Estimating the FOMC's Interest Rate Rule with Variable Selection and Partial Regime Switching

Adam Check*

August 10, 2019

Abstract

In many recent empirical studies of the Federal Open Market Committee's (FOMC's) interest rate rule, the parameters of the rule are allowed to change over time. However, within this literature, there is no consensus about the nature of the parameter change. Some authors, such as Sims and Zha (2006) only find evidence for a change in the variance of the interest rate rule, while others such as Gonzalez-Astudillo (2018) find evidence for changes in inflation and output gap responses. In this paper, I develop a new two-regime Markov-switching model that probabilistically performs variable selection and identification of parameter change for each variable in the model. After performing Bayesian estimation of this model and allowing for stochastic volatility, I find substantial evidence that there have been changes in the FOMC's response to the unemployment gap and in the volatility of the rule, but a low probability that there have been changes in the response to the inflation gap or any of the other parameters.

JEL Codes: C22, C24, C51, C52, E52

*I can be contacted at ajc@stthomas.edu

†I would like to thank Jeremy Piger, Dean Croushore, Mark Kuperberg, and Eric Gaus as well as participants in the Liberal Arts Macro, Midwest Economic Research Group, and WEAI conferences for helpful comments and suggestions on previous versions of this paper.

1 Introduction

The modern view of monetary policy is that the Federal Open Market Committee (FOMC) adjusts the nominal Federal Funds rate based on measures of economic performance. Mathematically, this is typically formulated as a version of the the Taylor Rule, first described in Taylor (1993), in which the target nominal Federal Funds rate is a linear function of output and inflation. This policy rule, and others taking very similar forms, are the foundation of the past two decades of empirical analysis of historical FOMC behavior.

Many researchers allow for the possibility that the coefficients of the policy rule change over time. In part, this is due to observed macroeconomic variables, such as inflation, changing substantially over time, suggesting possible changes in FOMC priorities. Additionally, narrative approaches such as Romer and Romer (2004) have found systematic differences in policy implementation under different Federal Reserve chairmen. The possibility of changing coefficients in the policy rule has been modeled many different ways, and many different measures of inflation, output, and employment have been used when estimating the policy rule. This has lead to a proliferation of results, with four different findings appearing commonly in this literature:

1. The FOMC's inflation response has changed over time, becoming more aggressive against inflation after the appointment of Paul Volker in 1979 (e.g. Clarida et al. (2000)).

2. The FOMC's output response has changed over time, becoming less responsive to changes in the output gap after the 1970's (e.g. Orphanides (2004)).

3. The FOMC's inflation and output response have both changed over time, typically according to the patterns identified in (1) and (2) (e.g. Gonzalez-Astudillo (2018)).

4. The coefficients of the FOMC's interest rate rule have not changed over time, but the variance has (e.g. Sims and Zha (2006)).

Addressing the nature and timing of structural change in FOMC behavior is of importance to both academics and policymakers. As has been demonstrated in many different types of DSGE models, coefficients in the interest rate rule help to determine inflation volatility and persistence, as well as short-term output growth rates and volatility. If monetary policy in the United States has changed, it is crucial that we document how, as it will eventually allow us to attribute changes in economic performance to changes in policy. This is especially important in light of the claim made by Taylor (2013) that a weak inflation response returned in the mid-2000s and engendered the housing bubble.

While the early papers in this literature such as Clarida et al. (2000) and Orphanides (2004) used split-sample regressions to model possible coefficient change, later work has used models explicitly designed to detect structural change. Some authors have used a time-varying parameter model, allowing the interest rate rule parameters to slowly change over time.[1] In more recent work, it has been common to model changes in the FOMC's interest rate rule as a Markov-Switching model.[2] However, the exact nature and timing of coefficient change is disputed, with many of these studies using different specifications of the interest rate rule, different time periods, or different information sets. In addition, in some studies such as Bennani et al. (2018) and Gonzalez-Astudillo (2018), the authors have searched for parsimonious models by "pre-testing" — performing frequentist testing of coefficient constancy in an unrestricted Markov-switching model, and then re-estimating a smaller model with a subset of coefficients restricted to be identical across regimes. Other studies, such as Murray et al. (2015), have simply estimated the unrestricted Markov-switching model.

In this paper, I introduce a novel Bayesian econometric model that is able to probabilistically determine the specification of the FOMC's interest rate rule in the presence of a two-state Markov-Switching process. This Markov-Switching Stochastic Search Variable Selection (MS-SSVS) model nests both a constant coefficient model, consistent with the find-

[1]Examples include Boivin (2006), Kim and Nelson (2006), and Primiceri (2005). A somewhat different approach that also allows for the FOMC to shift its policy horizon is considered in Lee et al. (2015).

[2]Examples include Soques (2019), Bennani et al. (2018), Gonzalez-Astudillo (2018), Alba and Wang (2017), Murray et al. (2015), Castelnuovo et al. (2014), Bianchi (2013), and Davig and Leeper (2011)

ings of Sims and Zha (2006), and a full Markov-Switching model, consistent with Murray et al.'s (2015) as special cases. In addition, the MS-SSVS model can probabilistically restrict a subset of coefficients to remain constant across regimes, nesting FOMC behavior like that identified by Orphanides (2004), Bennani et al. (2018), or Gonzalez-Astudillo (2018) where only a subset of coefficients change over time. The MS-SSVS model can also restrict coefficients in either one or both regimes to be zero, so that a variable may be completely excluded from the regression in either one regime or in both regimes. In short, for each coefficient in each regime, there are three possibilities: (1) the coefficient is restricted to zero; (2) the coefficient is restricted to be the same as the coefficient in the other regime; (3) the coefficient is freely estimated independently of the coefficient in the other regime. The restrictions and coefficients are estimated at the same time in a unified model, which avoids pre-testing and makes the estimated coefficients easily interpretable.[3] In addition, MS-SSVS averages over the uncertainty associated with model choice, so rather than only estimating a final "best" model, inference is performed by weighing estimates across models according to their posterior probabilities. This can expose features such as bi-modal coefficient densities, which would likely be missed under traditional estimation methods such as maximum likelihood estimation of a single Markov-switching model.

The MS-SSVS model builds on the work of George and McCulloch (1993) and George et al. (2008), who developed Stochastic Search Variable Selection (SSVS) in order to perform variable selection in linear regression models and linear Vector Autoregressions (VARs). SSVS has some differences with competing methodologies such as Bayesian Model Averaging (BMA) that make it especially attractive in a Markov-Switching environment. One major advantage of SSVS is that it is not necessary to directly compute or approximate the marginal likelihood, which is a computationally intensive task in Markov-Switching models. Instead, the uncertainty associated with variable inclusion and variable switching is nested within a unified hierarchical model.

[3]See Giles and Giles (1993) for an overview of some of the problems associated with pre-testing.

Through Monte-Carlo exercises with simulated data, I show that the newly developed MS-SSVS model is able to correctly identify when coefficients should be freely estimated and when they should be restricted — this is true both when the actual coefficients are zero, and when the coefficients are the same across regimes. As expected, the MS-SSVS model is better able to identify coefficient restrictions as the signal to noise ratio increases. This increase in signal is modeled in two ways, either as a reduction in error volatility or as an increase in sample size. The MS-SSVS model is particularly adept at identifying linear cross-regime restrictions with a high degree of accuracy, even as the amount of noise increases.

When I apply this new methodology to Federal Funds Rate data from 1970-2008, I find three things. First, I find that the point-estimate of the FOMC's inflation response coefficient is nearly identical across both regimes. Second, I find evidence for two distinct regimes in coefficient values, with the unemployment gap response coefficient differing across the regimes. I find that there was a relatively strong unemployment response coefficient in the mid 1970s, the late 1980s and early 1990s, and between 2004-2006. This strong unemployment response corresponds to a heightened probability that the inflation response coefficient was relatively low, as the inflation response coefficient has a bimodal distribution in this regime. Finally, similar to Sims and Zha (2006), I find strong evidence of changes in the volatility of the Federal Funds rate.

2 Previous Studies of the FOMC Interest Rate Rule

In an attempt to distill previous findings on changes in the FOMC's interest rate rule, I have grouped papers according to their main conclusion as it relates to monetary policy. Broadly speaking, these papers typically fall into one of four categories:

1. The FOMC's inflation response has changed over time, but the output response has not.

2. The FOMC's output response has changed over time, but the inflation response has

not.

3. The FOMC's inflation and output response have both changed over time.

4. None of the coefficients of the FOMC's interest rate rule have changed over time, but the variance has.

The findings of some recent and well-known studies of the FOMC's interest rate rule are highlighted in Table 1. This table is not meant to be an exhaustive list of all papers in this strand of literature, but instead it is meant to provide enough examples that the reader can get a sense of the types of differences seen across studies. In some of these studies, the findings are more subtle; for instance, Boivin (2006) finds that the point-estimates of both the output and inflation response coefficients change substantially over time, but that the confidence interval at any given point in time is very wide. In ambiguous cases like this I did my best to summarize the author's views about their findings.

Table 1: Evidence of Parameter Change in Previous Studies

Study	Smoothing	Output	Inflation	Variance
Clarida et al. (2000)	–	–	Y	–
Orphanides (2004)	–	Y	–	–
Boivin (2006)	–	Y	Y	–
Sims and Zha (2006)	–	–	–	Y
Murray et al. (2015)	Y	Y	Y	Y
Alba and Wang (2017)	–	Y	Y	Y
Gonzalez-Astudillo (2018)	–	Y	Y	Y
Bennani et al. (2018)	Y	Y	–	–

Note: Smoothing, output, inflation, and variance indicate whether coefficients in each respective variable were found to change across regimes. "Y" means that there was evidence that it switched, "–" means that the variable was excluded from consideration, was not allowed to switch, or was not found to switch.

