Performance of Simple Interest Rate Rules Subject to Fiscal Policy

Nigel McClung

3 October 2018
(Link to Most Recent Draft)

Abstract

This paper examines the performance and robustness of simple interest rate rules in models with rational agents or learning agents subject to: (1) permanent or recurring active fiscal policy; and/or (2) the presence of long-term government debt. My analysis indicates that the "global" responsiveness of fiscal policy to debt determines the optimal monetary policy response. When fiscal policy is globally passive or globally active the optimal monetary policy rule typically features time-invariant coefficients with high inflation reaction coefficients in globally passive models and interest rate pegs in globally active models. In cases where fiscal policy features balanced or strong switching between active and passive fiscal policy stances, the optimal monetary policy rule features switching coefficients. These results extend to models with adaptive learning, including a hidden Markov model of learning never seen before in the literature.

1. Introduction

This paper examines the performance and robustness of simple monetary policy rules in New Keynesian models with: (1) permanent or occasionally non-Ricardian fiscal policy; (2) long-term government debt. In these models, time-variation in the fiscal policy stance on debt is captured by Markov-switching in fiscal policy rule coefficients. Policy performance is measured in terms of a loss function that equals some weighted average of the variance of inflation and the variance of the output gap, and central banks are tasked with selecting implementable interest rate rules that minimize loss taking fiscal policy as given. While a primary goal of this project is to identify optimal interest rate rules in models with rational expectations, we endeavor to construct policies that are robust to parameter and model uncertainty, and that also perform well in models with constant gain learning.

Our contributions provide answers to three questions. First, should monetary policymakers respond to time-varying fiscal policy stances on the debt by implementing time-varying monetary policy? A growing body of work argues that fiscal policy stances on the debt are time-varying.[1] By offering a potential answer to this question, we may better understand the importance of precisely identifying the timing and magnitude of fiscal policy regime shifts.

We find that the long-run or "global" responsiveness of fiscal policy to government debt determines whether the optimal interest rate rule is time-varying. When fiscal policy is "globally active" or "globally passive", central banks typically lack reason to track the fiscal policy stance, and should instead implement time-invariant interest rate rules. In the case of globally active policy, we find that there are strikingly large regions of the parameter space for which time-invariant interest rate pegs are optimal. For "globally switching" policies that feature more balanced fiscal regimes, the optimal policy is both time-varying and parameter-dependent. Hence, we identify certain cases where monetary policymakers should track the timing and magnitude of fiscal policy regime changes.

Second, do the optimal policy rules under rational expectations perform well in models with adaptive learning? By answering this question, we hope to identify conditions under which the optimal rational expectations policy is robust to misspecifications about the true model of expectations formation. We find that interest rate pegs are also optimal in globally active models with least squares learning agents. Moreover, our learning agents in globally switching and passive models tend to prefer inflation reaction coefficients in passive fiscal regimes that are higher than the optimal rational expectations coefficients.

These last points are demonstrated in a constant gain learning model with ob-

[1] see Bianchi (2012), Bianchi and Ilut (2017), Davig and Leeper (2006, 2011), Gonzalez-Astudillo (2013), Kleim, Kriwoluzky, and Sarferaz (2015, 2016)

served states, and in a novel hidden Markov model of learning. To the best of our knowledge, this is the first paper to study a model with least squares learning agents who jointly estimate the equilibrium law of motion and state probabilities in a self-referential framework. In this environment, it is unclear whether optimal policies are robust to the exclusion of contemporaneous variables from agents' information sets. This is because the non-linear structure of Markov-switching DSGE models prevents agents with lagged information sets from learning the MSV solution.[2] This calls into question the relevant class of equilibria in our policy analysis. Fortunately, in this environment, agents generally succeed in identifying policy states and learning equilibrium coefficients if they receive some contemporaneous signal of current policy. We emphasize that certain implications of the nowcasting problem just described, as well as the broader analysis of stability under learning in hidden Markov models, are left for future work.

Third, we ask: how does optimal monetary policy vary with the average maturity of debt? McClung (2017b) shows that the maturity structure of debt matters for determinacy in models with switching fiscal policy, which is not the case in analogous models with fixed fiscal regimes. We expand on this by showing how the menu of potential optimal policies available to central banks is particularly susceptible to the effects of maturity when fiscal policy is globally switching. Moreover, we show that uncertainty over fiscal policy variables can generate the kind of tradeoff between minimizing loss and maximizing probability of determinacy and E-stability that Evans and McGough (2007) discusses. A companion project to this paper explores whether the monetary authority ought to make balance sheet decisions that target the optimal debt maturity structure.

This paper most directly builds on the works of three optimal policy papers in the New Keynesian literature. First, this paper extends Schmitt-Grohe and Uribe (2007) which studies optimal simple monetary policy rules in fixed regime New Keynesian models with active and passive fiscal policy. We extend this paper by allowing for time-variation in fiscal policy stances. Second, this paper borrows heavily from Orphanides and Williams (2007), which studies the performance and robustness of simple monetary rules in simple monetary models with price-stickiness and learning agents. We build on their work by studying the robustness of optimal rational expectations monetary policy rules to misspecifications about private sector expectations in models with fiscal policy. Finally, we extend Chen et al (2015), which studies joint optimal monetary and fiscal policy in a model with switching policies. Their paper derives fully optimal joint policy rules in a Stackelberg game. In contrast, we look for optimal simple, implementable policy rules that are robust to misspecifications about private sector expectations. We also take the complementary view that fiscal

[2]Here, a MSV solution *only* depends on lagged endogenous variables, and contemporaneous exogenous shocks including the Markov state. We argue that our agents with lagged information sets can only learn equilibria that also depend on past Markov states.

policymakers do not engage in a sophisticated optimization routine to determine fiscal surpluses.

This paper also contributes to a growing literature on regime-switching policy in the context of the Fiscal Theory of the Price Level. For instance, Davig and Leeper (2011), Bianchi (2012, 2013), and Bianchi and Melosi (2014) estimate switching New Keynesian models such as the models in this paper and find evidence of fiscal and monetary policy switches in the U.S. Ascari et al. (2017) and Cho and Moreno (2016) attempt to generalize the determinacy conditions from Leeper (1991) to environments with switching coefficients. We extend their work by showing how maturity impacts determinacy in the presence of fiscal policy switching, and by further studying how the across-regimes behavior of fiscal policymakers constrains the menu of monetary policies consistent with determinacy. Additionally this paper borrows heavily from, and attempts to contribute to, a learning literature involving models with monetary-fiscal policy interactions that includes Eusepi and Preston (2011, 2012, 2013) and Bianchi (2013) and Bianchi and Melosi (2014).

Finally, this paper contributes a hidden Markov model of least squares learning to the learning literature. While many hidden Markov models of learning are introduced by papers such as Bullard and Singh (2012) and Davig (2005), relatively few of them study agents who jointly estimate model parameters and state probabilities. Exceptions that do exist, such as Hansen and Sargent (2010), Hansen, Polson and Sargent (2010), and Johannes et al (2013) invariably involve Bayesian learners, and, almost invariably,[3] Bayesian learners in models without self-referential feedbacks.

Our approach differs from these other approaches in many ways. First, we study conditional least squares learners in a hidden Markov model, whereas other papers study Bayesian learning. Second, we concern ourselves with stability of rational expectations equilibrium in a model with hidden states. Because we can appeal to stochastic approximation papers that study convergence properties of our conditional least squares algorithm, future research may develop stability conditions that help extend the intuition of Evans and Honkapohja (2001) to our regime-switching models. Third, our agents estimate a Markov-switching VAR law of motion for endogenous variables.

Section 2 briefly introduces the model and estimation routine; section 3 derives the optimal interest rate rules under rational expectations in models with short-term debt; section 4 discusses the optimal rules under adaptive learning; section 5 discusses optimal policy in the presence of long-term debt; section 6 concludes.

[3]Hansen and Sargent (2010) is an exception

2. Model and Method

We consider a class of log-linearized New Keynesian models that is augmented to include time-varying fiscal policy as in Davig and Leeper (2011), long-term maturity structure of debt as in Woodford (2001) and Eusepi and Preston (2013). In this class of models, private sector behavior is given by two equations of the form:

$$
\begin{aligned}
\tilde{y}_t &= E_t\tilde{y}_{t+1} - \sigma^{-1}(\hat{i}_t - E_t\hat{\pi}_{t+1}) + \sum_d u_t^d \\
\hat{\pi}_t &= \beta E_t\hat{\pi}_{t+1} + \kappa\tilde{y}_t + \sum_s u_t^s
\end{aligned}
$$

where all variables are expressed as percentage deviations from steady state, $\tilde{y}$ is the output gap, $\hat{\pi}$ is inflation, and $\hat{i}$ is the deviation of nominal interest rates from the nominal interest rate target. $\sum_d u_t^d$ and $\sum_s u_t^s$ are demand and supply shocks, respectively, that may include any number of exogenous processes acting on technology, preferences, market power, etc. To introduce fiscal policy into the model, we consider the log-linearized versions of the following equations:

$$
\begin{aligned}
P_t^m b_t + \tau_t &= \frac{b_{t-1}}{\pi_t}(1 + \rho P_t^m) + G_t \\
P_t^m &= \frac{1}{1+i_t}(1 + \rho E_t P_{t+1}^m)
\end{aligned}
$$

where b is real debt, G is real government spending, and τ is a surplus rule. $\hat{P}_t^m$ is the price of the bond portfolio at time t and $\rho \in [0,1]$ captures the maturity structure of the government debt. While we relegate the derivation of these equations to the appendix, the intuition behind the bond portfolio is fairly simple: the government issues b_t units of a nominal bond portfolio at time t that pays 1 unit of nominal income at time $t+1$, ρ units at time $t+2$, ρ^2 units at $t+3$ and so forth. In other words, government debt exhibits a geometrically decaying maturity structure. This structure allows us to introduce long-term debt into our model by using a single state variable that captures the average maturity of debt, ρ. The limiting cases of ρ illuminate how larger values of ρ correspond to longer average maturities: when $\rho = 0$, all debt is short term, and when $\rho = 1$, all debt is in the form of consols.

In this paper, we are primarily interested in how optimal monetary policy depends on τ, which is characterized by a rule of the form:

$$
\begin{aligned}
\tau_t &= \bar{\tau}\left(b_t(1 + \rho P_t^m)\right)^{\gamma(s_t)} f_t \\
f_t &= f_{t-1}^{\rho_f} e^{\epsilon_t^f}
\end{aligned}
$$

where ϵ_t^f is some mean-zero i.i.d shock. τ adjusts some lump-sum component of the government's structural surplus in response to government liabilities, $b_t(1 + \rho P_t^m)$.

The responsiveness of this fiscal rule is determined by $\gamma(s_t)$, which is assumed to follow a two-state Markov process given by s_t. As we will discuss shortly, the value of γ determines which model variables need to stabilize government debt. Because the parameterization of this rule has general equilibrium implications for inflation and output, we allow the monetary policymaker to employ a log-linearized switching rule of the form:

$$\hat{i}_t = \rho_i \hat{i}_{t-1} + (1-\rho_i)(\phi_\pi(s_t)\hat{\pi}_t + \phi_y(s_t)\tilde{y}_t) + \epsilon_t^R$$

where s_t is the same process that drives variation in γ, $\hat{i}$ is the deviation of the nominal interest rate from its target. To impose structure on G_t and the private sector shocks in the model, we derive a model that is similar in spirit to the simple New Keynesian model in An and Schorfheide (2007). Specifically, government spending is given by

$$\begin{aligned} G_t &= \xi_t Y_t \\ g_t &= \frac{1}{1-\xi_t} \\ ln(g_t) &= \rho_g ln(g_{t-1}) + \epsilon_t^g \end{aligned}$$

According to this specification of government spending, a time varying fraction of output is consumed by the government. If we substitute this into the government budget constraint, it is straightforward to see that government debt depends directly on output. Therefore, government spending in our model introduces an output channel that may have implications for our results. In simple cases where we want to abstract from this output channel, we simply set $\bar{G} = 0$ and $\epsilon_t^g = 0$ and model fiscal disturbances through f_t. To bring demand and supply shocks into the model, we assume that there are both markup shocks and shocks to household preferences. The model derivation is left for the appendix, but we can write our model in the form

$$\begin{aligned} \tilde{y}_t &= E_t\tilde{y}_{t+1} - \sigma^{-1}(i_t - E_t\hat{\pi}_{t+1}) + \sigma^{-1}\rho_{\hat{z}}\hat{z}_t \\ \hat{\pi}_t &= \beta E_t\hat{\pi}_{t+1} + \kappa\tilde{y}_t + \hat{\mu}_t \\ \hat{i}_t &= \rho_i\hat{i}_{t-1} + (1-\rho_i)(\phi_\pi(s_t)\hat{\pi}_t + \phi_y(s_t)\hat{y}_t) + \epsilon_t^R \\ \hat{b}_t &= \beta^{-1}(\hat{b}_{t-1} - \pi_t) - (1-\rho)\hat{P}_{m,t} + \beta^{-1}\frac{\bar{G}}{\bar{b}}\hat{y}_t \\ &\quad -\beta^{-1}((1-\beta) + (1-\beta\rho)\frac{\bar{G}}{\bar{B}})\hat{\tau}_t + \beta^{-1}\hat{g}_t \\ \hat{\tau}_t &= \gamma(s_t)(\hat{b}_{t-1} + \beta\rho\hat{P}_{m,t}) + \hat{f}_t \\ \hat{P}_{m,t} &= -\hat{i}_t + \beta\rho E_t\hat{P}_{m,t+1} \\ \hat{f}_t &= \rho_{\hat{f}}\hat{f}_{t-1} + \epsilon_t^f \\ \hat{g}_t &= \rho_g\hat{g}_{t-1} + \epsilon_t^g \\ \hat{z}_t &= \rho_{\hat{z}}\hat{z}_{t-1} + \epsilon_t^z \\ \hat{\mu}_t &= \rho_\mu\hat{\mu}_{t-1} + \epsilon_t^\mu \end{aligned}$$

where z is the technology shock, μ is the cost-push shock, P is the transition probability matrix and $p_{ij} = Pr(s_t = j|s_{t-1} = i)$. All other variables are defined as before. In the baseline analysis, we calibrate our model so that the steady state government liabilities, $\bar{b}(1 + \rho\bar{P})$, equals the steady state level of output. We also set $\bar{G}/\bar{b}$ so that $\bar{G}/\bar{Y} = .2$ conditional of ρ. My main results do not seem to depend on these assumptions, except in rare special cases we discuss in section 3. Having written the model, we are now in a position to define the fiscal policy stance on debt, and the monetary policy rule. To that end, we use the following two definitions.

