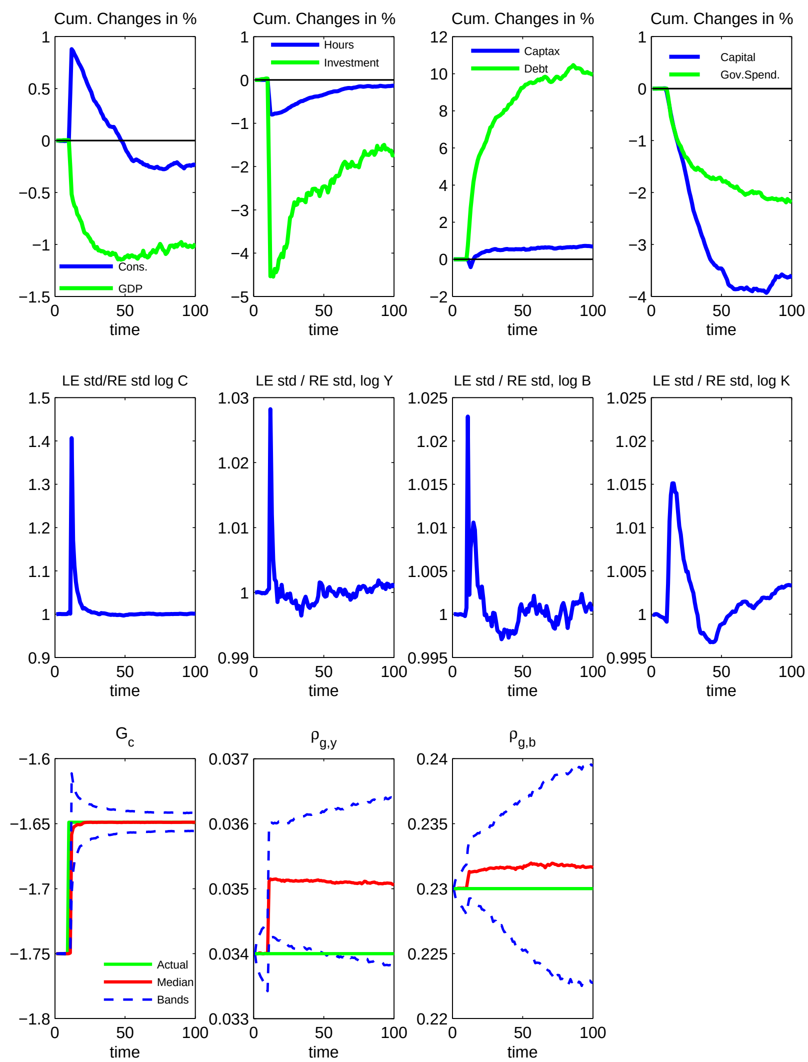

Figure 5: Summary of outcomes under learning, 1-standard-deviation case

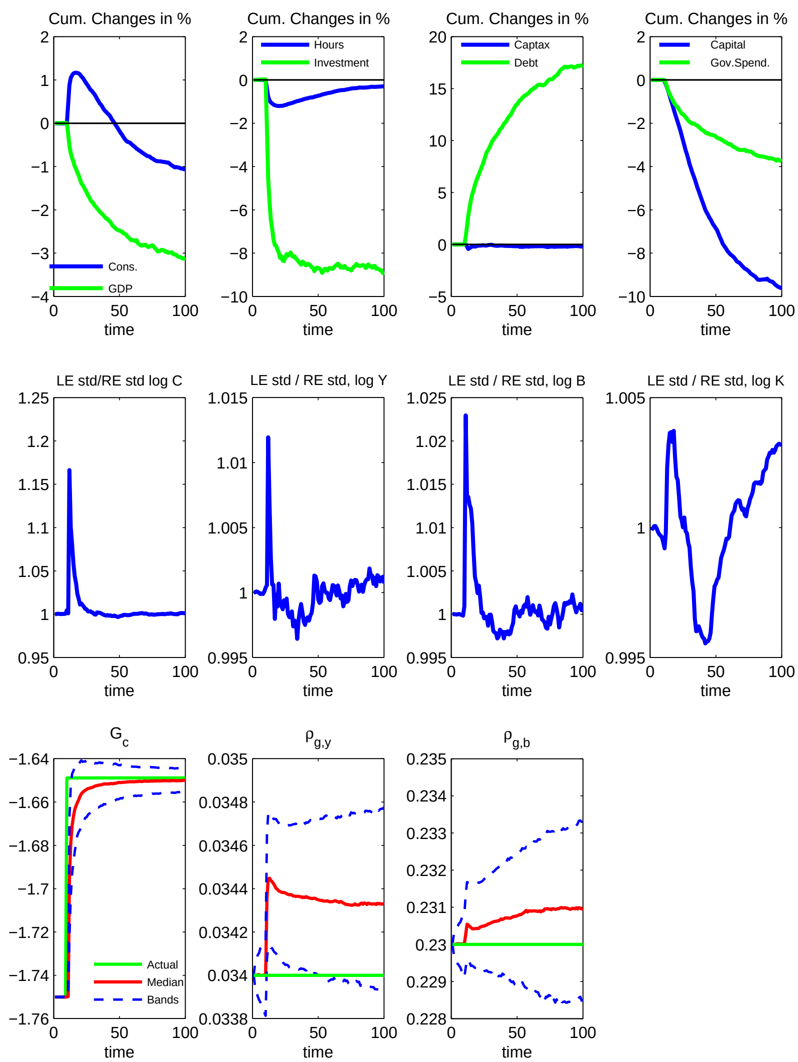