From Table 1, we see that conclusions about changes in the policy rule are mixed. For example, two papers find strong evidence for changes in the interest rate smoothing parameter, while other papers find weak or no evidence for changes in this parameter. In many papers, the exact variables used (e.g. output gap vs. unemployment gap) affect the results, even when the methodology is the same. Furthermore, even when two papers find that the same parameter switched, the qualitative conclusions could differ. For instance, Clarida et al. (2000) and Alba and Wang (2017) both find evidence that the FOMC's inflation response has changed over time. However Clarida et al. (2000) find that the Taylor principle was likely not satisfied during the 1970's, while Alba and Wang (2017) find that the Taylor principle was always satisfied.

3 Data and Model Outline

For simplicity, I will first proceed as if the fully unrestricted Markov-switching model is being estimated. In that case, the model set-up is similar to others in the recent literature:

$$\begin{aligned} i_t &= \mu_{s_t} + \rho_{s_t} i_{t-1} + \phi_{\pi_{s_t}}(\pi_t^e - \pi_t^T) + \phi_{u_{s_t}} u_t^e + \phi_{\Delta u_{s_t}} \Delta u_t^e + \sigma_t \varepsilon_t \\ s_t &\in \{0, 1\} \\ P(s_t = j | s_{t-1} = i) &= p_{ij} \\ \varepsilon_t &\sim \mathcal{N}(0, 1) \end{aligned}$$

where i_t is the nominal Federal Funds rate in period t, π_t^e is the average three-quarter-ahead expected inflation rate at time t, u_t^e is the average three-quarter-ahead expected unemployment gap at time t, and Δu_t^e is the total expected change in the unemployment rate between time t and time $t+3$, and p_{ij} is the transition probability between state i and state j. The data is recorded quarterly and is gathered from the FOMC Greenbook, and for inflation I use the relevant GNP or GDP Deflator inflation. I use "meeting-based timing", discussed in Check (2019), in which the Federal Funds rate is averaged between meeting dates. In the baseline estimation, I assume a constant natural rate of unemployment and I use the methodology of Chan et al. (2013) to estimate the inflation target, using real-time GDP Deflator data. I restrict the sample from late 1969 through the end of 2007. The start date coincides to when the FOMC began regularly reporting three-quarter-ahead forecasts, and the end date is chosen to avoid the zero-lower-bound.

Finally, I assume that the volatility of the error term follows a random walk (i.e. the model exhibits "stochastic volatility"). Let $\sigma_t = \exp(\frac{h_t}{2})$. Then:

$$h_t = h_{t-1} + v_t$$
$$v_t \sim \mathcal{N}(0, Q)$$

There are a few differences between the model above and some of the other models in the literature. First, while most papers have included interest rate smoothing ($\rho_{s_t} i_{t-1}$), not all have.[4] Second, most papers in the literature use a measure of the output gap rather than the unemployment rate. However, as argued in Kozicki and Tinsley (2006) and Kozicki and Tinsley (2009), the historical narrative evidence from the 1970's and 1980's is much more consistent with a FOMC that responded to unemployment rather than output. In addition, Kozicki and Tinsley (2009) note that the Federal Funds rate target may depend on *changes* in real activity rather than (or in addition to) *gaps* in real activity. Therefore I include measures of both the real-time unemployment gap and the change in the unemployment rate. Since the MS-SSVS model performs variable selection, the model should be able to identify if one or more of these variables does not belong in the policy rule.

In addition to differences in data and timing, I believe that the use of a stochastic volatility process is unique in the single-equation interest rate rule literature.[5] Since the main focus of this study is to identify possible change in coefficient values, it is important to allow the variance to evolve separately. If the regimes covered both the coefficients as well as the variance, relatively large changes in the variance could drive the estimated regimes, which would then mainly identify periods of high and low volatility, rather than identifying changes in coefficients. While modeling the variance as a separate regime switching process

[4]For example, Alba and Wang (2017) does not allow for interest rate smoothing.
[5]Primiceri (2005) allowed for stochastic volatility in a VAR set-up.

would also break this link, identifying variance regimes is not the main focus of the study, and stochastic volatility is robust to different types of parameter change.[6]

4 Full Econometric Model

The model I introduce is based on a Markov-Switching model with switching in coefficients:

$$
\begin{aligned}
&y_t = X_t \beta_{St} + \varepsilon_t \\
&\varepsilon_t \sim \mathcal{N}(0, \sigma^2) \\
&S_t \in \{0, 1\}
\end{aligned}
$$

The regime, S_t follows a first order Markov process:

$$
\begin{aligned}
Pr(S_t = j | S_{t-1} = i) &= p_{ij} \\
i, j &\in \{0, 1\}
\end{aligned}
$$

The model as detailed above has been well studied, and there exist well known frequentist and Bayesian procedures to estimate it. These methods are described in Hamilton (1989), Kim and Nelson (1999), and Frühwirth-Schnatter (2006), among others. While these techniques make it feasible to estimate the model, model comparison remains relatively cumbersome. The estimation process itself can be time consuming, making the estimation of more than a handful of models potentially infeasible. This problem is only exacerbated when performing Bayesian Model Averaging, as this requires estimation of the marginal likelihood of each model — this is an extra, complicated, and time consuming step that needs to be undertaken

[6] While policy changes by the FOMC may impact the volatility of a monetary policy rule, it is also likely impacted by structural changes in the banking system and policy implementation. Changes in volatility are therefore less likely to be well described by a two-state Markov-Switching process.

after estimation of the model.[7] Finally, the total number of models to consider grows more rapidly in this class of models than in linear models. In linear models, there are only two possibilities for each regressor — either it belongs in the model or it does not. However, in a Markov-switching model with two states, there are four possibilities for each regressor: (1) it does not belong in either regime, (2) it belongs in each regime, and is the true effect is distinct under each regime, (3) it belongs in each regime but the true effect is the same regardless of regime, (4) the variable belongs in only one of the two regimes. This only increases the burden of model comparison, as the number of models to consider expands more rapidly when the possibility of switching is properly accounted for.

Despite the numerical difficulties with marginal likelihood calculations, model comparison in this class of models remains important. In this paper, I build an econometric model that allows for the four possibilities elicited above. It does so in a computationally feasible manner by utilizing a hierarchical prior to nest all of these possibilities within a single model. To do this, I build on Stochastic Search Variable Selection (SSVS), which was developed in George and McCulloch (1993) and further studied and implemented in George et al. (2008) and Koop and Korobilis (2010). In SSVS, a mixture of Normals prior is used on the regression coefficients to allow researchers to place prior weight on the possibility that the coefficient may be exactly equal to zero (i.e. that the variable does not belong in the model). This prior impacts the model likelihood and allows the data to inform whether the variable belongs in the model. Since variable selection is built into the model likelihood, only one model needs to be estimated. This framework is therefore simpler and more efficient than performing Bayesian model averaging by estimating hundreds, thousands, or more models and then comparing them based on their marginal likelihoods.

To help elicit the mathematical details of this model, let $\beta^k = [\beta_i^k \ \beta_j^k]$ be a vector that contains coefficient k in state i and state j. Under the MS-SSVS model the prior $p\big(\beta^k\big)$ is

[7] Additionally, some authors such as Kruschke (2015) argue against using the marginal likelihood for model comparison due to its high sensitivity to parameter priors. This problem was originally pointed out in Jeffreys (1939) and received more attention after Lindley (1957) named it a "paradox."

assumed to be a mixture Normals with five components:

$$\beta^k \sim \gamma_1^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \tau_0 \begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}\right) + \gamma_2^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \tau_1 \begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}\right) + \gamma_3^k \mathcal{N}\left(\begin{bmatrix}\hat{\beta}_{OLS}^k\\\hat{\beta}_{OLS}^k\end{bmatrix}, \tau_1 \begin{bmatrix}1 & 1-\epsilon\\1-\epsilon & 1\end{bmatrix}\right)$$
$$+ \gamma_4^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}\tau_1 & 0\\0 & \tau_0\end{bmatrix}\right) + \gamma_5^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}\tau_0 & 0\\0 & \tau_1\end{bmatrix}\right)$$
$$\gamma^k = (\gamma_1^k, \gamma_2^k, \gamma_3^k, \gamma_4^k, \gamma_5^k) \in \{(1,0,0,0,0),(0,1,0,0,0),(0,0,1,0,0),(0,0,0,1,0),(0,0,0,0,1)\}$$
$$\tau_0 = 0.1 \times \sqrt{\text{Var}\left(\hat{\beta}_{OLS}^k\right)}$$
$$\tau_1 = 15 \times \sqrt{\text{Var}\left(\hat{\beta}_{OLS}^k\right)}$$

where $\hat{\beta}_{OLS}^k$ is the OLS estimate of β^k under the assumption of no regime-switching, $\text{Var}\left(\hat{\beta}_{OLS}^k\right)$ is the OLS estimate of the variance, and τ_0, τ_1, and ϵ are parameters chosen by the researcher that control the variances of each distribution in the prior.[8]

The mixture distribution described above represents five distinct possibilities:

1. The coefficient is restricted to be approximately zero in each state, i.e. the variable is excluded in both states.

2. The coefficient estimates are freely estimated, independently of each other, i.e. the variable is included in each state, and the effect in each state is different.

3. The coefficient estimates are freely estimated, but identical, i.e. the variable is included in each state, and the effect in each state is the same.