Definition 1 *A fiscal policy is defined by the following parameters:* $\{p_{11}, p_{22}, \gamma(1), \gamma(2)\}$*. Stated more thoroughly, a switching fiscal policy is fully characterized by within-regime responses to outstanding debt, given by* $\{\gamma(1), \gamma(2)\}$ *and by the transition probabilities* $\{p_{11}, p_{22}\}$*.*

Definition 2 *A monetary policy is defined by the parameters of the interest rate rule:* $\{\phi_\pi(1), \phi_\pi(2), \phi_y(1), \phi_y(2), \rho_i\}$*. In section 5, we may allow a monetary policy to be indexed by the average maturity of debt,* ρ*.*

We subject our policy rule to a monetary policy shock to help account for fluctuations of i around its target value, or to capture any short-lived deviation of policy from the rule that might be caused by dissension between policymakers. In our simple model without debt, ϵ^R is isomorphic to a demand shock in the IS curve. As such, we do not need monetary policy to explore optimal monetary responses to demand shocks in a model with Ricardian dynamics. In our model, however, ϵ^R shows up in both the IS curve and the government budget constraint and this will have implications for optimal policy.

To help distinguish between policy regimes, we follow Leeper (1991) and describe an "active" policymaker as one who determines inflation without concern for the stability of debt, and a "passive" policymaker as one who directly acts to stabilize the evolution of debt. With respect to fiscal policy, $\gamma > 1$ characterizes a "passive" policy regime. When $\gamma > 1$, bonds evolve according to a stable autoregressive process so that changes in i and π are not needed to keep debt from exploding. Intuitively, $\gamma > 1$ means that surpluses adjust endogenously by an amount that is sufficient to pay down interest and principal on new debt issuance over an infinite horizon. In such an environment, forward looking agents recognize that any wealth effects stemming from debt issuance will be offset by future taxes and this renders policy Ricardian. Because fiscal policy stabilizes debt, the central bank is free to contain inflation as it pleases – ideally by employing an active monetary policy that satisfies the Taylor Principle.

When fiscal policy is active (i.e. $\gamma < 1$), surpluses do not rise by enough to offset any wealth effects coming from any new debt issuance and this causes consumption and inflation to rise in response to higher debt. Any rise in inflation that results from these wealth effects must be accommodated by central banks; if central banks raise interest rates by more than one-for-one in response to higher inflation, they will raise real debt service costs, leading to higher debt and therefore higher inflation in the future, and so on. Central banks therefore must respond weakly or passively to inflation so that inflation may erode the outstanding debt stock without generating additional debt service costs. Such monetary policy is said to be "passive" and is characterized by a violation of the Taylor Principle (e.g. $\phi_\pi < 1$).

Parameter values are chosen to so that the resulting model is determinate. While there are no simple analytical conditions for determinacy in our switching model, Woodford (1998a) gives simple conditions for determinacy in the case of non-switching (assuming $\bar{G} = 0$):[4]

	$\phi_\pi > 1 - \frac{1-\beta}{\kappa}\phi_y$	$\phi_\pi < 1 - \frac{1-\beta}{\kappa}\phi_y$
$\gamma \in (1, \frac{\beta^{-1}+1}{\beta^{-1}-1})$	determinate	indeterminate
$\gamma \notin (1, \frac{\beta^{-1}+1}{\beta^{-1}-1})$	no stable solution	determinate

We say that the economy is in Regime M when $\phi_\pi > 1 - \frac{1-\beta}{\kappa}\phi_y$ and $\gamma \in (1, \frac{\beta^{-1}+1}{\beta^{-1}-1})$. and that the economy is in Regime F when $\phi_\pi < 1 - \frac{1-\beta}{\kappa}\phi_y$ and $\gamma \notin (1, \frac{\beta^{-1}+1}{\beta^{-1}-1})$. In Regime M, fiscal policy is passive while monetary policy is active. This is the standard assumption in most New Keynesian research. In Regime F, fiscal policy is active while monetary policy is passive. Our model features switching between Regime F and Regime M policy configurations. That is, our model features 2 states (i.e. $S = 2$) where each state is consistent with determinacy in the analogous fixed regime model.[5] We solve the model and check for determinacy in the mean-square-stable sense using techniques from Cho (2016).

We now specify the optimization problem. The central bank chooses $\phi = (\phi_\pi(1), \phi_y(1), \phi_\pi(2), \phi_y(2), \rho_i, \rho)\ \forall s_t$ to minimize:

$$l(\phi) = var(\pi) + \lambda var(y)$$

The choice of ϕ that minimizes l is referred to as the optimal policy or "optimized" simple interest rate rule as in Orphanides and Williams (2007). The optimal policy is said to be time-invariant if $\phi_\pi(1) = \phi_\pi(2)$ and $\phi_y(1) = \phi_y(2)$, and is said to be time-varying otherwise. Though the average maturity of debt appears to be a fiscal policy tool, central banks are able to engage in large-scale asset purchase (LSAP) programs that "twist" the maturity structure of the debt held by the public. Consequently,

[4]We assume that $\phi_\pi(s_t) \geq 0$ for all s_t

[5]Despite the fact that each regime induces determinacy in a fixed regime model, the model with switching between determinate regimes is often explosive or yields indeterminacy

we include this parameter in the central's bank's choice set in specific exercises. In addition to minimizing loss, the optimal policy should satisfy two criteria: (1) the optimal policy should implement a unique mean-square stable rational expectations equilibrium; (2) optimal inflation reaction coefficients must be non-negative.

As implied by our discussion of the Leeper (1991) conditions, the value of the fiscal policy parameter, γ, impacts the menu of policy options that central banks must choose from to contain inflation. When $\gamma < 1$, fiscal policy is active and monetary responses to inflation must be dovish; when γ is high and policy is passive, monetary responses must be aggressive. To help characterize how fiscal policy constrains central bankers in our model with time-varying policy stances, we employ a generalization of these conditions similar in spirit to conditions developed by Ascari et al (2017). Our taxonomy considers three types of fiscal policy stances: (1) "globally passive" policies that support a stable Ricardian equivalent equilibrium; (2) "globally active" policies that are more active than passive across regimes; (3) "globally switching" or "balanced" policies that are neither more active nor more passive across regimes. We note that both globally active and globally switching policies feature non-Ricardian dynamics; only globally passive policies are Ricardian. The following definitions help us characterize our three categories of switching fiscal policy and provide valuable intuition.

Definition 3 *A fiscal policy is* ***globally passive*** *if* $\phi_\pi(1) = \phi_\pi(2) = \alpha^P$ *for all* $\alpha^P > 1$ *yields a determinate equilibrium.*

A globally passive policy can be paired with any time-invariant interest rate rule that satisfies the Taylor Principle. In order for this to be true, fiscal policy must be Ricardian. Otherwise, we could choose a time-invariant active monetary policy that places debt on an explosive path. Because globally passive implies Ricardian equivalence and vice versa we can determine if a policy is globally passive using the following conditions passive (assuming $p_{11} + p_{22} > 1$):[6]

$$(p_{11} + p_{22} - 1)h_1^2 h_2^2 \quad < \quad 1 \tag{1}$$
$$p_{11}h_1^2(1 - h_2^2) + p_{22}h_2^2(1 - h_1^2) + h_1^2 h_2^2 \quad < \quad 1 \tag{2}$$

where $h_i = \beta^{-1}(1-(1-\beta)\gamma(i))$ for $i = 1, 2$. These conditions, which Ascari et al (2017) present, tell us when the budget constraint implies a mean-square stable autoregressive process for debt. If a fiscal policy satisfies these conditions, then debt evolves according to a mean-square stable autoregressive process without accommodation from the monetary authority and this allows monetary policymakers to determine inflation and output in the non-policy block of the New Keynesian model. Determinacy then requires that interest rates respond aggressively to inflation. Figure 1 shows regions of determinacy, indeterminacy and non-existence of stable solutions for

[6] We are interested in highly persistent regimes, which makes this a harmless assumption

a model with globally passive policy. As argued in Ascari et al. (2017), the determinacy region in Figure 1 presents something akin to the Long-Run Taylor Principle in Davig and Leeper (2007): fiscal policy can be very active for short amount of times, or modestly active with persistence, and the resulting equilibrium may still be Ricardian and determinate if policy is mostly passive overall.[7] Note that a fixed passive fiscal policy regime is merely a special case of a globally passive policy.

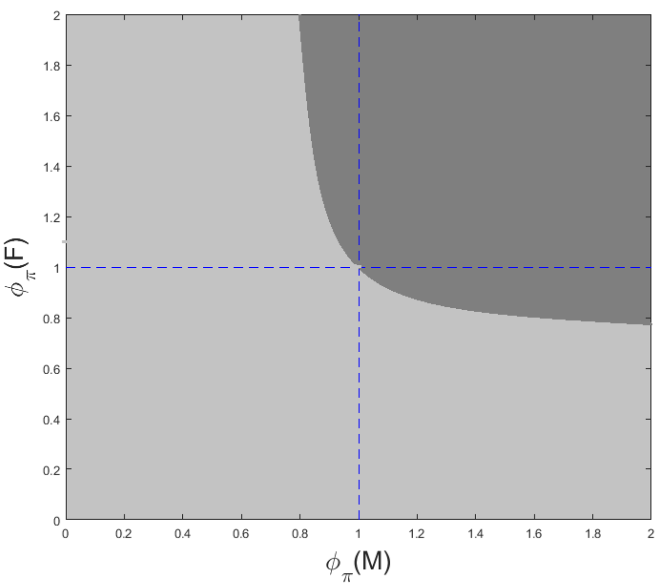

Figure 1: Globally Passive Policy: the determinacy region is dark gray; indeterminacy is light gray; no stable solutions is white

Definition 4 *A fiscal policy is* ***globally active*** *if $\phi_\pi(1) = \phi_\pi(2) = \alpha^A$ for all $\alpha < 1$ yields a determinate equilibrium.*

Stated equivalently, policy is globally active if a unique equilibrium exists when the monetary authority employs a time-invariant passive monetary policy. For a permanent passive monetary policy to be consistent with determinacy in our model, fiscal policy must be more active than passive overall. Figure 2 shows regions of determinacy, indeterminacy and explosiveness region for active fiscal policy. Note that a fixed active fiscal policy is merely a special case of a globally active policy.