Figure 6: Summary of outcomes under learning, 3-standard-deviations case

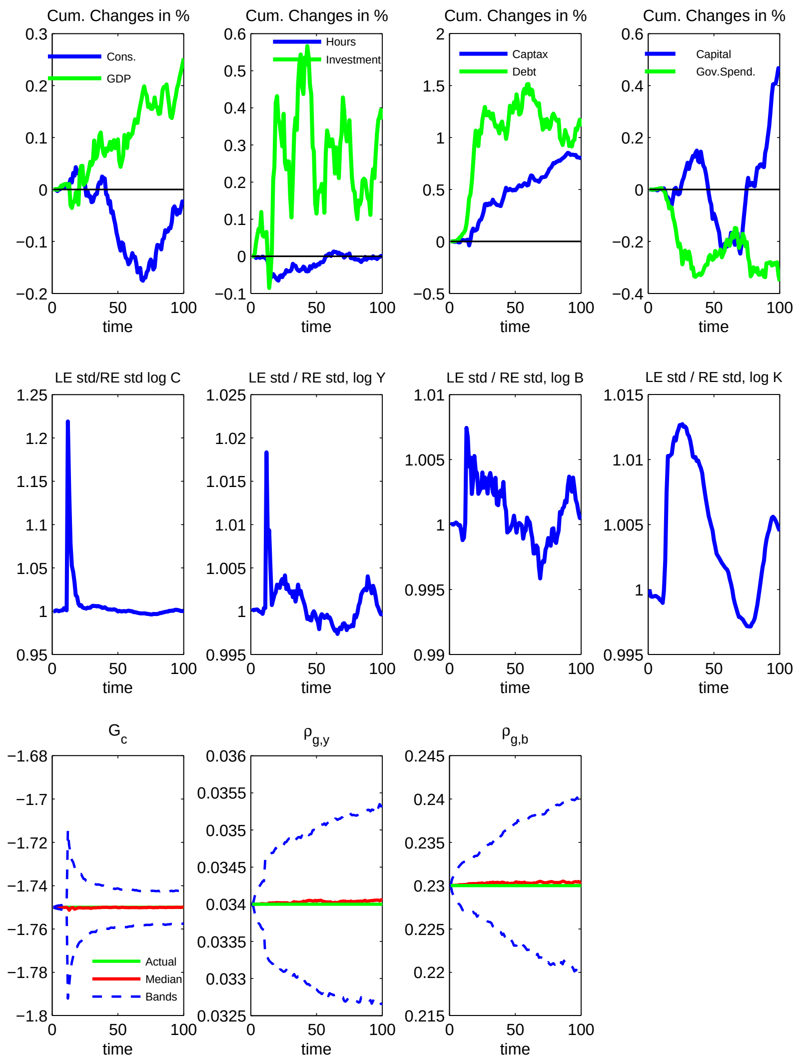

Figure 7: Summary of outcomes under learning when there is no fiscal policy change