4. The coefficient in state 0 is freely estimated, but the coefficient in state 1 is restricted

[8]The choices of τ_0, τ_1, and ϵ are similar in spirit to Bayesian linear regression, where the researcher typically chooses the variance of the prior distribution for the regression coefficients. Note that SSVS models are not technically fully Bayesian since they rely on the data to inform the priors. This can be avoided by running OLS on a pre-sample of data and using those pre-sample estimates instead. In addition, in SSVS, model restrictions can only enforced approximately. For more details, see the Technical Appendix.

to be near zero, i.e. the variable is excluded from state 1.

5. The coefficient in state 1 is freely estimated, but the coefficient in state 0 is restricted to be near zero, i.e. the variable is excluded from state 0.

I assume that each indicator vector γ^k comes from the following prior distribution:

$$\gamma_j^k = \begin{cases} (1,0,0,0,0) & \text{with probability } p_1 \\ (0,1,0,0,0) & \text{with probability } p_2 \\ (0,0,1,0,0) & \text{with probability } p_3 \\ (0,0,0,1,0) & \text{with probability } p_4 \\ (0,0,0,0,1) & \text{with probability } p_5 \end{cases}$$

$$\sum_{i=1}^{5} p_i = 1$$

$$0 < p_i < 1 \; \forall \; i \in \{1,2,3,4,5\}$$

where each prior mixture probability, p_i, is a fixed constant set by the researcher.[9]

5 Monte-Carlo Analysis

To test the power of this procedure to identify coefficient restrictions, I perform a Monte-Carlo analysis. I consider two data generating processes (DGPs): one in which data is generated from a process that has several types of restrictions, and another that corresponds to linear regression. The first Monte-Carlo exercise is analogous to a situation in which some parameters in the Taylor rule switch, some remain constant across regimes, and others are zero in both regimes. The second Monte-Carlo exercise is analogous to a situation in which the Taylor rule variables are correctly specified, but that there is actually no switching — the parameter values are all identical across regimes.

[9] For more details on the exact estimation procedure, please refer to the Technical Appendix.

In both cases, the true model can be written as:

$$y_t = X_t \beta_{St} + \varepsilon_t$$
$$\varepsilon_t \sim \mathcal{N}(0, \sigma^2)$$
$$Pr(S_t = j | S_{t-1} = i) = p_{ij}$$

I assume that there is an intercept and two independent regressors with mean zero:

$$X_t = \begin{bmatrix} 1 & X_{1,t} & X_{2,t} \end{bmatrix}$$
$$\begin{bmatrix} X_{1,t} \\ X_{2,t} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \right)$$

Since there are three columns in X_t, $K = 3$. Recall that $\beta^k = [\beta_0^k \ \beta_1^k]'$ for $k \in \{1, \cdots, K\}$.

5.1 DGP 1: Restricted Coefficients and Markov-Switching

For the first Monte-Carlo exercise, I assume that the intercept is different in each regime, the coefficient on the first regressor is identical across regimes, and the coefficient on the second regressor is zero in both regimes. This set-up is analogous to a situation where, for example, the real interest rate switches, the FOMC has the same response to inflation (x_1) across regimes, but does not respond to the output gap (x_2) in either regime. As defined above, we have:

$$\beta^1 = \begin{bmatrix} 1.0 \\ -0.5 \end{bmatrix} \quad \beta^2 = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix} \quad \beta^3 = \begin{bmatrix} 0.0 \\ 0.0 \end{bmatrix}$$

Finally, the transition probabilities for each regime are given by:

$$P = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} = \begin{bmatrix} 0.8 & 0.2 \\ 0.3 & 0.7 \end{bmatrix}$$

which implies that the system will spend roughly 60% of time periods in regime 0 and 40% of time periods in regime 1.

5.1.1 Priors

In this exercise, I assume that the variance term is constant. Therefore, I set an inverse-Gamma prior on the variance term. The priors are as follows:

Table 2: Monte-Carlo Priors

Parameter	Prior Mean	Prior S.D.
c_0	0.1	0.0
c_1	15	0.0
τ_0^k	$c_0\sqrt{\text{Var}\left(\hat{\beta}_{OLS}^k\right)}$	0.0
τ_1^k	$c_1\sqrt{\text{Var}\left(\hat{\beta}_{OLS}^k\right)}$	0.0
p_1	0.25	0.0
p_2	1/12	0.0
p_3	0.50	0.0
p_4	1/12	0.0
p_5	1/12	0.0
p_{00}	0.8	0.16
p_{11}	0.8	0.16
σ^2	improper	infinite

5.1.2 Results

In the Monte-Carlo exercise, I vary both the number of observations, T, and the standard deviation of the error term, σ. I consider $T \in \{50, 100, 150, 200, 250\}$ and $\sigma \in \{0.1, 0.5, 1.0\}$. Models estimated on DGPs with large T and small σ are most able to pick out the correct restrictions, as these models have the largest sample size and smallest variance.

In the tables below, I present the average accuracy of identification of the correct restriction by the estimation procedure. This number has been averaged over the results across 200 separate data generation and estimation procedures. For example, for the first column of table (2), I set $\sigma = 0.1$ and $T = 50$. I then generate 200 data sets and run the estimation procedure on each. For each data set, I calculate the percentage of draws in which the appropriate restriction was chosen, and I average this percentage across all 200 data sets. For the estimation procedure, I use 15,000 burn-in draws and 20,000 posterior draws. For ease of notation, let the regression coefficients on the intercept term be rewritten as $\begin{bmatrix} \beta_0^1 \\ \beta_1^1 \end{bmatrix} \equiv \begin{bmatrix} \mu_0 \\ \mu_1 \end{bmatrix}$. Finally, in the tables below, I have abused notation in the first two rows. For $\mu_0 \neq 0$ and $\mu_1 \neq 0$, I mean that each is freely estimated, so that they are not restricted to be equal to zero and not restricted to be identical to each other.

Table 3: Low Variance ($\sigma = 0.1$)

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 \neq 0$	100%	100%	100%	100%	100%
$\mu_1 \neq 0$	100%	100%	100%	100%	100%
$\beta_0^1 = \beta_1^1$	99.4%	99.7%	99.8%	99.9%	99.9%
$\beta_0^2 = 0$	94.4%	95.4%	95.8%	95.8%	95.9%
$\beta_1^2 = 0$	94.4%	95.4%	95.8%	95.8%	95.9%

Table 4: Medium Variance ($\sigma = 0.5$)

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 \neq 0$	75.3%	98.8%	100%	100%	100%
$\mu_1 \neq 0$	57.7%	95.2%	99.4%	100%	100%
$\beta_0^1 = \beta_1^1$	93.3%	97.8%	98.5%	98.9%	99.0%
$\beta_0^2 = 0$	79.5%	82.5%	83.9%	83.1%	82.7%
$\beta_1^2 = 0$	78.8%	81.8%	84.0%	83.3%	83.7%

Table 5: High Variance ($\sigma = 1.0$)

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 \neq 0$	31.5%	55.6%	71.0%	81.2%	90.3%
$\mu_1 \neq 0$	13.0%	25.7%	42.1%	52.6%	64.8%
$\beta_0^1 = \beta_1^1$	80.1%	88.1%	93.0%	94.2%	96.4%
$\beta_0^2 = 0$	74.3%	76.6%	78.2%	76.3%	76.6%
$\beta_1^2 = 0$	74.4%	75.9%	78.6%	76.3%	77.4%

Three things become apparent when looking at these tables. First, the MS-SSVS model is able to correctly identify all types of restrictions when the data has a high signal to noise ratio. This shows that the estimation procedure is well-behaved when the amount of noise in the data generating process is relatively small. Second, the model performs very well at detecting the linear regression restriction, $\beta_0^1 = \beta_1^1$, and fairly well at detection the zero-restriction for coefficients β^2. Third, as the noise increases, the model has a relatively more difficult time detecting that μ_1 is actually different than zero compared to μ_0. This is likely due to the fact that the absolute value of μ_1 is smaller than the absolute value of μ_1.

5.2 DGP 2: Linear Model with no switching

I repeat the same exercise as above, except I change the true value of the coefficients so that all regressors belong in the model, the effect is different than zero for all regressors, and the effect is the same in both regimes for all regressors. This corresponds to linear regression in which there is no misspecification — all the variables included in the model have a non-zero effect, and there are no excluded relevant variables. This set-up is analogous a situation in which the Taylor rule was correctly specified and there was no parameter change over time, similar to what was found by Sims and Zha (2006).

$$\beta^1 = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix} \quad \beta^2 = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix} \quad \beta^3 = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix}$$

I leave everything else, including the priors, unchanged and conduct the same analysis as above. I present the results in the tables below.