Definition 5 *A fiscal policy is* ***globally switching*** *if there exists $\alpha^A < 1$ and $\alpha^P > 1$ such that neither $\phi_\pi(1) = \phi_\pi(2) = \alpha^A$ nor $\phi_\pi(1) = \phi_\pi(2) = \alpha^P$ yield a determinate equilibrium.*

[7] In Figures 1-2, parameters are chosen to emphasize changes in determinacy regions that result from time-varying policy. Other parameterizations may yield smaller, qualitatively similar determinacy regions

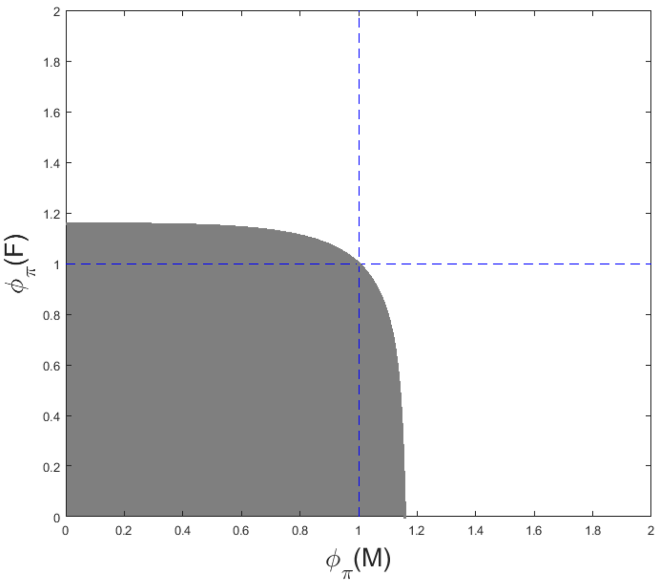

Figure 2: Globally Active Policy: the determinacy region is dark gray; indeterminacy is light gray; no stable solutions is white

The set of globally switching fiscal policies is the complement of the set of globally active and passive policies. Intuitively, a globally switching policy is neither active enough in the long-run to support all passive monetary policies nor passive enough in the long-run to support all active monetary policies. These fiscal policies are balanced in the sense that they are not obviously more active or passive overall. For example, a globally switching policy may feature slow-changing, strongly active and strongly passive regimes, or fast-changing weakly active and weakly passive fiscal policy regimes. Table 1 offers very rough qualitative examples of how fiscal policies may be assigned to certain categories. Figure 3 shows determinacy regions for policies that feature highly persistent and/or strongly active and passive regimes. This figure suggests that determinacy requires monetary authorities to be hawkish during passive fiscal regimes and dovish during active fiscal regimes. Crucially, central banks cannot implement time-invariant policies such as permanent interest rate pegs because the overall fiscal policy stance is no longer mostly active or mostly passive. Figure 4 shows determinacy regions for policies that feature fast-changing and/or weakly active and passive regimes. In these scenarios, central bankers face a meager menu of policy options. Typically, determinacy regions for globally switching policies will resemble either Figure 3 or 4 depending on the strength of switching regime fiscal policy responses to debt and the persistence of regimes. Table 1 offers very rough qualitative examples of how fiscal policies may be assigned to certain categories.

Before we present results we offer some final intuition about the fiscal policy tax-

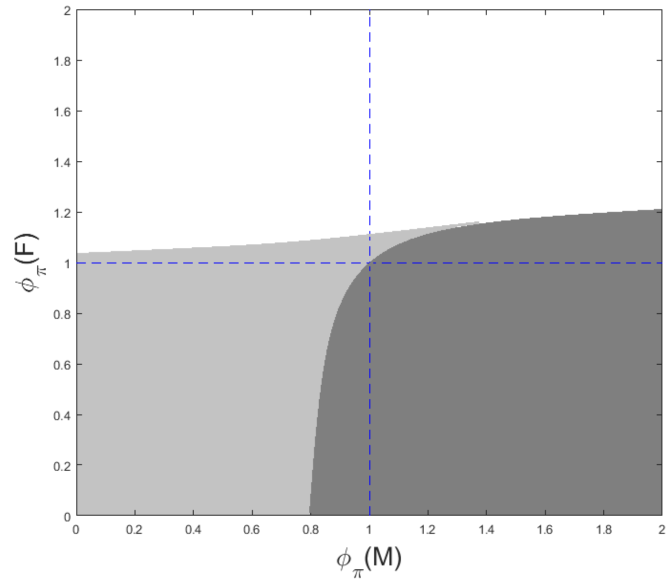

Figure 3: Globally Switching Policy: The determinacy region is dark gray; indeterminacy is light gray; no stable solutions is white

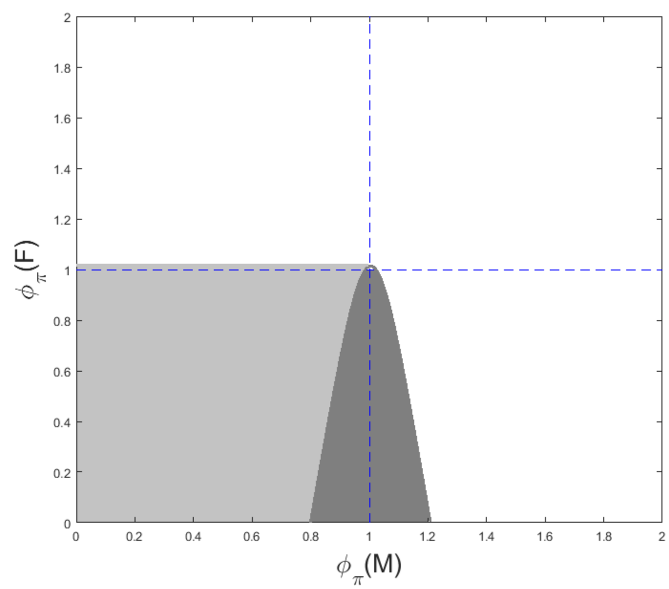

Figure 4: Globally Switching Policy: the determinacy region is dark gray; indeterminacy is light gray; no stable solutions is white

	low pers.; weak active	high pers.; weak active	low pers.; strong active	high pers.; strong active
low pers.; weak passive	**GS**	**GA**	**GA**	**GA**
high pers.; weak passive	**GP**	**GS**	**GS**	**GA**
low pers.; strong passive	**GP**	**GS**	**GS**	**GA**
high pers.; strong passive	**GP**	**GP**	**GP**	**GS**

Table 1: Strength and persistence of fiscal regime and the overall stance of fiscal policy. "pers." = persistence; GA = "globally active"; GP = "globally passive"; GS = "globally switching"

onomy. Globally active and globally passive policies can be coupled with a wide range of time-invariant policies to deliver a determinate model, while globally switching policies must be paired with time-varying monetary policies for determinacy. One practical benefit of using time-invariant policies is that their implementation does not require policymakers to actively track any changes in responsiveness of fiscal policy to debt. As we show next, time-invariant policies are going to perform well in models with globally active or passive policy.

3. Short-term Debt and Rational Expectations

In this section, I abstract from long-term debt by assuming $\rho = 0$. We also assume that $\lambda = 0$ so that the central bank loss function equals that variance of inflation, but the optimal policies discussed in this section appear to perform well in applications with small λ such as the λ weights typically found in microfounded loss functions. Because we prioritize inflation-targeting, $\phi_y(1) = \phi_y(2) = 0$ reduces loss in our numerical search. Accordingly, we restrict our attention the interest rate rules of the form:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)\phi_\pi(s_t)\pi_t$$

We present our numerical results through a series of claims contained in this section.

Claim 1 *If fiscal policy is globally passive then for all parameterizations the optimal monetary policy response is to employ an interest rate rule of the form*

$$\begin{aligned} \phi_\pi(s_t) &= \bar{\phi_\pi} \quad \forall s_t \\ \rho_i &= \bar{\rho}_i \geq 0 \end{aligned}$$

Since determinacy requires that inflation and output be determined in the non-policy block, the optimal simple policy rule is identical to the optimal rule used in small-scale 3-equation models that consist of an IS curve, Phillips Curve and interest rate rule (see Woodford (2003)). Intuitively, a globally passive policy supports a mean-square stable autoregressive process for debt. Consequently, central banks do not need to accommodate fiscal policy and this allows monetary policymakers to determine inflation through the non-policy block.

When $\lambda = 0$, $\bar{\phi_\pi} \to \infty$ (see Woodford (2003)). Two features of this result should be emphasized. First, the optimal policy is time-invariant despite switching in the fiscal policy stance. Second, the monetarist equilibrium can be stable in models with persistent active fiscal policy. For example, the monetarist equilibrium is stable when $p_{11} = p_{22} = .95$, $\gamma(1) = 5$, $\gamma(2) = 0$, and $\bar{G} = 0$ despite the fact that fiscal surpluses are entirely exogenous half of the time. See Orphanides and Williams (2007) for a treatment of the optimal ρ_i when interest rates are determined in this environment.

Claim 2 *If fiscal policy is globally active then for all reasonable parameterizations the monetary authorities should employ a permanent interest rate peg (i.e.* $\phi_\pi(1) = \phi_\pi(2) = 0$*) in order to minimize the variance of inflation.*

While we cannot prove Claim 3 formally, our claim relies on the following numerical support: for all globally active policies in $p_{11} \in [.9, 1]$, $p_{22} \in [.9, 1]$, $\gamma(1) \in [-10, 10]$, and $\gamma(2) \in [-10, 10]$, the interest rate peg is optimal. For the posterior mean calibration with added cost-push shocks and fiscal variables, we search over approximately 64,000 globally active policies and find that the interest rate peg is optimal for each one of them. We repeat this analysis for alternative reasonable calibrations (alternative shock covariance-variance structure, alternative persistence parameters for structural shocks, σ, κ) and find that this result is robust.

We add the word "reasonable" because non-Ricardian fiscal policy presents a tradeoff between stabilizing inflation in response to private sector shocks (i.e. demand and supply shocks), and stabilizing inflation in response to policy shocks. In conjunction with active fiscal policy, private sector shocks call for very high ρ_i (e.g. $\rho_i = .995$) and time-varying inflation reaction coefficients, while pegs perform best in response to monetary and fiscal policy shocks. As a result, the optimal monetary policy in a model with globally active fiscal policy depends on the net effect of these competing influences on inflation.

As it turns out, the private sector shock variances need to be very large relative to policy shock variances, or the private sector shocks need to be very persistent relative to policy shocks for the interest rate peg to be suboptimal. In particular, monetary policy shocks need to be very small relative to other shock variances. To illustrate this last point, we calibrate the model at the posterior mean, shut down each shock except for one private sector shock and ask: how large does the variance of the monetary policy shock need to be for the interest rate peg to be optimal?

When we set $\bar{G} = 0$, so that output no longer impacts debt through the budget constraint and set $\rho_u = .99$ where ρ_u is the supply shock persistence term, we need for the variance of the monetary policy shock to be greater than .014% of the variance of the i.i.d innovation to the supply shock to get an optimal interest rate peg. For $\rho_u = .9$, the monetary policy shock needs to be greater than .0017% of the variance of the same innovation to the supply shock. When we increase $\bar{G}$ to .2, the interest rate peg is optimal even when cost-push is the only shock in the model.

When we set $\bar{G} = 0$, so that output no longer impacts debt through the budget constraint and set $\rho_z = .99$ where ρ_z is the demand shock persistence term, the variance of the i.i.d. innovation to the demand shock must be less than 5 times the variance of the monetary policy shock for the peg to be optimal. This suggests that demand shocks are a bigger threat to the optimal interest rate peg. However, when $\rho_z = .9$, the monetary policy shock only needs to be greater than .025% of the variance of the i.i.d innovation to the demand shock for the peg to be optimal. Of course, these exercises exclude fiscal policy shocks, and those shocks help to select the peg. For example, if we shut down monetary policy shocks and set all remaining shock parameters to their posterior mean values, we can set $\rho_z = .99$ and still have an optimal peg.

Because intuition supports the inclusion of policy shocks in our model, and because it is highly unlikely that an estimated model will reject the inclusion of policy shocks, we regard cases where the peg is suboptimal as special cases involving potentially unreasonable parameterizations of the model. We also note that pegs are quite often *nearly* optimal in that loss is often close to 0% higher under the peg when compared to the optimum. However, we have found cases where loss is as much as 3% higher under the peg.

To understand why pegs perform so well in globally active models, it's important to recall the fact that debt both determines inflation and is stabilized by inflation in any non-Ricardian equilibrium. This means that any shock to government debt (i.e. any shock appearing in the budget constraint) will have an affect on inflation and output. To fix things, consider a shock which raises debt. Since agents perceive government debt as net wealth, this will raise consumption and inflation. This is one sense in which debt determines inflation under a globally active policy. The amount of inflation generated in general equilibrium depends on monetary policy, however. As such, monetary policy determines how inflation feeds back to stabilize debt. If an interest rate peg is in place, a large inflation will occur today, which pushes debt in the direction of its steady state value. On the other hand, if the central bank allows interest rates to respond positively, then debt service costs will increase today, which creates higher debt tomorrow and so on. The higher expected path of debt raises time t inflation expectations, so that inflation is both higher today and propagated into the future. In a similar thought experiment, Leeper and Leith (2016) show that the present value of inflation will be higher under the responsive interest rate than under the peg in their small-scale New Keynesian model. Once they solve for the equilibrium path of inflation, it's straightforward to show that the sharp, sudden responses of inflation under the peg are consistent with less volatility in inflation. The complexity of our non-linear model makes it very difficult to repeat a similar experiment in this paper. However, Claim 3 strongly suggests that their results generalize to models with time-varying fiscal stances – even models with recurring passive fiscal policy regimes.

While pegs are broadly consistent with stable inflation in our non-Ricardian model, $\phi_\pi(1) = \phi_\pi(2) = 0$ does not guarantee determinacy for all fiscal policies that violate the abovementioned conditions (i.e. mean-square stable common-factor sunspots may exist). In particular, indeterminacy obtains if fiscal policy is too passive in one regime. For example, if $p_{11} = p_{22} = .95$, $\gamma(1) = 2$, $\gamma(2) = -5$ then fiscal policy is sufficiently active for the interest rate peg to deliver determinacy. If, however, $\gamma(1) = 2$ is replaced by $\gamma(1) = 5$, then policy is too passive in regime 1 for the interest rate peg to deliver determinacy. These are globally switching equilibria and determinacy requires that they be paired with special optimal equilibria.