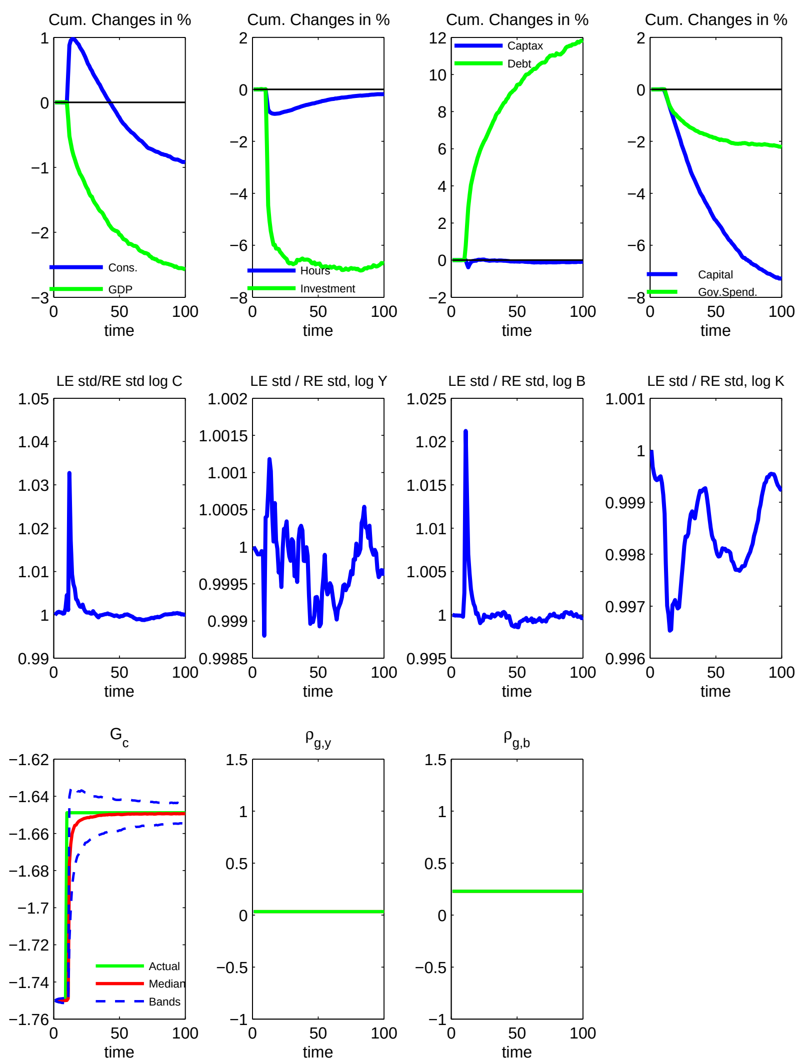

Figure 8: Summary of outcomes under learning when agents only need to learn about G_c

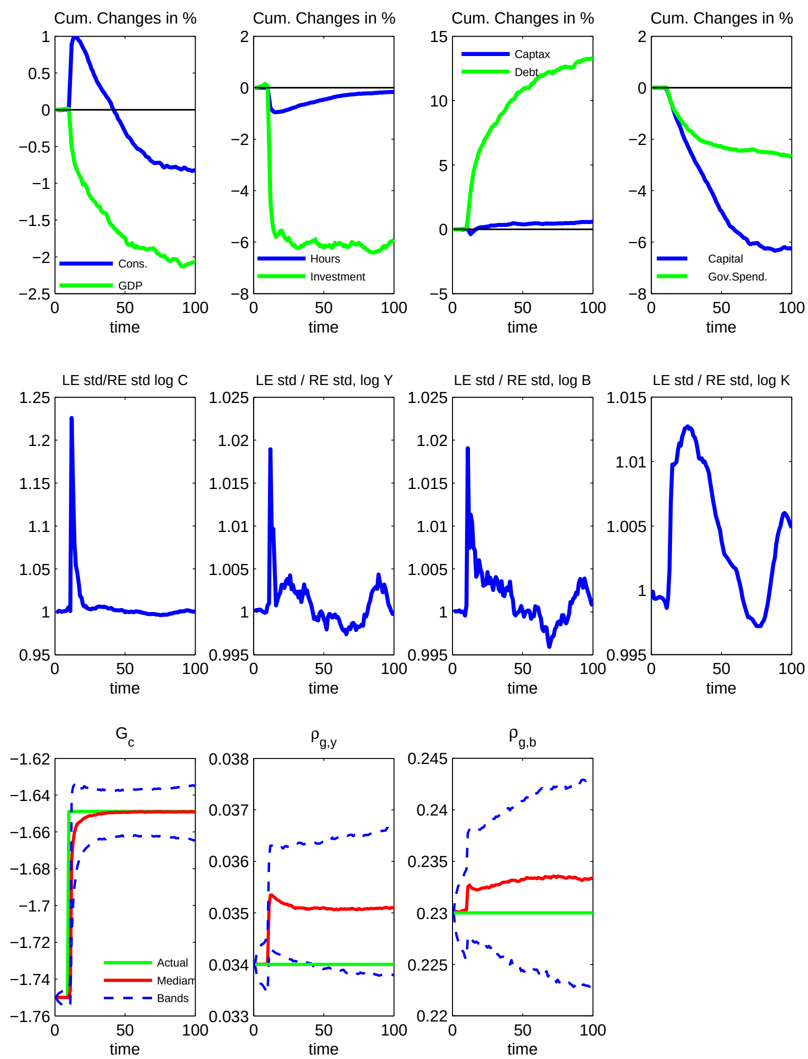

Figure 9: Summary of outcomes under learning when agents think that consumption enters the policy rule for government spending