Table 6: Low Variance ($\sigma = 0.1$), Linear Regression

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 = \mu_1$	100%	100%	100%	100%	100%
$\beta_0^1 = \beta_1^1$	100%	100%	100%	100%	100%
$\beta_0^2 = \beta_1^2$	100%	100%	100%	100%	100%

Table 7: Medium Variance ($\sigma = 0.5$), Linear Regression

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 = \mu_1$	92.8%	96.5%	97.9%	99.4%	99.3%
$\beta_0^1 = \beta_1^1$	92.4%	96.0%	98.1%	99.0%	99.7%
$\beta_0^2 = \beta_1^2$	92.8%	96.6%	97.6%	99.2%	99.6%

Table 8: High Variance ($\sigma = 1.0$), Linear Regression

	$T = 50$	$T = 100$	$T = 150$	$T = 200$	$T = 250$
$\mu_0 = \mu_1$	81.8%	85.7%	87.5%	90.6%	91.8%
$\beta_0^1 = \beta_1^1$	84.1%	86.1%	89.3%	92.4%	93.4%
$\beta_0^2 = \beta_1^2$	84.2%	88.3%	90.3%	92.3%	93.3%

Under the linear regression DGP, the model performs remarkably well under all sample sizes and all variance sizes. This may be partially driven by the fact that there is a 50% prior probability placed on each of these coefficients being identical. It is important to note that while I have fixed the prior mixture probabilities to be identical for all sets of parameters, in general a researcher could relax this assumption, placing different prior mixture probabilities on each pair of coefficients.[10]

6 Application: Interest Rate Rules

Next, I apply the MS-SSVS model to the interest rate rule equation and data described in section three. Recall that I estimate a rule of the form:

$$i_t = \mu_{s_t} + \rho_{s_t} i_{t-1} + \phi_{\pi_{s_t}}(\pi_t^e - \pi_t^T) + \phi_{u_{s_t}} u_t^e + \phi_{\Delta u_{s_t}} \Delta u_t^e + \sigma_t \varepsilon_t$$

$$\varepsilon_t \sim \mathcal{N}(0,1)$$

$$\sigma_t = \exp\left(\frac{h_t}{2}\right)$$

$$h_t = h_{t-1} + v_t$$

$$v_t \sim N(0,Q)$$

$$s_t \in \{0,1\}$$

$$P(s_t = j | s_{t-1} = i) = p_{ij}$$

while also allowing for MS-SSVS coefficient restrictions on μ_{s_t}, ρ_{s_t}, $\phi_{\pi_{s_t}}$, $\phi_{u_{s_t}}$, and $\phi_{\Delta u_{s_t}}$.

I have three major findings: (1) I find evidence of distinct regimes in monetary policy, with the response to the unemployment gap changing substantially between regimes; (2) at the posterior mean, the FOMC's response to the inflation gap is similar across both regimes; (3) the volatility of the error term varies substantially across time and the timing does not appear to coincide with the estimated regimes.

Detailed findings for the regression parameters are presented in Tables 9, 10 and 11, as well as in Figure 1. In the last column of Tables 9 and 10, we see that there is substantial evidence that four of the five regression coefficients have entered the Taylor rule linearly,

[10] For example, if a researcher suspected that one pair of coefficients would be different in each regime, she could increase the prior probability on mixture 2 and reduce the prior probability on mixture 3.

remaining constant regardless of regime. There is only strong evidence of regime switching in the unemployment gap response coefficient, which has only a 9.4% posterior probability of remaining constant across regimes. This is further evidenced in Table 11, which shows that the posterior mean for all coefficients are very similar across regimes, with the exception of the unemployment gap response coefficient . Finally, this is visualized in Figure 1, which plots the entire estimated density of coefficients in each regime. While most densities lie largely on top of each other, the densities of the unemployment gap response coefficient are largely separated, indicating that it is distinct across regimes.

Table 9: Estimated Restrictions in the "Strong" Unemployment Response Regime

	Zero-Restriction	Freely Estimated	"Identical" Restriction
μ_1	0.0%	0.0%	100%
ρ_1	0.0%	0.0%	100%
ϕ_{π_1}	13.1%	4.5%	82.3%
ϕ_{u_1}	0.0%	89.9%	9.4%
$\phi_{\Delta u_1}$	0.3%	7.4%	89.9%

Table 10: Estimated Restrictions in the "Weak" Unemployment Response Regime

	Zero-Restriction	Freely Estimated	"Identical" Restriction
μ_0	0.0%	0.0%	100%
ρ_0	0.0%	0.0%	100%
$\phi_{\pi,0}$	4.0%	13.7%	82.3%
$\phi_{UN,0}$	67.2%	23.5%	9.4%
$\phi_{\Delta UN,0}$	4.5%	5.5%	89.9%

Table 11: Mean Coefficient Values in Each Regime

	"Weak" UN Regime	"Strong" UN Regime
μ	6.622	6.622
ρ	0.949	0.949
ϕ_π	0.245	0.218
ϕ_{UN}	-0.070	-0.331
$\phi_{\Delta UN}$	-0.185	-0.195

Note: Inflation, the unemployment gap, and the change in unemployment were standardized by subtracting the mean and dividing by the standard deviation before the model was estimated. The coefficients on these variables therefore represent the short-run response of a one-standard-deviation change in the respective regressor. The long-run response coefficients can be recovered by transforming the coefficients from each posterior draw.

Figure 1: Regression Parameters

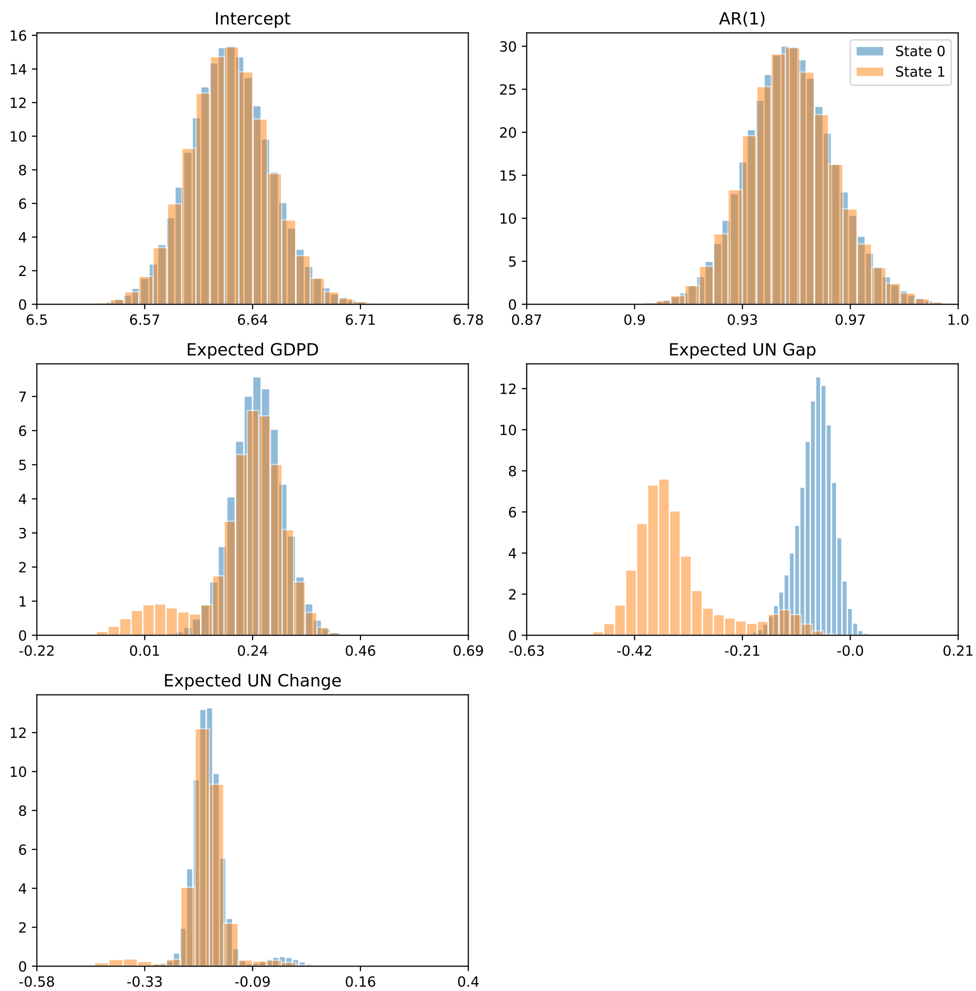

To get a better sense of the magnitude of the difference across regimes, Figure 2 plots the FOMC's expected response to a persistent one-percentage-point unemployment gap up to 40 meetings (roughly 4–5 years) in the future. Clearly, the responses differ across the regimes. In the strong unemployment response regime, if this unemployment gap persisted for 40 meetings, the FOMC would have decreased the Federal Funds rate by a total of roughly 4.2 percentage points, vs. only a 0.9 percentage point reduction in the weak unemployment response regime.

On the contrary, in Figure 3, we see that even in the long-run, the FOMC's response to expected inflation is roughly constant across regimes. After a one-percentage-point inflation gap lasting 40 meetings, the FOMC would likely increase the Federal Funds rate by between roughly 2.0 and 2.5 percentage-points under either regime. For all periods, the credible intervals from each regime overlap substantially, indicating that there is very weak evidence of any meaningful difference between the inflation response across regimes.