Claim 3 *Optimal globally switching monetary policies are time-varying and parameter dependent*

For example, the optimized inflation reaction coefficients for the policy given by $p_{11} = p_{22} = .95$, $\gamma(1) = 5$, $\gamma(2) = -5$ and for the policy given by $p_{11} = p_{22} = .95$, $\gamma(1) = 2$, $\gamma(2) = 0$, are $(\phi_\pi(1), \phi_\pi(2), \rho_i) = (3.3, 0, .99)$ and $(\phi_\pi(1), \phi_\pi(2), \rho_i) = (2.97, .73, .99)$, respectively. These particular optimized coefficients come from the expected posterior loss exercise we introduce in the next paragraph. Since the optimized policy favors large swings in inflation responses, policy inertia is undesired (i.e. $\rho_i = 0$ is frequently optimal).

While results in the globally active and globally passive settings hinge only on fiscal policy parameters (for reasonable parameterizations of shock processes), the optimal policy in globally switching models depends on any model parameter that impacts determinacy conditions. This means that we need to choose parameter values in order to draw conclusions about optimal policy in the globally switching models. To help inform our selection of model parameter values in a manner that mitigates problems associated with parameter uncertainty, we estimate the following truncated model using Bayesian techniques:[8]

$$\begin{aligned}
\hat{y}_t &= E_t\hat{y}_{t+1} - \sigma^{-1}(i_t - E_t\hat{\pi}_{t+1}) + (1-\rho_g)\hat{g}_t + \sigma^{-1}\rho_{\hat{z}}\hat{z}_t \\
\hat{\pi}_t &= \beta E_t\hat{\pi}_{t+1} + \kappa(\hat{y}_t - \hat{g}_t) + \hat{\mu}_t \\
\hat{i}_t &= \rho_i\hat{i}_{t-1} + (1-\rho_i)(\phi_\pi\hat{\pi}_t + \phi_y\hat{y}_t) + \epsilon_t^R \\
\hat{g}_t &= \rho_g\hat{g}_{t-1} + \epsilon_t^g \\
\hat{z}_t &= \rho_{\hat{z}}\hat{z}_{t-1} + \epsilon_t^z \\
\hat{\mu}_t &= \rho_\mu\hat{\mu}_{t-1} + \epsilon_t^\mu \\
\tilde{y}_t &= \hat{y}_t - \hat{g}_t
\end{aligned}$$

where the final equations reflects the fact the natural rate of output equals $\hat{g}_t$ in our model, thus allowing us to glean information about government shocks from estimates of the IS and Phillips Curves. Using the separated partial means test to test the convergence of our estimates, we believe that the best results emerge when we place dogmatic priors over ρ_u and σ_μ and estimate only the non-policy block with the interest rate. Since this exercise intends to consider counterfactural policies, estimates of the underlying fiscal policy stance are unnecessary. However, estimates pertaining to the shock processes and other private sector coefficients help to discipline our analysis towards parameter regions that agree better with the data. In extensions of the present work, we intend to estimate a fuller DSGE model.

After sampling from the posterior distribution, we follow Cogley et al. (2011) and compute the expected posterior loss associated with each policy parameterization. A Monte Carlo average of the following expected posterior loss function is computed:

$$\int l(\phi)P(\tilde{\theta}|Y)d\tilde{\theta}$$

where $\tilde{\theta} = (\kappa, \sigma, \rho_g, \rho_z, \rho_\mu, \sigma_g, \sigma_z, \sigma_\mu, \sigma_r)$. For the previously mentioned case where $\gamma(1) = 5$ and $\gamma(2) = -5$, $p_{11} = p_{22} = .95$, the optimal policy is given by $\phi_\pi(1) = 3.33$, $\phi_\pi(2) = 0$, $\rho_i = 0.99$. Intuitively, monetary policy should be active in the passive fiscal regime, and very passive in the active fiscal regime.

In Figure 5, we calibrate non-policy parameters at their estimated posterior mean values to show regions of the fiscal policy parameter space that correspond to globally active, globally switching and globally passive fiscal policy. While these results are obtained numerically, we can analytically distinguish between our three fiscal policy categories in a simpler model with flexible prices. Estimates of $\gamma(M)$ and $\gamma(F)$ taken from some of the abovementioned papers such as Davig and Leeper (2006, 2011) suggest that pre-Great Recession U.S. fiscal policy lies somewhere near the dividing line between Globally Active and Globally Switching fiscal policies.

[8]see the Appendix for tables containing information about our prior and posterior distributions

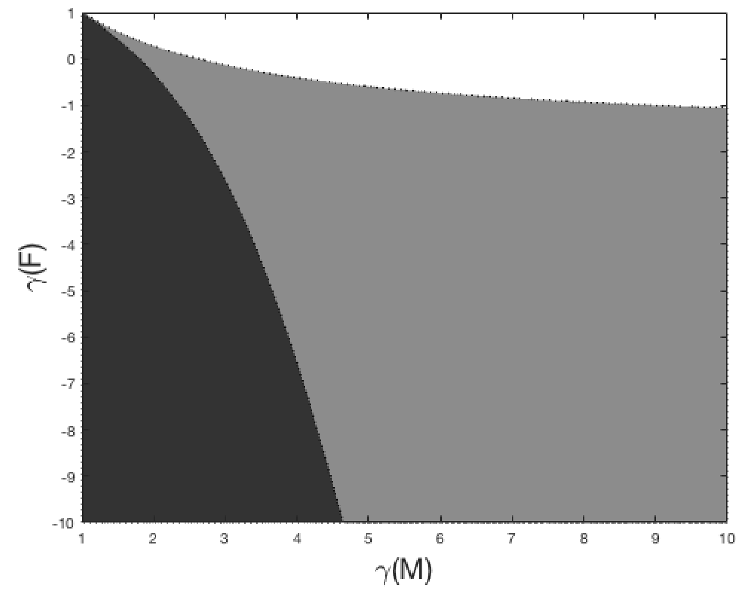

Figure 5: **Fiscal Policy Taxonomy and Fiscal Policy Parameter Space:** dark gray corresponds to Globally Active; light gray corresponds to Globally Switching; white corresponds to Globally Passive

To sum up, the optimal policy response depends on whether fiscal policy is globally active, globally passive or globally switching. If fiscal policy is globally passive, then optimal policy is time-invariant and calls for large inflation reaction coefficients. If fiscal policy is not globally passive, then interest rate pegs deliver the fundamental solutions that minimize loss. If, however, fiscal policy is globally switching then interest rate pegs lead to indeterminacy. In those settings, and those settings alone, the optimal policy is time-varying.

Type	$(\gamma(1),\ \gamma(2))$	**optimal** $(\phi_\pi(1),\ \phi_\pi(2))$
GP	$(5,\ 0)$	$(\infty,\ \infty)$
GS	$(5,\ -5)$	$(3.33,\ 0)$
GS	$(2,\ 0)$	$(2.73,\ .72)$
GA	$(2,\ -5)$	$(0,\ 0)$

Table 2: Optimal inflation reaction coefficients under rational expectations when $\rho = 0$

4. Adaptive Learning

In this section we relax the assumption that agents form rational expectations, and study policy performance in a model where agents attempt to learn the equilibrium law of motion for the model's endogenous variables, and form forecasts of future variables according to an estimated perceived law of motion. Relative to rational expectations models, models with learning agents feature instabilities that arise from agents' forecast errors. Specifically, agents' forecast errors affect the model's data generating process, thereby changing future data points and future estimates of the model's coefficients. This self-referential feature of our model fundamentally changes the way in which policy interacts with expectations to contain inflation and output. As such, the inclusion of adaptive learning in our analysis provides an important robustness check. Our main conclusion is that the optimized simple policy rules studied in section 3 are robust to misspecifications of the underlying model of expectations employed by agents. That is, the optimized policy rules under rational expectations are optimal or nearly optimal in models with adaptive learning agents, with exceptions in the case of globally switching policy.

We present our results in two subsections. 4.1. studies a learning model in which agents observe the model's endogenous variables (with a reasonable lag), exogenous driving processes and the underlying Markov state that drives variation in fiscal and monetary policy rules. When agents observe the underlying Markov state, they can easily update parameter estimates using a *within-state* recursive least squares algorithm that resembles the least squares algorithm developed and discussed in Evans and Honkapohja (2001). While this learning specification provides a natural first step away from the rather strong assumption that agents form rational expectations, it still assumes that agents easily observe something an applied econometrician would not: the underlying state of policy. Section 4.2 therefore backs away from this assumption.

In 4.2, agents estimate the same perceived law of motion used in 4.1, but do not observe the underlying Markov state (i.e. agents find themselves in a hidden Markov model). Because agents do not observe the stance of fiscal and monetary policy, they cannot use the recursive least squares algorithm employed in 4.1. Instead, we allow agents to use the recursive MLE algorithm and the recursive conditional least squares algorithm developed in Krishnamurthy and Yin (2002) and LeGland and Mevel (1997) to update parameter estimates after observing the model's endogenous variables and exogenous driving processes. We emphasize two main results from this section. First, this is, to the best of our knowledge, the first paper to study least squares learning agents who estimate a Markov-switching autoregressive equilibrium law of motion in a self-referential model with hidden Markov states.[9] That is, previous research does not jointly estimate perceived laws of motion and the Markov state probabilities. We therefore regard this section as a springboard for future research on the use of hidden

[9] see literature review for details on related papers

markov models of learning. Second, the exogeneity[10] of policy rule coefficients makes it possible for agents to infer the underlying state with some reasonable accuracy. Hence, model dynamics in section 4.2 are very similar to model dynamics observed in 4.1. We conclude that it is potentially reasonable to assume that agents observe policy switches, as in section 4.1, but it remains to be seen whether with assumption is strong in models where Markov-switching affects non-policy variables such as trend growth.

4.1 Observed Markov States

We now develop a model of learning in which agents observe the underlying Markov state (i.e. they observe the underlying policy stance). The model dynamics are still given by an actual law of motion, which can be constructed from the log-linearized equilibrium conditions in section 2:

$$x_t = A(s_t)E_t x_{t+1} + B(s_t)x_{t-1} + C(s_t)z_t \tag{3}$$

where $x = (\pi \ y \ i \ b \ \tau \ P)'$ and $z = (g \ \hat{z} \ \epsilon^R \ \mu \ f)$. Under rational expectations, agents know the full structure given by (3) and can solve for the rational expectations equilibrium. Under adaptive learning, however, agents do not know (3) and are therefore incapable of computing the true mathematical expectations of tomorrow's variables. Despite the fact that agents are not fully rational, we still endow agents with sophisticated beliefs about the law of motion governing inflation, output, etc., in equilibrium. Specifically, we give agents the following perceived law of motion (PLM):

$$x_t = a(s_t) + b(s_t)x_{t-1} + c(s_t)z_t \tag{4}$$

Notice that this perceived law of motion has the same functional form as the MSV law of motion, which implies that agents may conceivably learn the rational expectations equilibrium law of motion if their estimates of $a(s_t)$, $b(s_t)$, and $c(s_t)$ converge to their rational expectations values (i.e. if $a(s_t) \to 0_{n\times 1}$, $b(s_t) \to \Omega(s_t)$ and $c(s_t) \to \Gamma(s_t)$ for all s_t). If a rational expectations equilibrium can be learned, it is said to be "stable under learning" or "expectationally stable" ("E-stable") (see McClung (2016, 2017b) for more about E-stability in this class of models). E-stable rational expectations equilibria are easier to rationalize in the sense that adaptive learning supports a coordination story for their realization, and in the sense that they are robust to the strong assumptions that undergird rational expectations. Our task in this section is to study the volatility of inflation and output when agents beliefs about the structure of the economy are close to the unique rational expectations equilibrium implemented by the monetary policy rule.

[10]By "exogeneity" we mean that the coefficients in the equilibrium law of motion for policy variables τ and i do not depend on agents' beliefs

To make our model of learning fully operational, we must specify agents' information set, their estimation strategy, and the full process through which expectations interact with predetermined variables to pin down the endogenous variable values. We begin by specifying agents' time t information set, I_t, which includes all past observations of x, and all past and current observations of z and s. Formally: $I_t = \{y_{t-1}, y_{t-2}, \ldots, y_0; z_t, z_{t-1}, \ldots, z_0; s_t, s_{t-1}, \ldots, s_0\}$. We could exclude z_t from the information set (i.e. only include past values of z) and obtain similar results. Using observations in I_t, agents will update their estimates of the coefficients in (4) using the following *within-regime* learning algorithm:

$$\Phi(s_t)_{st} = \Phi(s_t)_{st-1} + \psi_{s_t} R(s_t)_{st}^{-1} u_t (x_t - \Phi(s_t)'_{st-1} u_t) \tag{5}$$

$$R(s_t)_{st} = R(s_t)_{st-1} + \psi_{s_t} (u_t u_t' - R(s_t)_{st-1}) \tag{6}$$

where $\Phi(s_t)_{st} = (a_t(s_t), b_t(s_t), c_t(s_t))'$ are the time-t estimates of regime s_t coefficients, $u_t = (1, x'_{t-1}, z'_t)'$, and st is the number of realizations of state s_t up until and including time t. Alternatively, we might use a learning algorithm that estimates a dummy variable regression where elements in u are interacted with dummy variables that take on values of 1 or 0 depending on the underlying Markov state. The last feature of the algorithm we need to define is the gain parameter, ψ_{s_t}. Intuitively, ψ_{s_t} attaches a weight to each new observation and therefore determines the extent to which new information impacts parameter estimates. If we give each observation equal weight by setting $\psi = 1/t_{s_t}$, where t_{s_t} is the number of realizations of s_t up until time t, then our learning algorithm becomes the conditional recursive least squares estimator of Φ. Clearly, as $t \to \infty$ the estimates converge to some value, which may be the rational expectations equilibrium coefficients depending on initial beliefs and the E-stability of the equilibrium under study. Alternatively, we might allow agents to give more weight to recent observations by using a constant gain parameter, $\psi = \bar{\psi}$, where $\bar{\psi}$ is some scalar. In constant gain learning algorithms, beliefs will never converge, but may converge to some distribution centered on the rational expectations equilibrium. These algorithms are considered appropriate in settings where agents may expected structural changes in the model, or in settings where agents simply value recent data more than older data.