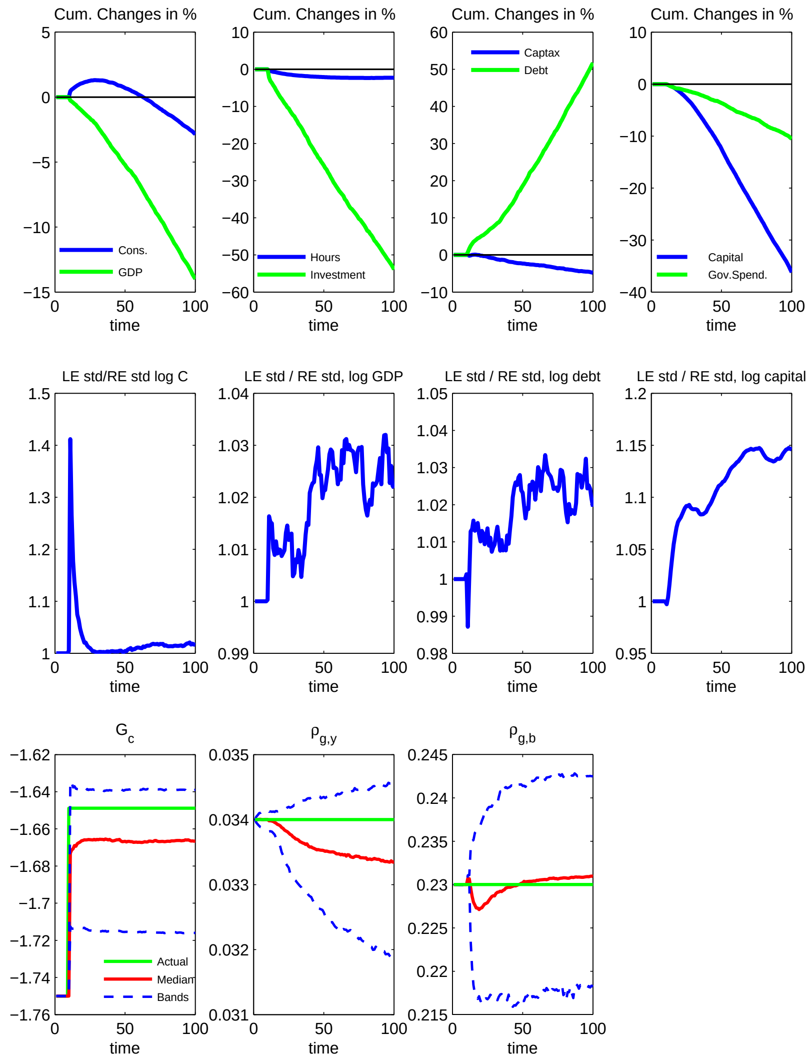

Figure 10: Summary of outcomes under learning when volatilities are unknown

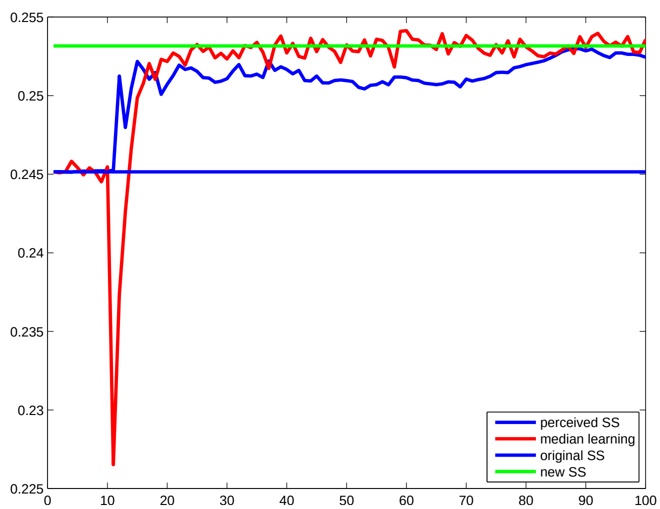

Figure 11: The capital tax rate