Figure 2: Response to 100 Basis Point Unemployment Gap

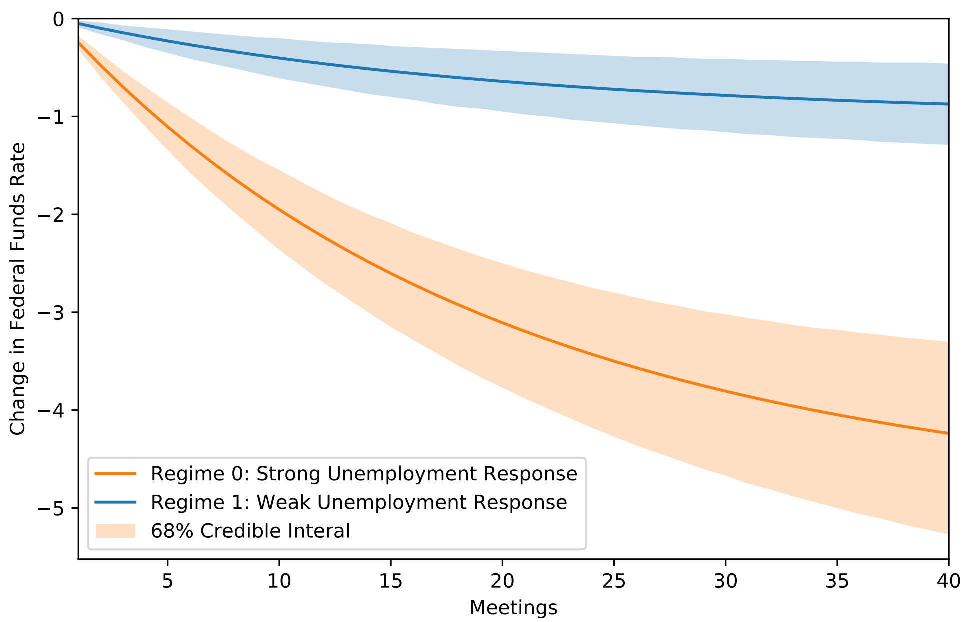

Despite the fact that the posterior mean for the inflation response is roughly equal across regimes, the inflation response coefficient has a bimodal density in the strong unemployment response regime, as can be seen in Figure 1. In this regime, there is roughly a 13% probability that the FOMC was responding only very weakly to the inflation gap. This finding would be missed if parameter constancy had been "pre-tested" in an unrestricted model, with the final estimated model only allowing for changes in the unemployment gap. The bimodal inflation response shows the importance of averaging over uncertainty with respect to parameter switching, and underlines a strength of the MS-SSVS approach taken in this

Figure 3: Response to 100 Basis Point Inflation Gap

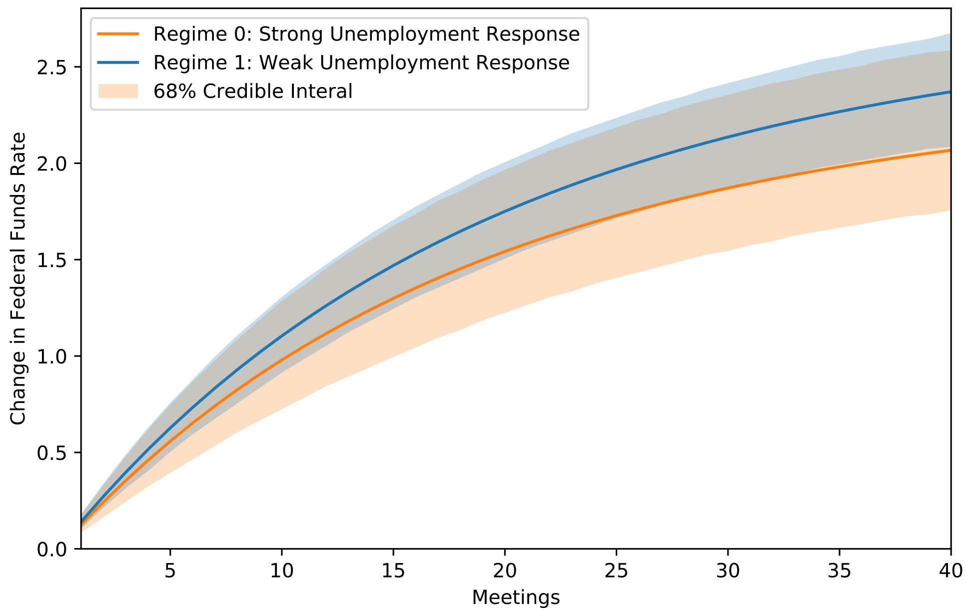

paper.

I also find strong evidence of change in the volatility of the interest rate rule. This change can be seen in Figure 4, with the standard deviation of the error in the interest rate rule peaking in the early 1980's at around 150 basis points and falling substantially since then, to below 25 basis points. The spike in volatility in the interest rate rule in the late 1970's and early 1980's is expected, as the Federal Reserve was on record as targeting the money supply rather than the interest rate rule during this time period. Once they returned to targeting the Federal Funds rate, the volatility steadily declined, as the interest rate rule matched actual FOMC behavior much more accurately.

Along with the estimated parameter values, the timing of regimes and regime changes are of interest. These are presented in Figure 5, which plots the posterior mean of being in the strong unemployment gap response at each meeting date. Also displayed in Figure 5 are the dates of changes in Fed chair, and for each Fed chair, the average probability that the FOMC was in the strong unemployment gap response regime during his time as chairman.

Figure 4: Stochastic Volatility

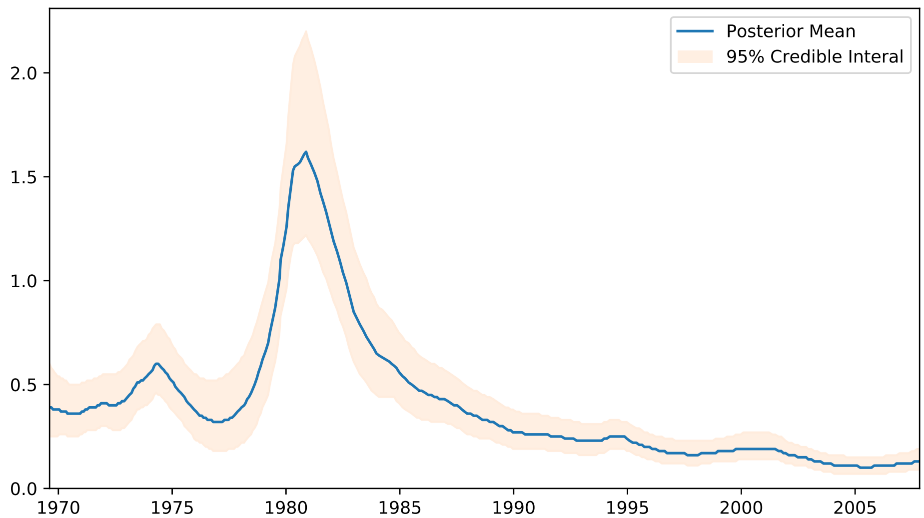

While there are certainly swings in estimated regime probabilities within each chairman's tenure, especially under Burns and Greenspan, the transition period between chairmen also seems to correspond to regime change. For instance, the probability of being in the strong unemployment response regime begins declining from over 80% very late in Burns' tenure, continues sliding throughout Miller's brief chairmanship, and bottoms out below 20% as Volcker takes control in 1979. Interestingly, despite the fact that I am using a very different proxy of the output gap than Orphanides (2004), my findings are largely consistent with his — the FOMC under Burns was more likely to respond strongly to the output gap, with the FOMC under Volcker much less likely to do so. Under Burns, the average probability of being in the strong unemployment response regime was roughly 60%, while under Volcker it was roughly 20%. Because the inflation response coefficient is largely similar across regimes, this means that the FOMC under Volcker put *relatively* more weight on its inflation response, which is consistent with narrative historical evidence (e.g. Kaya et al. (2019)).

Figure 5: Probability of Strong Unemployment Response Regime

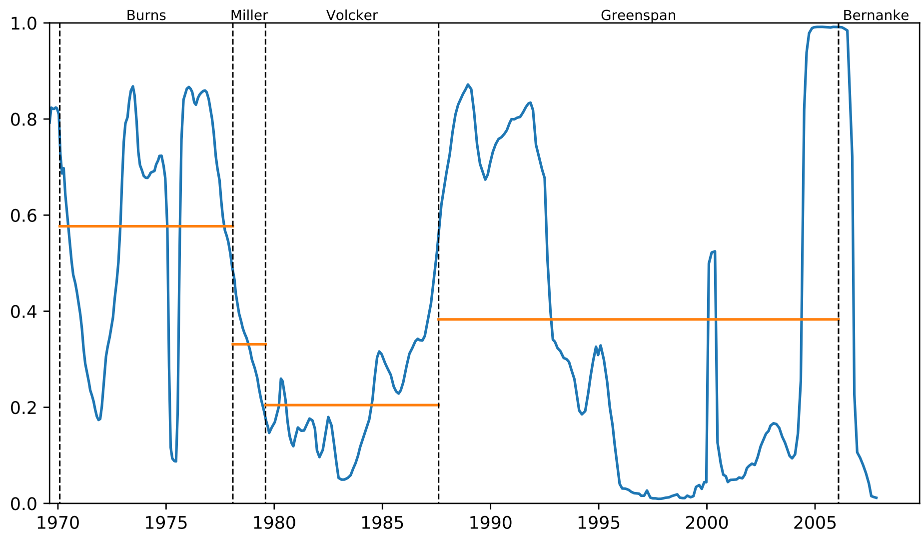

Note: The solid blue line represents the probability of the "strong unemployment response." The vertical dashed lines denote changes in Fed chairman, and the orange horizontal lines represent the average probability of being in the "strong unemployment response" regime under the different chairmen.

While my results are consistent with some of the previous literature, such as Orphanides (2004) and Sims and Zha (2006), they stand in contrast to some others. For example, in a recent application of a Markov-Switching model to interest rate rules in Murray et al. (2015) the authors find that the inflation response in one regime was much lower than in the other regime, and there were periods in which it was highly probable that the FOMC failed to satisfy the "Taylor Principle." However, they find that this weak inflationary response occurred during the Volcker years, 1979-1985 (among other times). This seems highly counterfactual, and goes against historical accounts, narrative evidence from the FOMC meetings during this time period, and other statistical evidence such as Clarida et al. (2000), all of which attribute the fall in inflation after 1980 to the *strong* inflation response during Volcker's tenure. My findings are more consistent with this latter strand of literature — although I do not find evidence of a direct change in the inflation response coefficient, I do find an increase in the *relative* inflation response during the Volcker years, as I find the FOMC under Volcker largely stopped responding to the unemployment gap.