Having specified the learning algorithm, we now outline the sequence of events that lead to an equilibrium at time t:

1. Agents observe z_t and s_t and add those to their information sets.

2. Using I_t and time $t-1$ estimates $a_{t-1}(s_t)$, $b_{t-1}(s_t)$, $c_{t-1}(s_t)$. Agents form fore-

casts, $\hat{E}_t x_{t+1}$:

$$\begin{aligned}\hat{E}_t x_{t+1} &= (p_{s_t 1} a(1)_{t-1} + p_{s_t 2} a(2)_{t-1}) + \\ &\quad (p_{s_t 1} b(1)_{t-1} + p_{s_t 2} b(2)_{t-1})(a(s_t)_{t-1} + b(s_t)_{t-1} x_{t-1} + c(s_t)_{t-1} z_t) + \\ &\quad (p_{s_t 1} c(1)_{t-1} + p_{s_t 2} c(2)_{t-1}) \rho z_t\end{aligned}$$

3. x_t is generated from the actual law of motion, (3), which gives us time t endogenous variables as a function of beliefs and predetermined variables

4. Agents observe x_t and add it to their information sets

5. Agents use (5)-(6) to update their estimates

6. Forward t to $t+1$ and repeat steps 1-5.

Before studying policy performance in this environment, we first use a decreasing gain parameter see whether agents can learn the rational expectations equilibrium corresponding to each of the parameterizations we consider. Initial beliefs about $a(s_t)$, $c(s_t)$ for $s_t = 1, 2$ are set to zero, while initial beliefs about $b(s_t)$ are perturbed around $\Omega(s_t)$.[11] For all parameterizations we consider here, beliefs eventually converge to their rational expectations equilibrium values. Figure 6 illustrates the convergence of beliefs for the posterior mean calibration with $\gamma(1) = 5$, $\gamma(2) = -5$, $\phi_\pi(1) = 3$, $\phi_\pi(2) = 0$. In this figure, as well Figure 9, we plot the difference of actual beliefs and rational expectations equilibrium beliefs over time (i.e. a value of 0 means that beliefs equal the rational expectations equilibrium).

To help better understand the impact that learning has on model dynamics, we study policy performance in a model with constant gain learning and a gain parameter equal to .01.[12] In such a model, we cannot compute the unconditional variance of inflation and output. We therefore approximate the variance of inflation and output by simulating the model for 100,000 periods and computing sample variances.[13] Because these simulations are more computationally intensive, we do not compute expected posterior losses as in section 3. Instead, we set non-policy parameters equal

[11]We set initial beliefs about the VAR coefficients away from zero (but still far from their REE values) to help improve the rate of convergence of beliefs to the REE. We also want to mention that beliefs may not converge to the rational expectations equilibrium for all initial values; E-stability is a local stability concept that only applies to beliefs that are in some neighborhood of their potential convergence points.

[12]The learning algorithm is also augmented with a ridge correction mechanism as in Slobydan and Wouters (2012), and projection facility that prevents estimates from updating if the updated parameters imply a Markov-switching VAR that is not mean-square-stable. Intuitively, the projection facility formalizes the notion that agents reject unstable models. We invoke the projection facility and ridge correction mechanism in far less than 1% of simulated periods

[13]Each simulation uses the same 100,000 realizations of shocks. We do this to help mitigate the potential for large outlier shocks to bias our sample variances.

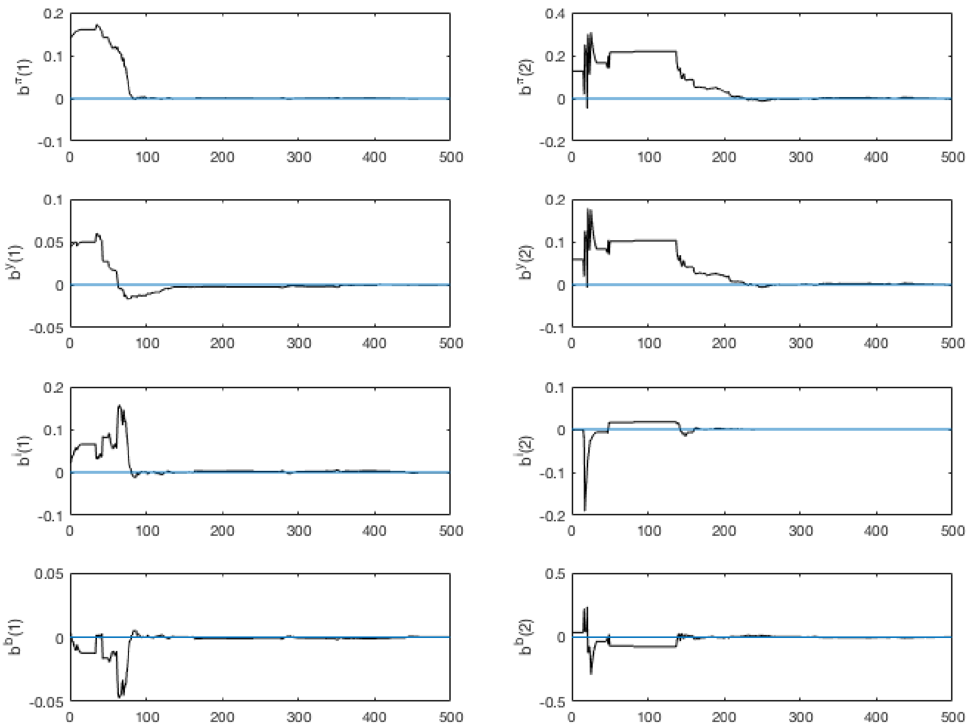

Figure 6: **Coefficient Estimate Errors and Observed State Learning:** the left-hand column features the VAR-coefficients on independent variable lagged debt in regime 1; right-hand column features the VAR-coefficients on independent variable lagged debt in regime 2. Notice that beliefs are held fixed when they correspond to an inactive state (e.g. notice the flat, "mesas" in the state 2 coefficients between t=50 and t=150).

to their posterior mean (with added cost-push shock), and make inferences based on this model calibration. Otherwise, the procedure for measuring performance is the same as the procedure in section 4: we search over monetary policy parameters and find the set of interest rate rule coefficients that minimizes the variance of inflation and output. Table 3 presents our main findings.

4.2 Unobserved Markov States (Hidden Markov Model)

The learning model in 4.1. provides valuable evidence that the optimized simple rules in section 3 are robust to misspecifications of private sector expectations. However, that model makes one potentially unreasonable assumption: agents observe the state of fiscal and monetary policy. In practice, applied econometricians do not observe

Type	$(\gamma(1),\ \gamma(2))$	**optimal RE coeff.** $(\phi_\pi(1),\ \phi_\pi(2),\ \rho_i)$	**optimal AL coeff.** $\phi_\pi(1),\ \phi_\pi(2),\ \rho_i)$	**Projection Facility (per 100,000)**
GS	$(5,\ -5)$	(3.3, 0, .99)	(3.68, 0, .99)	110
GA	$(2,\ -5)$	(0, 0, 0)	(0, 0, 0)	86

Table 3: Optimal coefficients under adaptive learning. The larger inflation reaction coefficients under learning echoes a result from Orphanides and Williams (2007). 4 is the largest inflation reaction coefficient used in this particular numerical search. Despite the small gain parameter and infrequent use of the projection facility, the model is frequently unstable for $\psi > .02$

the stance of fiscal and monetary policy. Instead, econometricians use techniques developed in papers such as Hamilton (1989) to identify the probable state of the economy at any point in time. Because a lot of adaptive learning research begins with the premise that our models' agents should be no more informed and rational than the econometricians among us, we endeavor in this section to remove s_t from the information set, I_t, and study the model-implied dynamics of inflation and output. We refer to the model developed in 4.2 as the hidden Markov model of learning. Before deriving the hidden Markov model of learning, we emphasize that self-referential feedback in this model not only poses the risk of destabilizing agents' beliefs about model coefficients; forecast errors act on both future coefficient estimates and agents' inferences about the underlying state. One may therefore expect additional expectations-induced volatility in this model.

As it turns out, the structure of our model makes it possible for agents to infer the underlying state with reasonable accuracy so that the removal of Markov states from agents' information set only raises the volatility of inflation and output slightly. This last point is partly explained by an argument made in Bianchi (2013) which states that fully rational agents can perfectly infer today's state if they observe contemporaneous and past x, z. Their argument relies on the fact that rational agents know all of the S within-regime systems of equations (i.e. $x_t = \Omega(s_t)x_{t-1} + \Gamma(s_t)z_t$) that may determine x_t. All agents in their model need to do to perfectly infer the state is compute each of the S equations until they find the correct system of equations. Their argument does not apply to our framework; if agents hold incorrect beliefs about the economy – as they always do in a model of constant gain learning, or before beliefs converge – they may make horrible inferences about the state of the economy. Despite this limitation, the equilibrium coefficients of the policy rules are exogenous to beliefs, which makes it easy for agents to learn the rational expectations equilibrium law of motion for fiscal surpluses and infer from it the underlying state of the economy with reasonable but far from perfect accuracy. We emphasize that other equilibrium coefficients do depend on agents' beliefs, so that our model is still self-referential.

As before, agents beliefs about the law of motion for endogenous variables is given by the PLM in (4). In what follows, we consider two information structures. First, we assume that $I_t = \{y_{t-1}, y_{t-2}, \ldots, y_0; z_{t-1}, \ldots, z_0\}$. After examining the potential convergence points of beliefs, and pointing out the exogeneity of the surplus law of motion, we then add surpluses, τ_t, to I_t and demonstrate that agents' beliefs can converge to the rational expectations equilibrium. Under both information structures, agents do not observe s_t, which implies that they cannot use the learning algorithm in 4.1 to update their beliefs. To get around the difficulty presented by the hidden Markov process, we rely on techniques from Krishnamurthy and Yin (2002) and LeGland and Mevel (1997), which present "online" or recursive algorithms for learning the coefficients of an exogenous Markov-switching autoregression. Specifically, we use the recursive maximum likelihood estimator (RMLE) from both papers, and the recursive conditional least squares estimator (RCLS) from LeGland and Mevel (1997). While newer alternatives to these algorithms exist outside of the stochastic approximation literature, we rely on these papers because they present convergence results that may prove useful in extensions of the current analysis.

The algorithms described in both papers make inferences about the coefficients, $\Phi(s_t)$, and the Markov process, s_t, using two related recursive processes. First, agents make inferences about s_t using a prediction filter of the form introduced by Hamilton (1989). To develop this filter we first define *within-regime* conditional densities for x, $f_{s_t} = f(x_t|x_{t-1}, x_{t-2}, ..., z_t, z_{t-1}, ..., s_t; \Phi(s_t)_{t-1})$. In a model with normally distributed *i.i.d* innovations to our exogenous driving process, f_{s_t} assumes the following form:

$$f_{s_t} = (2\pi)^{-t/2}|\Sigma|^{-.5} exp\{-.5(x_t - \mu(s_t)_{t-1})'\Sigma^{-1}(x_t - \mu(s_t)_{t-1})\}$$

where $\mu(s_t)_{t-1} = a_{t-1}(s_t) + b_{t-1}(s_t)x_{t-1} + c_{t-1}(s_t)z_t$ and Σ is the covariance-variance matrix for the *i.i.d* innovations to z. To make future calculations easier, we define the following matrices:

$$f_t = (f_{1t}, f_{2t} \ldots, f_{St})'$$

$$F_t = diag(f_{1t}, f_{2t} \ldots, f_{St})$$

Let $\hat{p}_{i,t|t-1} = Pr(s_t = i|I_t)$, and $\hat{p}_{t|t-1} = (\hat{p}_{1t|t-1}, \hat{p}_{2t|t-1}, \ldots, \hat{p}_{St|t-1})'$. $\hat{p}_t$ follows the recursion:

$$\hat{p}_{t+1|t} = \frac{P' F_t \hat{p}_{t|t-1}}{f_t' \hat{p}_{t|t-1}} \tag{7}$$

where it is assumed that agents know the true transition probabilities in P. The prediction filter in the last equation completely describes how agents recursively compute their predictions for today's state. Because inferences about s_t are made prior to time t, agents can, at best, infer s_{t-1} perfectly. As we show below, this feature of our model makes it impossible for agents' beliefs to converge to the rational expectations equilibrium studied in previous sections, and is the primary reason why we argue for the addition of τ_t to I_t. The second recursive process in the algorithms presented by Krishnamurthy and Yin (2002) and LeGland and Mevel (1997) updates the parameter estimates, $\Phi(s_t)$, according to:

$$\Phi_t = \Phi_{t-1} + \gamma S(x_t, I_t; \Phi_{t-1}) + \epsilon_t M_t$$

where Φ_t is a $k \times 1$ vector[14] that contains the elements of $\Phi(s_t)$ for all s_t, γ is the gain parameter and M_t is a correction term that brings Φ_t into some reasonably defined constraint set G (i.e. we use a projection facility in our implementation of their algorithms). Let Φ_t^l denote the l-th element of Φ_t. The function $S(x_t, I_t; \Phi_{t-1})$ is the only thing that varies across the two algorithms we use in the paper. For the RMLE algorithm, $S(x_t, I_t, \Phi_{t-1})$ is given by the following equations:

$$S(x_t, I_t, \Phi_{t-1}) \;\; = \;\; (S^1(x_t, I_t, \Phi_{t-1}), \ldots, S^k(x_t, I_t, \Phi_{t-1}))'$$

where

$$S^l(x_t, I_t, \Phi_{t-1}) = \frac{f_t' \omega_t^l}{f_t' \hat{p}_{t|t-1}} + \frac{(\partial f_t' / \partial \Phi_t^l) \hat{p}_{t|t-1}}{f_t' \hat{p}_{t|t-1}} \tag{8}$$

for all $l \in \{1, ..., k\}$ and $\omega_t^l = \frac{\partial \hat{p}_{t|t-1}}{\partial \Phi_t^l}$. We update ω_t^l recursively as follows:

$$\omega_{t+1}^l = R_{1t} \omega_t^l + R_{2t} \tag{9}$$

where

$$\begin{aligned} R_{1t} &= P'(I - \frac{F_t \hat{p}_{t|t-1} \mathbf{1}_s'}{f_t' \hat{p}_{t|t-1}}) \frac{F_t}{f_t' \hat{p}_{t|t-1}} \\ R_{2t} &= P'(I - \frac{F_t \hat{p}_{t|t-1} \mathbf{1}_s'}{f_t' \hat{p}_{t|t-1}}) \frac{(\partial F_t)/(\partial \Phi_t^l) \hat{p}_{t|t-1}}{f_t' \hat{p}_{t|t-1}} \end{aligned}$$

Equation (7), (8) and (9), plus initial conditions, give us the RMLE algorithm. To derive the RCLS we only need to change our definition of $S^l(x_t, I_{t-1}, \Phi_{t-1})$ as follows:

$$S^l(x_t, I_t, \Phi_{t-1}) = (\phi_{\Phi_{t-1}}(x_t - \phi_{\Phi_{t-1}}' \hat{p}_{t|t-1}))' \omega_t^l + (\frac{\partial \phi_{\Phi_{t-1}}}{\partial \Phi_{t-1}^l}(x_t - \phi_{\Phi_{t-1}}' \hat{p}_{t|t-1}))' \hat{p}_{t|t-1} \tag{10}$$

where $\phi_{\Phi_{t-1}}$ is a matrix that collects the conditional mean for each state (i.e. $\mu(s)_{t-1}$ for each $s \in \{1, \ldots, S\}$). Before outlining the events leading to a temporary equilibrium, we emphasize that this algorithm is very similar to the algorithm presented in section 4.1. Specifically, if agents observe the state so that ω becomes a vector of zeros (since $(I - \frac{F_t \hat{p}_{t|t-1} \mathbf{1}_s'}{f_t' \hat{p}_{t|t-1}} \to 0_S)$, and they replace $\hat{p}_{t|t-1}$ with $\hat{p}_{t|t} = (1\ 0)'$ or $\hat{p}_{t|t} = (0\ 1)'$ to reflect this knowledge, then this algorithm becomes the recursive estimator of section 4.1 with $R(s_t)_{s_t} = I$. We can now outline the sequence of events that lead to an equilibrium at time t:

[14] in our model with $S = 2$, n endogenous variables and m endogenous variables, $k = 2(n(n+1) + nm)$

1. Agents update information sets.

2. Using I_t and time $t-1$ estimates $a_{t-1}(s_t)$, $b_{t-1}(s_t)$, $c_{t-1}(s_t)$. Agents form forecasts, $\hat{E}_t x_{t+1}$:

$$\begin{aligned}
\hat{E}_t x_{t+1} \quad = \quad & (\hat{p}_{1t|t-1}p_{11} + \hat{p}_{2t|t-1}p_{21})a(1) + (\hat{p}_{1t|t-1}p_{12} + \hat{p}_{2t|t-1}p_{22})a(2) + \\
& \hat{p}_{1t|t-1}p_{11}b(1)(a(1) + b(1)x_{t-1} + c(1)z_t) + \\
& \hat{p}_{1t|t-1}p_{12}b(2)(a(1) + b(1)x_{t-1} + c(1)z_t) + \\
& \hat{p}_{2t|t-1}p_{21}b(1)(a(2) + b(2)x_{t-1} + c(2)z_t) + \\
& \hat{p}_{2t|t-1}p_{22}b(2)(a(2) + b(2)x_{t-1} + c(2)z_t) + \\
& ((\hat{p}_{1t|t-1}p_{11} + \hat{p}_{2t|t-1}p_{21})c(1) + (\hat{p}_{1t|t-1}p_{12} + \hat{p}_{2t|t-1}p_{22})c(2))\rho z_t
\end{aligned}$$

3. x_t is generated from the actual law of motion, (3), which gives us time t endogenous variables as a function of beliefs and predetermined variables

4. Agents observe x_t and add it to their information sets

5. Agents use (7), (8), (9), or (7), (9), and (10) to update their coefficient estimates and prediction filter

6. Forward t to $t+1$ and repeat steps 1-5.

Before presenting results, it is important to note that our information structure in 4.2 prevents agents from learning the rational expectations equilibrium studied in all previous sections. This is because agents only form $\hat{p}_t$ using $t-1$ information. Hence, if agents perfectly infer s_{t-1} – which is the best they can do – they still hold the following beliefs about s_t: $\hat{p}_{t|t-1} = (p_{s_{t-1}1}, p_{s_{t-1}2})' < (1,1)'$. In this best case scenario, agents' beliefs about the VAR coefficients, $b(s_t)$, will not converge to a solution of (3). If, instead, agents allow their beliefs about PLM coefficients to depend on both s_t and s_{t-1} then this information structure may allow agents to learn solutions to the following fixed point condition:

$$b(s_t, s_{t-1}) = A(s_t) \sum_{j=1}^{2} \sum_{h=1}^{2} p_{s_{t-1}j} p_{jh} b(h,j) b(j, s_{t-1}) + B(s_t) \tag{11}$$

These solutions, which we refer to as history-dependent equilibria, do solve (3). However, they do not satisfy the following fixed point condition:

$$b(s_t) = A(s_t)(p_{s_t 1} b(1) + p_{s_t 2} b(2)) b(s_t) + B(s_t) \tag{12}$$

which is a necessary condition for solutions of the form, $b(s_t)$. While beliefs are no longer consistent with the rational expectations equilibria we examined up until now, we nonetheless find that beliefs can converge.[15] Hence, while beliefs never converge to the rational expectations equilibrium, they may nonetheless be stable over time and converge to values that may be relatively close to the original rational expectations equilibrium.

To identify potential convergence points consistent with (11), we use the Gröbner basis approach from Foerster et al (2016). We then explore issues of uniqueness and E-stability pertaining to this class of equilibria. Initial evidence suggests that policy parameters widely associated with determinacy in the preceding analysis may admit multiple mean-square stable history dependent equilibria that satisfy the fixed point condition in (11). Moreover, these equilibria do not appear to be stable under learning. Since this class of equilibria is arguably relevant in settings where agents cannot observe contemporaneous variables, we intend to further explore these issues of uniqueness and expectational stability in future work.

Figure 7 plots $\hat{p}_1$ over time. In our calibration $p_{11} = .95$ so that oscillation in their beliefs between .05 and .95 implies that they're inferring s_{t-1} almost perfectly. To better understand how agents so successfully infer the underlying state of the economy, despite initial incorrect beliefs about the structure of the economy, we redefine $x = (\tilde{x}, \tau)'$ where $\tilde{x} = (y, \pi, i, b, P)'$ and point out that the actual law of motion for x (after beliefs are substituted in) may be written as:

$$\begin{pmatrix} \tilde{x}_t \\ \tau_t \end{pmatrix} = \begin{pmatrix} \tilde{\Omega}(s_t; \Phi_{t-1}) \\ \Omega_\tau(s_t) \end{pmatrix} \begin{pmatrix} \tilde{x}_{t-1} \\ \tau_{t-1} \end{pmatrix} + \begin{pmatrix} \tilde{\Gamma}(s_t; \Phi_{t-1}) \\ \mathbf{e}_6' \end{pmatrix} z_t \tag{13}$$

where $\Omega_\tau(s_t) = (0\ 0\ 0\ \gamma(s_t)\ 0\ 0)$, and $\mathbf{e}_6' = (0\ 0\ 0\ 0\ 0\ 1)$. Clearly, the evolution of τ_t is only endogenous to beliefs through b_{t-1}; the coefficients governing the evolution of τ are exogenous, which suggests that agents will quickly learn the law of motion for τ and then make accurate inferences for $\hat{p}_t$ that rely on the marginal density:

$$f_{s_t}^\tau = f(\tau_t | x_{t-1}, \Phi_{t-1}) \tag{14}$$

The marginal density in (14) is so essential for correct inference of s_{t-1} that we can redefine our prediction filter using only the marginal densities for surpluses and get results that are nearly identical to the results displayed in Figure 7. The fact that surpluses are determined at the beginning of t (i.e. all shocks and b_{t-1} have been realized by beginning of t, so that τ_t is fixed before agents form expectations), begs an important question about timing: should agents be able to observe τ_t at the beginning of t? That is, should I_t include τ_t? If agents observe τ_t at t, they may be able to perfectly infer s_t. This allows for the fixed conditions in (11) and (12) to coincide, so that agents may actually learn the rational expectations equilibrium

[15] Even for constant gain parameters beliefs appear to converge to a distribution around a fixed point

under study. To support this idea numerically, we first redefine the prediction filter:

$$
\begin{aligned}
f_t^\tau &= (f_{1t}^\tau, f_{2t}^\tau \ldots, f_{St}^\tau)' \\
F_t^\tau &= diag(f_{1t}^\tau, f_{2t}^\tau \ldots, f_{St}^\tau) \\
\hat{p}_{t|t}^\tau &= \frac{F_t^\tau \hat{p}_{t|t-1}^\tau}{f_t'^\tau \hat{p}_{t|t-1}^\tau} \\
\hat{p}_{t+1|t}^\tau &= P' \hat{p}_{t|t}^\tau
\end{aligned}
$$

Now agents use $\hat{p}_{t|t}^\tau$ instead of $\hat{p}_{t|t-1}$ when forming expectations at time t. As shown in Figure 8 agents can now infer the current state very effectively, which allows them to learn the rational expectations equilibrium under study in the previous section, as demonstrated by Figure 9. In Figure 9, we initialize beliefs away from the rational expectations equilibrium[16], set $\psi = t^{-2/3}$ (as in LeGland and Mevel (1997)) and estimate the model using the RCLS algorithm. We also use a projection facility that prevents agents from accepting a mean-square-unstable PLM, but this facility is invoked in far less than .1% of periods simulated. Compared to Figure 6, the rate of convergence is slow under RCLS, but this may be driven the errors in the prediction filter (Figure 7) and the large decreasing gain parameter $t^{-2/3}$. We find that the optimal policy results from section 4.1 generalize to the hidden Markov model of learning.

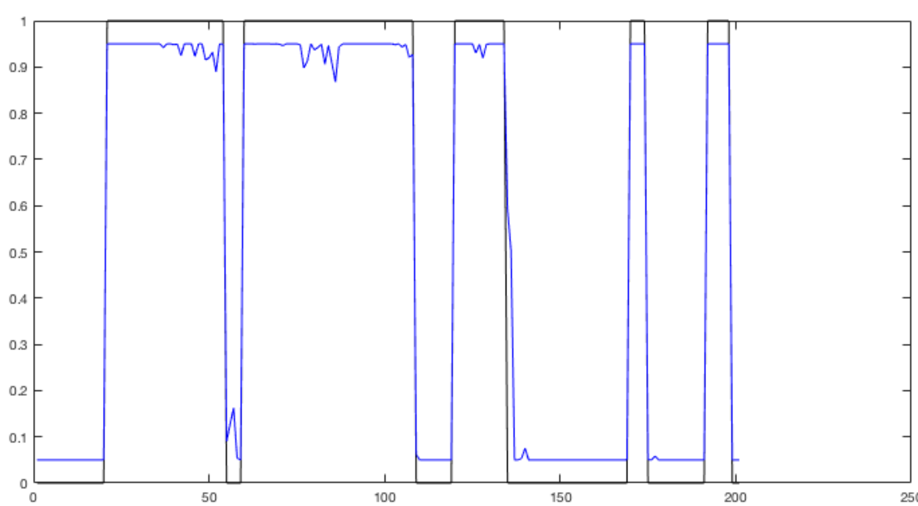

Figure 7: Blue line is $\hat{p}_{1,t|t-1}$; black line equals 1 if $s_t = 1$ and 0 otherwise

[16]As seen in the third subplot in the second column of Figure 9, initial beliefs about the dependence of i on b in regime F are unintentionally close to 0. Our results do not depend on this initial belief.