7 Conclusion

Over the past 15 years, there has been considerable disagreement about the existence and nature of changes in the coefficients in the FOMC's interest rate rule. In an attempt to clarify the nature of these changes, I build a Markov-Switching model that can endogenously determine the existence of two types of restrictions: (1) zero-restrictions, in which a variable may be excluded from one or both of the regimes and (2) identity-restrictions, in which the regression coefficient on the same variable may be restricted to be identical across both regimes. The estimation procedure blends and extends the Gibbs samplers that were previously derived for estimation of Markov-Switching models and Stochastic Search Variable Selection models. I call this unified model an MS-SSVS model.

I find that the MS-SSVS model performs well at identifying true restrictions in a Monte-Carlo exercise using simulated data. In general, the MS-SSVS model performs best in data-sets that have a relatively small amount of noise. In these data sets, it is able to detect zero-restrictions, "identical" restrictions, and switching in the coefficients with high probability. The MS-SSVS is still able to identify these restrictions as the amount of noise grows, and it is able to detect linear cross-regime restrictions with a surprisingly high degree of accuracy in even the noisiest data sets.

When I apply this model to Federal Funds rate data I find three things. First, there is relatively little evidence that there have been economically meaningful shifts in inflation response over the period 1970-2007. Second, consistent with Orphanides (2004), I find substantial evidence that there have been changes in the unemployment gap coefficient over time. I find that the periods least likely to have had a strong response to the unemployment gap are the 1980s, and the period from roughly 1995-2004. The first period mostly corresponds to the chairmanship of Paul Volcker, and suggests that the FOMC focused relatively more on responding to changes in inflation under his leadership. This is consistent with much, but not all, of the previous literature. Finally, I find strong evidence that there have been changes in the volatility of interest rate rule. This adds to a relatively strong body of existing evidence; for example, Sims and Zha (2006), finds that models that allow for a change in variance substantially out-perform models with constant variance.

References

Alba, J. D. and Wang, P. (2017). Taylor Rule And Discretionary Regimes In The United States: Evidence From A k-State Markov Regime-Switching Model. *Macroeconomic Dynamics*, 21(03):817–833.

Bennani, H., Kranz, T., and Neuenkirch, M. (2018). Disagreement between FOMC Members and the Fed's Staff: New Insights Based on a Counterfactual Interest Rate. *Journal of Macroeconomics*, 58:139–53.

Bianchi, F. (2013). Regime switches, agents' beliefs, and post-world war ii u.s. macroeconomic dynamics. *The Review of Economic Studies*, 80(2):463–490.

Boivin, J. (2006). Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data. *Journal of Money, Credit and Banking*, 38(5):1149–1173.

Castelnuovo, E., Greco, L., and Raggi, D. (2014). Policy Rules, Regime Switches, And Trend Inflation: An Empirical Investigation For The United States. *Macroeconomic Dynamics*, 18(04):920–942.

Chan, J. C. C., Koop, G., and Potter, S. M. (2013). A New Model of Trend Inflation. *Journal of Business & Economic Statistics*, 31(1):94–106.

Check, A. (2019). Interest Rate Rules in Practice — A Taylor Rule or a Tailor-Made Rule? . *Unpublished Manuscript*.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory. *The Quarterly Journal of Economics*, 115(1):147–180.

Davig, T. and Leeper, E. M. (2011). Monetary-fiscal policy interactions and fiscal stimulus. *European Economic Review*, 55:211–27.

Frühwirth-Schnatter, S. (2006). *Finite Mixture and Markov Switching Models*, volume 1 of *Spring Series in Statistics*. Springer Science & Business Media.

George, E. I. and McCulloch, R. E. (1993). Variable Selection Via Gibbs Sampling. *Journal of the American Statistical Association*, 88(423):881–89.

George, E. I., Sun, D., and Ni, S. (2008). Bayesian stochastic search for VAR model restrictions. *Journal of Econometrics*, 142(1):553–580.

Giles, J. A. and Giles, D. E. A. (1993). Pre-test Estimation and Testing in Econometrics: Recent Developments. *Journal of Economic Surveys*, 7(2):145–197.

Gonzalez-Astudillo, M. (2018). Indentifying the Stance of Monetary POlicy at the Zero Lower Bound: A Markov-Switching Estimation Exploiting Monetary-Fiscal Policy Interdependence. *Journal of Monetary, Credit, and Banking*, 50(1):115–154.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–84.

Jeffreys, H. (1939). *Theory of Probability*. Oxford University Press, 4 edition.

Kaya, A., Golub, S., Kuperberg, M., and Lin, F. (2019). The federal reserve's dual mandate and the inflation-unemployment tradeoff. *Contemporary Economic Policy*, Forthcoming.

Kim, C.-J. and Nelson, C. R. (1999). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, volume 1 of *MIT Press Books*. The MIT Press.

Kim, C.-J. and Nelson, C. R. (2006). Estimation of a forward-looking monetary policy rule: A time-varying parameter model using ex post data. *Journal of Monetary Economics*, 53(8):1949–1966.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–93.

Koop, G. and Korobilis, D. (2010). Bayesian Multivariate Time Series Methods for Empirical Macroeconomics. *Foundations and Trends(R) in Econometrics*, 3(4):267–358.

Kozicki, S. and Tinsley, P. (2006). Minding the Gap: Central Bank Estimates of the Unemployment Natural Rate. *Computational Economics*, 27(2):295–327.

Kozicki, S. and Tinsley, P. (2009). Perhaps the 1970s FOMC did what it said it did. *Journal of Monetary Economics*, 56(6):842–855.

Kruschke, J. K. (2015). *Doing Bayesian Data Analysis*. Academic Press / Elsevier, 2 edition.

Lee, K., Morley, J., and Shields, K. (2015). The Meta Taylor Rule. *Journal of Money, Credit and Banking*, 47(1):73–98.

Lindley, D. V. (1957). A statistical paradox. *Biometrica*, 44(1-2):187–192.

Murray, C. J., Nikolsko-Rzhevskyy, A., and Papell, D. H. (2015). Markov Switching And The Taylor Principle. *Macroeconomic Dynamics*, 19(04):913–930.

Orphanides, A. (2004). Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches. *Journal of Money, Credit and Banking*, 36(2):151–75.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, 72(3):821–852.

Romer, C. D. and Romer, D. H. (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review*, 94(4):1055–1084.

Sims, C. A. and Zha, T. (2006). Were There Regime Switches in U.S. Monetary Policy? *American Economic Review*, 96(1):54–81.

Soques, D. (2019). Timing and Signals of Monetary Regime Switching. *Unpublished Manuscript*.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39(1):195–214.

Taylor, J. B. (2013). A Review of Recent Monetary Policy. Economics Working Papers 13103, Hoover Institution, Stanford University.

Technical Appendix — More information about SSVS and the Estimation Procedure

Enforcing parameter restrictions in SSVS

In an SSVS model, the posterior probability of each restriction is proportional to the prior probability of the restriction times the prior density for β^k under that restriction evaluated at the posterior draw of β^k. Therefore, the "restrictions" are only enforced approximately. To see why this is necessary, consider a prior density for the zero-restriction in which $\tau_0 = 0$, so that the coefficient was literally restricted to equal zero. Unless this restriction was chosen, our estimate of β^k will almost surely never be exactly zero. Therefore, the posterior density under this restriction will always be zero, since $\beta^k \neq 0$; therefore this restriction will never be enforced. Our goal when choosing τ_0 is to choose a sensible value that will enforce this zero-restriction when appropriate, while keeping the estimate of β^k near zero in the event that this restriction is chosen. Our goal when choosing ϵ is similar. We want to choose a number small enough that the prior density under this restriction will be high when both values of β^k are approximately equal, but we need to be careful to not choose a value for ϵ that is so small that the restriction will never be enforced.

Using Data to Inform SSVS Priors

The prior for the regression coefficients is data dependent since $\widehat{Var(b_k)}$ depends on the dependent variable. Therefore, it does not adhere to the requirement, in a Bayesian approach, that the prior be independent of the observed dependent data. However, as discussed in George and McCulloch (1993), the zero-restriction region depends on the values for both τ_0 and τ_1. George and McCulloch (1993) find that using $\widehat{Var(b_k)}$ in the choices of τ_0 and τ_1 helps to ensure that this zero-restriction region lies over a sensible space so that the coefficients are restricted to be close to zero only where appropriate. This prior, although not technically valid due to its dependence on the observed data, remains popular in the literature, as evidenced by its use in Koop and Korobilis (2010).[11]

Estimation Procedure

I set independent priors across the hierarchical parameters:

$$p(p_{00}, p_{11}, \sigma^2, \tau_0, \tau_1, p_1, p_2, p_3, p_4, p_5, \epsilon) = \\ p(p_{00})p(p_{11})p(\sigma^2)p(\tau_0)p(\tau_1)p(p_1)p(p_2)p(p_3)p(p_4)p(p_5)p(\epsilon)$$

I assume that the prior parameters $\tau_0, \tau_1, \epsilon, p_1, p_2, p_3, p_4, p_5$ are each set by the researcher, i.e. their prior is a point-mass at a particular value. This is a common assumption in

[11] Priors of this form are sometimes called "empirical Bayes" methods. One argument for their use, although not mathematically rigorous, is that the goal of empirics is to discover features of the data. Priors of this form should be considered if they can be shown to be well behaved and able to uncover features of the data, even if they do not technically adhere to proper Bayesian theory.

the SSVS literature. The parameters τ_0, τ_1, and ϵ control the variance of the each prior mixture distribution. The probabilities, p_1, p_2, p_3, p_4, and p_5, control the weights for each prior distribution.