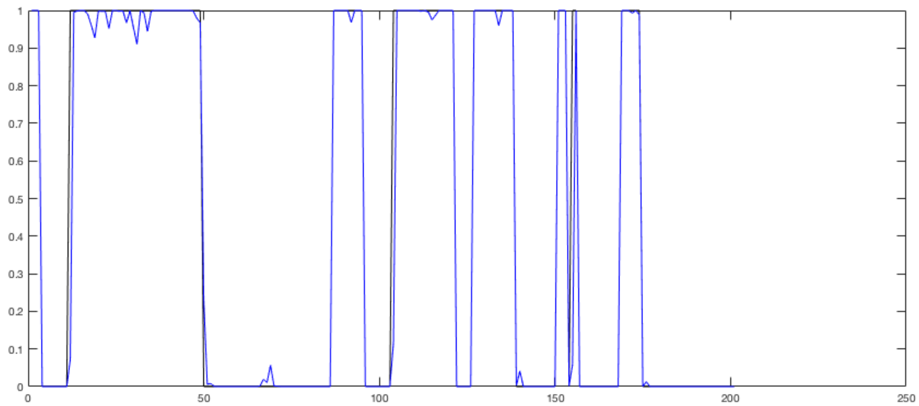

Figure 8: Blue line is $\hat{p}_{1,t|t}$; black line equals 1 if $s_t = 1$ and 0 otherwise

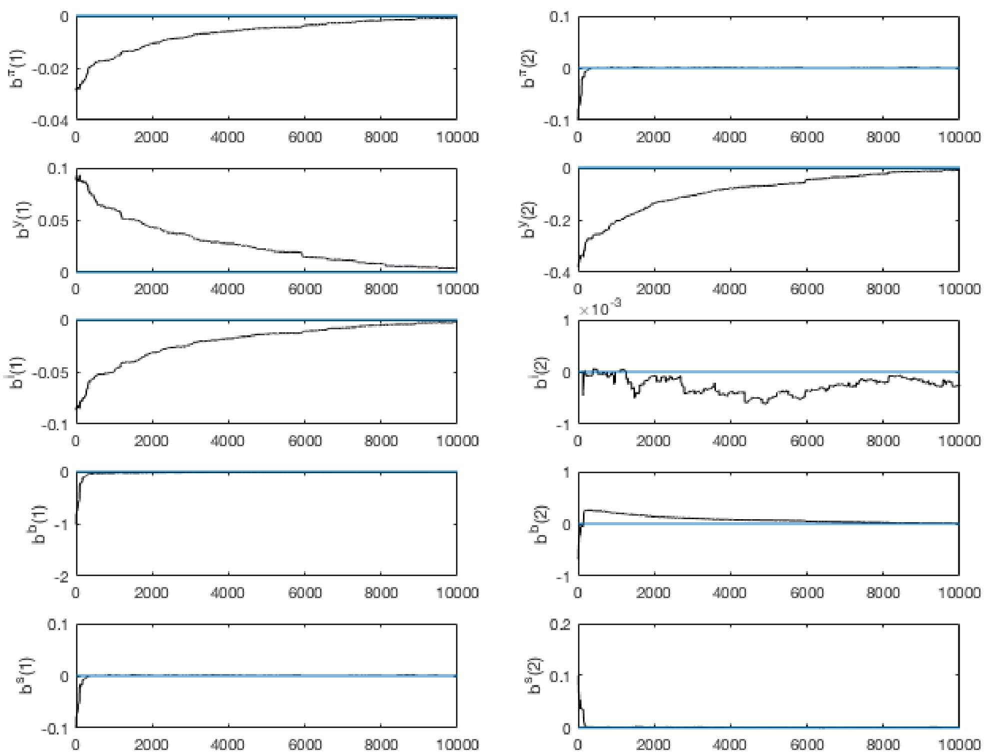

Figure 9: **Coefficient Estimate Errors in Hidden Markov Model Learning:** the left-hand column features the VAR-coefficients on independent variable lagged debt in regime 1; right-hand column features the VAR-coefficients on independent variable lagged debt in regime 2.

5. Long-term debt

In this section, we relax the assumption that $\rho = 0$ and introduce long-term debt into our model. While this innovation helps to bring our model closer to reality, it also creates a debt revaluation channel through which monetary and fiscal policy interact to affect agents' perceptions of bond wealth in non-Ricardian economies. This debt revaluation works as follows: if interest rates are reduced (increased), then the price of outstanding debt, given by $\hat{P^m}$, increases (decreases) and this positively (negatively) affects agents' perception of their own net wealth. This revaluation channel can often lean against the wealth effects created by movements in debt service costs, a tendency demonstrated in a host of papers (for example, see Leeper and Leith (2016)).

In addition to the creation of a revaluation channel, the introduction of long-term debt can alter the menu of monetary policies that induce a determinate equilibrium when fiscal policy switches and is non-Ricardian (see McClung (2017b)). The fact that maturity matters for determinacy in our simple switching DSGE model is a novel result insofar as the average maturity of debt does not matter for determinacy in the corresponding fixed regime model (see Jin (2013), for example). To illustrate the impact that maturity has on determinacy consider Figures 10 and 11.

The fact that maturity matters for determinacy complicates the monetary policymaker's problem in at least two ways. First, the policymaker now has an incentive to identify the steady state average maturity given by ρ when fiscal policy is non-Ricardian. Without knowledge of ρ, the policymaker cannot properly identify the menu of policies that induce a unique equilibrium, which may prevent them from finding the optimized policy. Second, the policymaker now has an incentive to consider balance sheet decisions that affect the value of ρ when solving their optimization problem. In our model, the relevant measure of government debt is government debt held by the household, not purely the debt issued by the fiscal authority itself. As such, central banks can impact the maturity structure of debt held by households by engaging in Operation Twist-style policies in which households and the monetary authority exchange short-term debt for long-term debt. Figure 10 illustrates a case where monetary policymakers may realize an incentive to lengthen the maturity debt held by households, while Figure 11 illustrates the opposite case. We hope to use ρ as a proxy for these debt operations by adding ρ to the central bank's choice set.

The search for an optimal ρ is further complicated by the fact that we face uncertainty over the true value of ρ.[17] We might address this uncertainty by assigning a prior distribution to ρ, adding ρ to θ then estimating the model using Bayesian techniques. Using simple priors over ρ we can generate a tradeoff between expected posterior loss and the probability that a given policy implements a unique mean-square

[17]the average maturity of debt, equal to $(1 - \beta\rho)^{-1}$ in our model has been estimated using U.S. data. However, because the actual maturity structure of U.S. debt does not decay geometrically, it is not clear whether or not such estimates should be used to select ρ

stable and E-stable equilibrium. For one simple prior over ρ, the policy that maximizes the probability of determinacy and E-stability (at .985) when $p_{11} = p_{22} = .95$, $\gamma(1) = 5$, $\gamma(2) = -1$, involves $\phi_\pi(1) = 1.2$, $\phi_\pi(2) = .9$ and an expected posterior loss of 4.15. If we replace $\phi_\pi(1) = 1.2$, $\phi_\pi(2) = .9$ with $\phi_\pi(1) = 1.3$, $\phi_\pi(2) = .8$, we reduce the probability of determinacy and E-stability to .917, but we also reduce expected posterior loss to 2.57. The addition of uncertainty over ρ therefore introduces a tradeoff between minimizing loss and maximizing the probability of determinacy and E-stability, a tradeoff first recognized by Evans and McGough (2007). Uncertainty over all other fiscal policy parameters will almost surely present a similar tradeoff.

We believe that parameter uncertainty and the addition of an extra dimension in ρ to our policy problem generates complications that are beyond the scope of the present analysis. However, we hope to fully explore issues pertaining to long-term debt in an estimated DSGE framework in the near future.

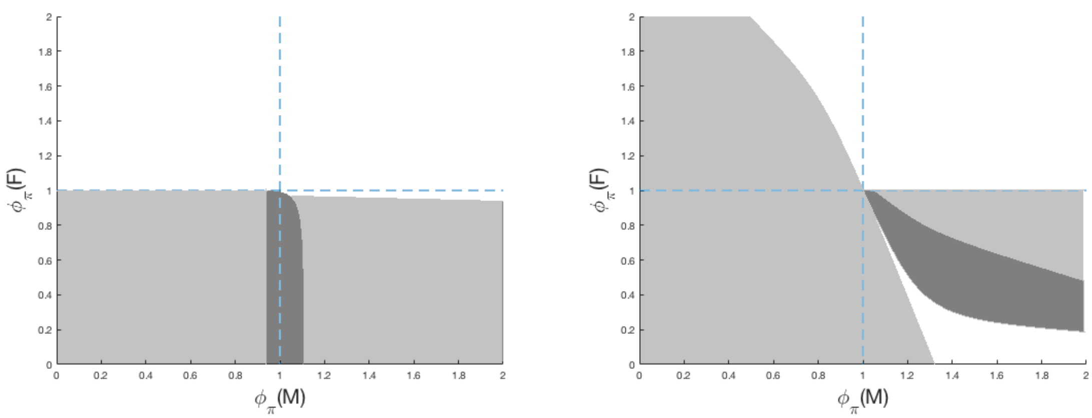

Figure 10: Left panel $\rho = 0$, right panel $\rho = .96$. $p_{mm} = .98$, $p_{ff} = .95$, $\gamma(M) = .02$, $\gamma(F) = -.01$. The determine region is dark gray; the indeterminate region is light gray; explosive region is white

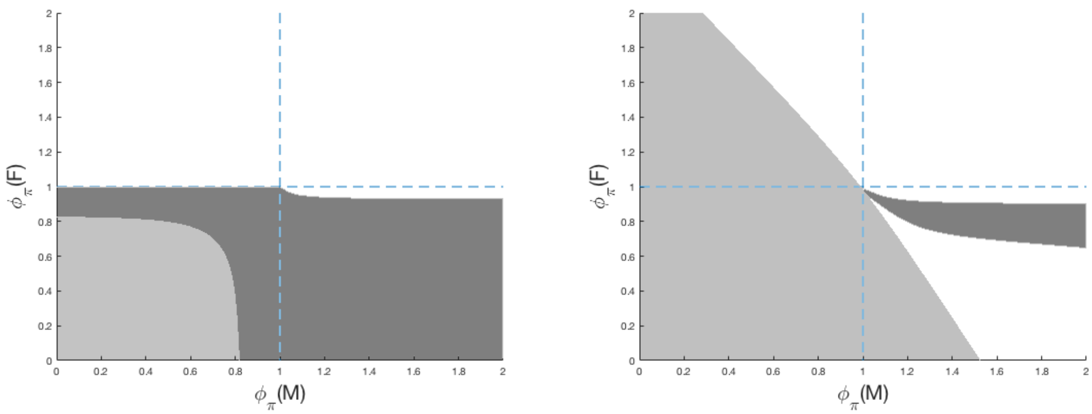

Figure 11: Left panel $\rho = 0$, right panel $\rho = .96$. $p_{mm} = .95$, $p_{ff} = .95$, $\gamma(M) = .05$, $\gamma(F) = -.05$. The determine region is dark gray; the indeterminate region is light gray; explosive region is white

6. Conclusion

This paper examines the performance and robustness of simple monetary policy rules in models with learning agents subject to: (1) permanent or occasionally non-Ricardian fiscal policy; and/or (2) the presence of long-term government debt. My analysis indicates that the "global" response of the fiscal policymaker to debt determines the optimal monetary policy response. When fiscal policy is globally passive or globally active the optimal monetary policy rule features time-invariant coefficients with high inflation reaction coefficients in globally passive models and interest rate pegs in globally active models. In cases where fiscal policy features balanced or strong switching between active and fiscal policy stances, the optimal monetary policy rule features switching coefficients. These results are robust to adaptive learning, including a novel hidden Markov model of learning we introduce in the paper. For this reason, we should want to better understand how the presence of long-term debt affects the optimal monetary policy in a model with switching fiscal policy stances.

Works Referenced

An, S., and F. Schorfheide (2007): “Bayesian Analysis of DSGE Models,” *Econometric Reviews*, 26(2-4), 113-172.

Ascari, G., Florio, A., and A. Gobbi (2017): “Controlling Inflation with switching monetary and fiscal policies: expectations, fiscal guidance and timid regime changes,” Working Paper.

Bianchi, F. (2012): “Evolving Monetary/Fiscal Policy Mix in the United States,” *American Economic Review Papers and Proceedings*, 101(3), 167-172.

——— (2013): “Regime Switches, Agents’ Beliefs, and Post-World War II U.S. Macroeconomic Dynamics,” *Review of Economic Studies*, 80(2), 463-490.

———, and L. Melosi (2014): “Dormant Shocks and Fiscal Virtue,” *NBER Macroeconomics Annual 2013*, 28, 1-46.

——— , and C. Ilut (2017): “Monetary/Fiscal Policy Mix and Agents’ Beliefs,” *Review of Economic Dynamics*, 26, 113-139.