For the other three hyper-parameters, p_{00}, p_{11}, and σ^2, I set prior distributions:

$$\begin{aligned} p(p_{00}) &= \text{Beta}(a_0, b_0) \\ p(p_{11}) &= \text{Beta}(a_1, b_1) \\ p(\sigma^2) &= \text{InverseGamma}(\alpha_Q, \beta_Q) \end{aligned}$$

Drawing from the full posterior directly is intractable. Instead, I draw from each of the conditional posteriors. This is called the Gibbs sampler. Let $\beta = [\beta_0, \beta_1]'$, $P_s = [p_{00}\ p_{11}]'$, $\tau = [\tau_0, \tau_1]'$, $P_\gamma = [p_1, p_2, p_3, p_4, p_5]'$, $\Gamma = \gamma^K$. The process is as follows:

1. Sample the indicators for the mixture of normals prior each variable:

$$\begin{aligned} p(\Gamma^{(z)}|Y, \beta^{(z-1)}, P_s^{(z-1)}, \sigma^{2,(z-1)}, S_T^{(z-1)}, \tau, P_\gamma) &= p(\Gamma^{(z)}|Y, \beta^{(z-1)}, S_T^{(z-1)}, \tau, P_\gamma) \\ p(\Gamma^{(z)}|Y, \beta^{(z-1)}, S_T^{(z-1)}, \tau, P_\gamma) &= \text{ Categorical} \\ \Gamma_k^{(z)} &= \text{ Categorical} \begin{pmatrix} \frac{p_1 f(N(0,\Sigma_1)|\beta^k)}{\sum_{i=1}^5 p_i f(N(0,\Sigma_i)|\beta^k)} \\ \frac{p_2 f(N(0,\Sigma_2)|\beta^k)}{\sum_{i=1}^5 p_i f(N(0,\Sigma_i)|\beta^k)} \\ \frac{p_3 f(N(0,\Sigma_3)|\beta^k)}{\sum_{i=1}^5 p_i f(N(0,\Sigma_i)|\beta^k)} \\ \frac{p_4 f(N(0,\Sigma_4)|\beta^k)}{\sum_{i=1}^5 p_i f(N(0,\Sigma_i)|\beta^k)} \\ \frac{p_5 f(N(0,\Sigma_5)|\beta^k)}{\sum_{i=1}^5 p_i f(N(0,\Sigma_i)|\beta^k)} \end{pmatrix} \end{aligned}$$

This procedure is based on George and McCulloch (1993). In the current paper, their procedure is slightly modified because I have a mixture of five normal distributions rather than two. Once the prior mixture distributions are selected, form the prior variance for β as:

$$D = \begin{bmatrix} \Sigma^{k=1} & 0 & \cdots & 0 & 0 \\ 0 & \Sigma^{k=2} & 0 & \cdots & 0 \\ 0 & 0 & \ddots & \cdots & 0 \\ 0 & 0 & 0 & \cdots & \Sigma^{k=K} \end{bmatrix}$$

D is block diagonal, with the Cholesky decomposition of the two-by-two mixture variance for each pair of coefficients, k, Σ^k along the diagonals, with zeros everywhere else.

2. Sample the regression coefficients:

$$p(\beta^{(z)}|Y, \Gamma^{(z)}, P_s^{(z-1)}, \sigma^{2,(z-1)}, S_T^{(z-1)}, \tau, P_\gamma) = p(\beta^{(z)}|Y, \Gamma^{(z)}, \sigma^{2,(z-1)}, S_T^{(z-1)}, \tau)$$
$$p(\beta^{(z)}|Y, \Gamma^{(z)}, \sigma^{2,(z-1)}, S_T^{(z-1)}, \tau) \sim \mathcal{N}\left(\hat{\beta}, V\right)$$
$$V = ((DRD')^{-1} + X'X)^{-1}$$
$$\hat{\beta} = VX'Y$$

where R is a prior correlation matrix. In this application, R is set to the identity matrix.

3. Sample the variance of the regression error:

$$p(\sigma^{2,(z)}|Y, \Gamma^{(z)}, \beta^{(z)}, P_s^{(z-1)}, S_T^{(z-1)}, \tau, P_\gamma) = p(\sigma^{2,(z)}|Y, \beta^{(z)}, S_T^{(z-1)})$$
$$p(\sigma^{2,(z)}|Y, \beta^{(z)}, S_T^{(z-1)}) = \text{ Inverse Gamma}$$
$$\sigma^{2,(z)} \sim \text{ IG}\left(a_Q + \frac{T}{2}, \beta_Q + \frac{SSE}{2}\right)$$

where T is the sample size and $SSE = (Y - X\beta)'(Y - X\beta)$. This step is replaced by sampling stochastic volatility via $Kimet\ al.$ (1998) in the interest rate rule application.

4. Sample the Markov State indicators:

$$p(S_T^{(z)}|Y, \Gamma^{(z)}, \beta^{(z)}, \sigma^{2,(z)}, , P_s^{(z-1)}, \tau, P_\gamma) = p(S_T^{(z)}|Y, \beta^{(z)}, \sigma^{2,(z)}, P_s^{(z-1)})$$

using the procedure described in Kim and Nelson (1999).

5. Sample the Markov transition probabilities:

$$p(P_s^{(z)}|Y, \Gamma^{(z)}, \beta^{(z)}, \sigma^{2,(z)}, S_T^{(z)}, \tau, P_\gamma) = p(P_s^{(z)}|Y, S_T^{(z)})$$
$$p(P_s^{(z)}|Y, S_T^{(z)}) = \text{ Beta}$$
$$P_s^{ii,(z)} = \text{ Beta}(a_i + N_{ii}, b_i + N_{ij})$$

where N_{ij} is the number of times that the regime transitioned from regime i to regime j in $S_T^{(z)}$.

Overview of the prior distribution for the regression parameters

Helicopter Tour of Prior for β^k

Recall that the prior for $\beta^k = [\beta_0^k \; \beta_1^k]'$ is given by a mixture of five Normal distributions:

$$\beta^k \sim \gamma_1^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \tau_0 \begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}\right) + \gamma_2^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \tau_1 \begin{bmatrix}1 & 0\\0 & 1\end{bmatrix}\right) + \gamma_3^k \mathcal{N}\left(\begin{bmatrix}\hat{\beta}_{\text{OLS}}^k\\\hat{\beta}_{\text{OLS}}^k\end{bmatrix}, \tau_1 \begin{bmatrix}1 & 1-\epsilon\\1-\epsilon & 1\end{bmatrix}\right)$$
$$+ \gamma_4^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}\tau_1 & 0\\0 & \tau_0\end{bmatrix}\right) + \gamma_5^k \mathcal{N}\left(\begin{bmatrix}0\\0\end{bmatrix}, \begin{bmatrix}\tau_0 & 0\\0 & \tau_1\end{bmatrix}\right)$$

In my application, I choose:

$$Pr(\gamma_1^k = 1) = 0.5$$
$$Pr(\gamma_2^k = 1) = \frac{.25}{3}$$
$$Pr(\gamma_3^k = 1) = 0.25$$
$$Pr(\gamma_4^k = 1) = \frac{.25}{3}$$
$$Pr(\gamma_5^k = 1) = \frac{.25}{3}$$

In addition, I choose:

$$\tau_0^k = c_0\sqrt{\widehat{var(\beta^k)}}$$
$$\tau_1^k = c_1\sqrt{\widehat{var(\beta^k)}}$$
$$c_0 = 0.1$$
$$c_1 = 15.0$$

where $\widehat{var(\beta^k)}$ is the OLS estimate of the variance of β^k under a no regime switching assumption. These priors are similar to ones suggested in George and McCulloch (1993) and Koop and Korobilis (2010). Finally, for the case of parameters restricted to be equal across regimes, I set $\epsilon = 1.0 - 0.99999$.

In figures (1)-(3), I plot the prior probability density function of β_0 and β_1, implicitly assuming that $\hat{\beta}_{\text{OLS}}^k = 0$. This prior density function has some striking features. It is strongly peaked near $\beta_0 = 0$ and $\beta_1 = 0$, so there is a relatively high prior probability that both coefficients are restricted to zero. If the estimated coefficients land in the orange region of figure 2 (or the yellow region of figure 3), it is almost a certainty that the priors for β_0 and β_1 will be centered on zero with a very tight prior variance. Additionally, there are three other regions which receive relatively large prior mass: both regions where one of the coefficients is restricted to be near zero, and the diagonal region representing coefficients that are (roughly) identical under each regime. Outside of these four relatively narrow but sharply peaked regions, the Normal distribution with the highest probability density function corresponds to both regimes being freely estimated.

Figure 6: Prior Probability Density Function for Different Values of β_0 and β_1

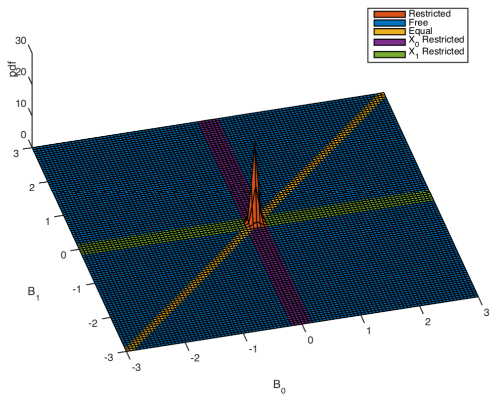

Figure 7: Prior Probability Density Function for Different Values of β_0 and β_1: View from Above

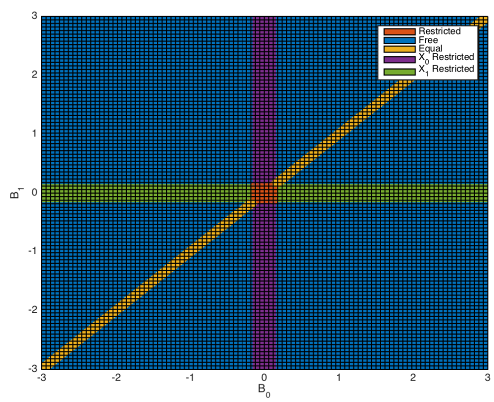

Figure 8: Contour Plot of the Prior Probability Density Function for Different Values of β_0 and β_1

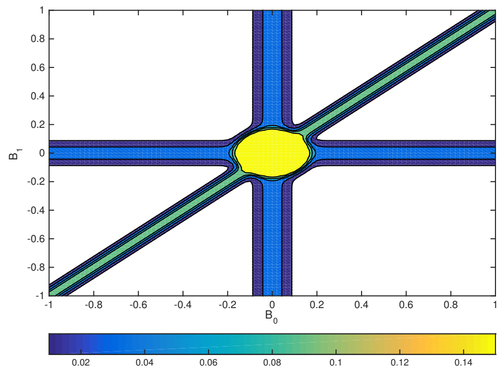

Convergence Diagnostics

Below, I present evidence that the estimator converges to a unique stationary distribution. I first present running mean plots throughout the samples that are discarded. If the sampler is converging to a stationary distribution, then the means of all of the parameters of the model should converge to their means in the stationary distribution. If it is not, then these means will be trending up, down, or bouncing around. Next, I present the autocorrelation functions for the parameters of the model. These functions show the correlation between the draw of the parameter at one iteration and the draw of the same parameter t iterations later. If the sampler is well-behaved, then the autocorrelation functions should fall towards zero as the number of iterations increases. A simple rule-of-thumb is that the number of discarded "burn-in" draws should be at least ten times larger than the maximum number of iterations that it takes the autocorrelation of any parameter to drop to zero.

Running Mean Plots

Figure 9: Regression Parameters

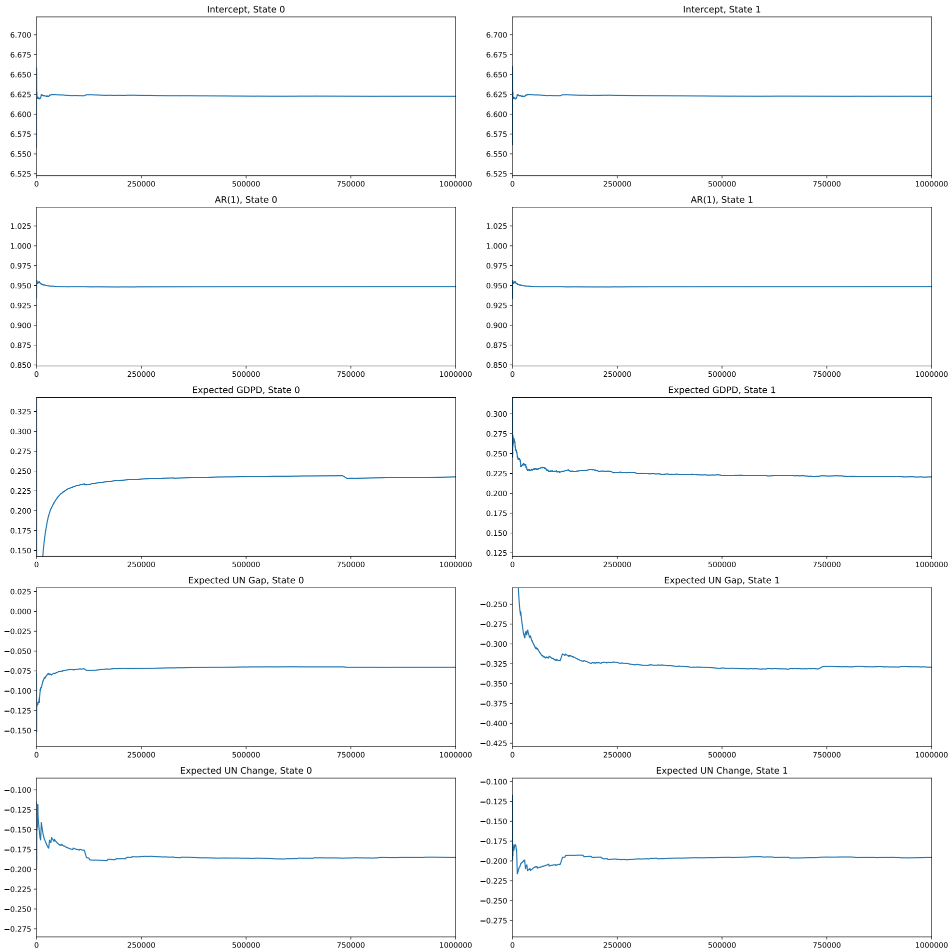

Figure 10: Regime Probabilities

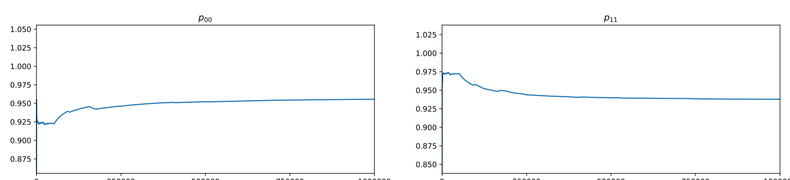

Autocorrelation Functions

Figure 11: Regression Parameters

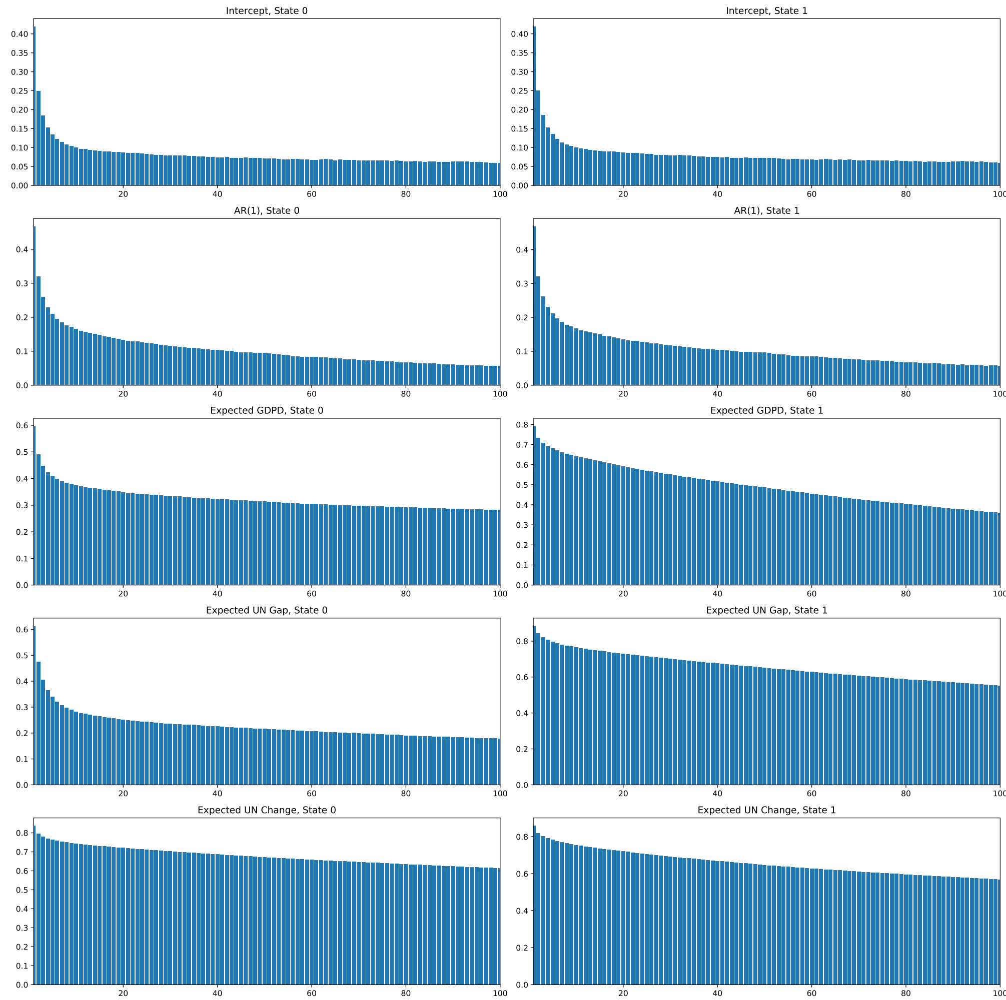

Figure 12: Regime Probabilities

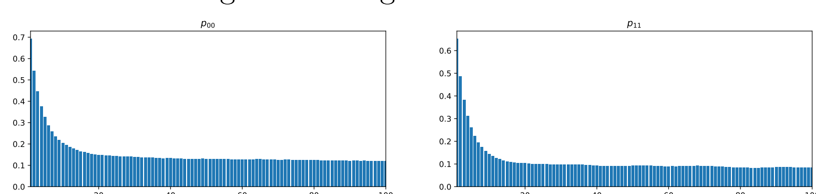

Both of these metrics suggest that the sampler is well-behaved, although the ACF function is declining fairly slowly. The running mean plots become flat towards the end of the discarded draws, suggesting that the sampler has converged to a stationary distribution. The autocorrelation plots for a couple of the regression parameters is still fairly high at 100 draws, suggesting that more than 1,000 burn-in draws are needed. I perform 1,000,000 burn-in draws in an abundance of caution. I keep the next 1,000,000 draws and use them to form posterior inference.