Bullard, L. and A Singh (2012): ”Learning And The Great Moderation,” *International Economic Review*, 53(2), 375-397.

Chen, X., E. M. Leeper, and C. Leith (2015): “U.S. Monetary and Fiscal Policy: Conflict or Cooperation?,” Manuscript, University of Glasgow.

Cho, S. (2016): “Sufficient Conditions for Determinacy in a Class of Markov-Switching Rational Expectations Models,” *Review of Economic Dynamics*, 21, 182-200.

———, A. Moreno (2016): “Global Determinacy Under Monetary and Fiscal Policy Switchings,” Working Paper.

Cogley, T., B. De Paoli, C. Matthes, K. Nikolov, and T. Yates (2011): “A Bayesian Approach to Optimal Monetary Policy with Parameter and Model Uncertainty,” *Journal of Economic Dynamics and Control*, 35(12): 2186-2212.

Davig, T. (2005): “Regime-switching debt and taxation,” *Journal of Monetary Economics* 51(4), 837-859.

Davig, T., and E.M. Leeper (2006): “Fluctuating Macro Policies ad the Fiscal Theory,” in *NBER Macroeconomic Annual 2006*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford. MIT Press, Cambridge, MA.

——— (2007): “Generalizing the Taylor Principle,” *American Economic Review*, 97(3), 607-635.

——— (2011): “Monetary-Fiscal Policy Interactions and Fiscal Stimulus,” *European Economic Review*, 55(2), 211-227.

Eusepi, S., and B. Preston (2011): “Learning about the Fiscal Theory of the Price

Level: Some Consequences of Debt-Management Policy," *Journal of the Japanese and International Economies*, 25(4), 358-379.

——— (2012): "Debt, Policy Uncertainty and Expectations Stabilization," *Journal of the European Economic Association*, 10(4), 860-886.

——— (2013): "Fiscal Foundations of Inflation: Imperfect Knowledge," Working Paper.

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

Evans, G.W. and B. McGough (2007): "Optimal Constrained Interest-Rate Rules," *Journal of Money, Credit and Banking* 39(6), 1335?1356.

Foerster, A., Rubio-Ramirez, J., Waggoner, D., and T. Zha (2016): "Perturbation Methods for Markov-switching Dynamic Stochastic General Equilibrium Models," *Quantitative Economics*, 7(2), 637-669.

Gonzalez-Astudillo, M. (2013): "Monetary-Fiscal Policy Interaction: Interdependent Policy Rule Coefficients," Finance and Economics Discussion Series No. 2013-58, Federal Reserve Board.

Hamilton, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica*, 57, 357-384.

Hansen, L. and N.G. Polson, (2006): "Tractible Filtering." Working paper, University of Chicago.

——— N.G. Polson, and T. Sargent (2010): "Nonlinear Filter and Robust Learning," Working Paper.

——— and T. Sargent (2010): "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics*, 1(1), 129-162.

Johannes, M, A. Korteweg and N. Polson (2014): "Sequential Learning, Predictability, and Optimal Portfolio Returns," *The Journal of Finance*, 69(2): 611-644.

Kim, C., and C. Nelson (1999). State-Space Models with Regime Switching. MIT Press: Cambridge, Massachusetts.

Kleim, M., A. Kriwoluzky, and S. Sarferaz (2015): "On the Low-Frequency Relationship Between Public Deficits and Inflation," *Journal of Applied Econometrics*, 31(3), 566?583.

——— (2016): "Monetary-Fiscal Policy Interaction and Fiscal Inflation: A Tale of Three Countries," *European Economic Review*, 88, 158-184.

Krishnamurthy, V., and G. Yin (2002): "Recursive Algorithms for Estimation of Hidden Markov Models and Autoregressive Models With Markov Regime," *IEEE Transactions of Information Theory*, 48(2), 458-476.

LeGland, F., and L. Mevel (1997): "Recursive estimation of hidden Markov models," *Proc. 36th IEEE Conf. Decision Control*, San Diego, CA, Dec. 1997.

Leeper, E.M. (1991): "Equilibria under 'active' and 'passive' monetary and fiscal policies," *Journal of Monetary Economics*, 27(1), 129-147,

————, and C. Leith (2016): "Understanding Inflation as a Joint Monetary-Fiscal Phenomenon" in *Handbook of Macroeconomics, Volume 2*, ed. by J. Taylor and H. Uhlig, pp. 2305-2415.

McClung, N. (2016): "E-Stability vis-a-vis Determinacy in Markov-Switching DSGE Models," Working Paper.

McClung, N. (2017a): "The Power of Forward Guidance and the Fiscal Theory of the Price Level," submitted.

McClung, N. (2017b): "Maturity, Determinacy and E-Stability in Models of Regime Switching Fiscal and Monetary Policy," Working Paper.

Orphanides, A., and J. Williams (2007): "Robust Monetary Policy with Imperfect Knowledge," *Journal of Monetary Economics*, 54(5), 1406-1435.

Schmitt-Grohe, S., and M. Uribe (2007): "Optimal simple and implementable monetary and fiscal rules," *Journal of Monetary Economics*, 54: 1702-1725.

Slobodyan, S., and R. Wouters. (2012): "Learning in a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," *American Economic Journal: Macroeconomics*, 4(2), 65-101.

Woodford, M. (1998a): "Control of the Public Debt: A Requirement for Price Stability," in *The Debt Burden and Its Consequences for Monetary Policy,* ed. by G. Calvo, and M. King, pp. 117-154. St. Martin's Pres, New York.

———— (2001): "Fiscal Requirement for Price Stability,? *Journal of Money, Credit, and Banking* 33(3), 669-728.

———— (2003). Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press: Princeton, New Jersey.

Appendix

A.1.

We present a simple model that is inspired by An and Schorfheide (2007). As in standard in the New Keynesian literature, the model consists of households, a competitive final goods producing firm, monopolistically competitive intermediate firms, a fiscal authority and a monetary authority. We briefly describe the optimization problems facing agents in this economy, then we collect the equilibrium conditions which are log-linearized and presented in section 2.

Households maximize a lifetime utility functions that depends positively on the level of consumption, C_t and negatively on labor supply, N_t. Additionally, households are subjected to a preference shock, Z_t that directly impacts the contribution of time t utility to overall lifetime utility. Formally:

$$\max_{\{C_t,N_t,W_t\}} E_o \sum_{t\geq 0} \left(\frac{C_t^{1-\sigma}-1}{1-\sigma} - \chi N_t \right) Z_t$$

subject to

$$P_t C_t + E_t(R_{t,t+1} W_{t+1}^j) \quad \leq \quad W_t^j + P_t \omega_t N_t - P_t \tau_t$$

and a tranversality condition of the form:

$$\lim_{t\to\infty} E_t[R_{t,T} W_T] = 0$$

where W_t is wealth at time t, ω is the competitive real wage paid to labor, τ is a lump-sum tax, C is consumption, and $R_{t,t+1}$ is a stochastic discount factor that equals $(C_{t+1}/C_t)^{-\sigma}$ in our model with complete markets. From the first order conditions for W_{t+1}, C_t and W_t we get the familiar necessary intertemporal and intratemporal conditions for the household optimization problem:

$$1 \quad = \quad \beta E_t \left\{ \frac{C_{t+1}}{C_t}^{-\sigma} \frac{Z_{t+1}}{Z_t} \frac{(1+i_t)}{\pi_{t+1}} \right\} \tag{15}$$

$$\omega_t \quad = \quad \chi C_t^{\sigma} \tag{16}$$

The perfectly competitive firm has technology described by:

$$Y_t = \left(\int_0^1 Y_t(j)^{1-\eta_t} dj \right)^{\frac{1}{1-\eta_t}}$$

where inputs, $Y_t(j)$, are goods produced by each intermediate firm $j \in [0,1]$, and η_t is a shock to markups. The perfectly competitive firm maximizes profits given by:

$$\Pi_t^{FIN} = P_t Y_t - \int_0^1 P_t(j) Y_t(j) dj$$

This implies the following demand schedule for each intermediate producer's good, $Y_t(j)$:

$$\begin{aligned} Y_t(j) &= \left(\frac{P_t(j)}{P_t}\right)^{-1/\eta_t} Y_t \\ P_t(j) &= \left(\int_0^1 P_t(j)^{\frac{\eta_t - 1}{\eta_t}}\right)^{\frac{\eta_t}{\eta_t - 1}} \end{aligned}$$

Intermediate firms are monopolistically competitive and utilize identical technologies that assume the form:

$$Y_t(j) = N_t(j)$$

To introduce nominal rigidities, we assume that firms face the following adjustment costs:

$$AC_t(j) = \frac{\phi}{2}\left(\frac{P_t(j)}{P_{t-1}(j)} - \pi\right)^2 Y_t(j)$$

Firms maximize the present value of firm profits taken real wages, ω_{t+s} as given. Formally, they choose labor inputs and prices to maximize the following:

$$\Pi^{INT} = E_0 \left\{ \sum_{t \geq 0} \beta^t R_{0,t} \left(\frac{P_t(j)}{P_t} Y_t(j) - \omega_t(j) N_t(j) - AC_t(j) \right) \right\}$$

Substituting the product demand schedule into the profits equation, then optimizing with respect to $P_{t+s}(j)$ and substituting for $\omega_{t+s} = c_t^\sigma$ and $R_{t|0} = (C_t/C_0)^{-\sigma}$ yields the following optimality condition:

$$\begin{aligned} \left(\frac{1}{\eta_t} - 1\right) &= \frac{C_t^\sigma}{\eta_t} - \frac{\phi}{2}\left(2(\pi_t - \pi) - \frac{(\pi_t - \pi)^2}{\eta_t}\right) && (17) \\ &+ \beta\phi\left(\left(\frac{C_{t+1}}{C_t}\right)^\sigma (\pi_{t+1} - \pi)\, \pi_{t+1} \frac{Y_{t+1}}{Y_t}\right) && (18) \end{aligned}$$

The fiscal authority only issues a bond portfolio, B_t^m, with a maturity that declines at a rate $\rho \in [0,1]$. Under this maturity structure, the quantity of government debt issued at $t-1$ that matures at $t+j$ is:

$$B_{t-1}(t+j) = B_{t-1}^m \rho^j$$

The evolution of the government's bond portfolio satisfies that following budget constraint:

$$B^m_{t-1}(1 - \sum_{j\geq 0} Q_t(t+j)\rho^j) + P_t G_t = P_t \tau_t + B^m_t \sum_{j\geq 0} Q_t(t+j)\rho^{j-1} \tag{19}$$

where $Q_t(t+j)$ is the price of debt that matures at time $t+j$ and is sold at t. To simplify the government budget constraint, we define the price of the bond portfolio, P^m_t, as:

$$P^m_t = E_t \sum_{j\geq 0} Q_t(t+j)\rho^{j-1}$$

which allows us to rewrite the government budget constraint as

$$B^m_{t-1}(1 + \rho P^m_t) + P_t G_t = P_t \tau_t + P^m_t B^m_t \tag{20}$$

Furthermore, we can show that bond prices follow a recursive formulation:

$$P^m_t = Q_t(t+1)(1 + \rho E_t P^m_{t+1}) \tag{21}$$

given B^m_{-1}. The government also implements a rule that adjusts real primary surpluses in response to the market value of real debt. In equilibrium, households hold all government debt which requires that the following condition hold $\forall t$:

$$W_t = B^m_{t-1}(1 + \rho P^m_t)$$

The processes for τ_t and G_t are specified. Finally, monetary policy follows the following rule:

$$R_t = R^{\rho_i}_{t-1}\left(R^* \left(\frac{\pi_t}{\pi^*}\right)^{\phi_\pi(s_t)} \left(\frac{Y_t}{Y^*_t}\right)^{\phi_y(s_t)}\right)^{1-\rho_i} \tag{22}$$

where $R_t = 1 + i_t$, $R^* = \beta^{-1}$, Y^*_t is potential output defined as the level of output that obtains without nominal rigidities and with constant markups. The log-linearized equilibrium conditions in section 2, are simply log-linearized versions of equations (15), (17), (20)-(22). μ_t is a composite of η_t from (17), and all other shocks and the fiscal policy rule are described in section 2.

A.2. Prior and Posterior Distributions

Table 4: We estimate the model as in An and Schorfheide (2007) using U.S. data from Q1:1983 to Q3:2007

Name	**Prior Density**	**Prior Param (1)**	**Prior Param (2)**	**Posterior Mean**
σ	Gamma	2.00	0.50	2.36
κ	Uniform	0.00	1.00	.91
ϕ_π	Gamma	1.50	0.25	2.16
ϕ_y	Gamma	0.50	0.25	.56
ρ_i	Uniform	0.00	1.00	.71
ρ_g	Uniform	0.00	1.00	.98
ρ_z	Uniform	0.00	1.00	.93
$100\sigma_m$	InvGamma	0.40	4.00	.2
$100\sigma_g$	InvGamma	1.00	4.00	.75
$100\sigma_z$	InvGamma	0.50	4.00	.2

Param (1) and Param (2) are the lower and upper bounds for the uniform distributions and the mean and standard deviation for the Gamma and Inverse Gamma distributions