How Optimal is US Monetary Policy?

Xiaoshan Chen[a,], Tatiana Kirsanova[b], Campbell Leith[b,*†]

[a] University of Durham; [b] University of Glasgow

Abstract

Using a small-scale microfounded DSGE model with Markov switching in shock variances and policy parameters, we show that the data-preferred description of US monetary policy is a time-consistent targeting rule with a marked increase in conservatism after the 1970s. However, the Fed lost its conservatism temporarily in the aftermath of the 1987 stock market crash, and again following the 2000 dot-com crash and has not subsequently regained it. The high inflation of the 1970s would have been avoided had the Fed been able to commit, even without the appointment of Paul Volcker or the reduction in shock volatilities.

Keywords: Bayesian Estimation, Interest Rate Rules, Optimal Monetary Policy, Great Moderation

JEL classification: E58, E32, C11, C51, C52, C54

*Corresponding author. Address: ASBS, University of Glasgow, University Avenue, Glasgow, G12 0QL, UK; phone: +441413303702; email: campbell.leith@glasgow.ac.uk

1. Introduction

It is common practice to adopt a simple Taylor (1999)-type instrument rule to describe monetary policy when estimating DSGE models.[1] This practice, however, is inconsistent with the claim of practitioners, that no central bank actually adopts such instrument rules, but rather prefer to set clear objectives and follow 'elaborate decision making-processes, in which huge amounts of data are elaborated and processed' (Svensson, 2003, pp. 428) in attempting to achieve those objectives. By specifying policy objectives the central bank adopts – using Svensson's terminology – a general targeting rule. This general targeting rule is then developed into a specific targeting rule by maximizing these objectives subject to the equations describing the decentralized equilibrium of the economy. The targeting rule that emerges is dependent on the degree of commitment the central bank possesses. What that degree of commitment is in practice, and whether or not we can develop data-coherent targeting rules, remains an open question, with the literature containing mixed results.

This paper considers various descriptions of policy – both instrument and targeting rules – and takes seriously the notion that policy making and the shocks hitting the US economy have been subject to shifts over the years. Doing so gives a far clearer indication as to which policy description best fits the data. This in turn has significant policy implications both in terms of designing monetary policy institutions and contributing to the debate on the source of the 'Great Moderation'.

The estimation demonstrates that the US monetary policy is best described by a time consistent targeting rule, labelled as discretion throughout the paper. This policy strongly dominates conventional simple instrument rules, as well as alternative forms of targeting rule with higher degrees of precommitment. This implies that during the post-WWII period the US Fed has not been making any credible policy commitments, either by following the Ramsey plan or following a simple instrument rule. The data also suggest that there have

---

[1]See e.g. Smets and Wouters, 2003.

been changes in the Fed's degree of anti-inflation conservatism and in the volatilities of shocks hitting the economy. Ignoring these changes reduces the models' ability to fit the data and distorts the ranking of models.

The results imply that the inferences about shock processes, habit persistence and inflation indexation change significantly across different policy specifications. Under targeting rules, relative to instrument rules, we find that there is a shift in emphasis away from preference shocks towards cost-push shocks in driving the US business cycle. Under discretion this greater emphasis on cost-push shocks is not implausible, but is dramatic under commitment. Differences in the estimates of structural parameters under targeting rules further reflect the need to generate a meaningful policy trade-off, resulting in the degree of habits and inflation indexation being higher under commitment. In contrast, discretion tends to downplay the extent of habits to prevent implausibly aggressive policy responses to the associated externality.

The findings contribute to the literature in two respects. First, they add to the small but growing research on the empirical validity of targeting rules. While there are papers which estimate models under commitment (Adolfson et al., 2011 and Ilbas, 2010), discretion (Dennis, 2004) and an intermediate case of limited commitment, also known as quasi-commitment, as in Debortoli and Lakdawala (2016), very few compare the empirical relevance across these different targeting rules and with simple instrument rules.[2] In contrast to these papers, we consider a wide range of policy descriptions, and allow for potential regime switches in the monetary policy specification. Doing so explains how different policies interact with inferences about shock processes and structural parameters of the model.

Second, the analysis presented extends the 'good luck' and 'good policy' debate to the framework of targeting rules. There is a large literature on the 'Great Moderation' based

---

[2] Adolfson et al. (2011) find that commitment is preferred to a simple instrument rule using Swedish data. Givens (2012) and Le Roux and Kirsanova (2013) suggest that discretion is marginally preferred to commitment in the US and UK respectively.

on simple instrument rules, which finds that breaks in estimated policy rules (Lubik and Schorfheide, 2005, and Boivin and Giannoni, 2006), the implicit inflation target (Favero and Rovelli, 2003, Erceg and Levin, 2003 and Ireland, 2007) and/or the volatility of the underlying shock processes (Sims and Zha, 2006) help to explain the evolution of inflation dynamics across time. Given these findings, we allow for variation in the policy-maker's degree of anti-inflation conservatism, and for switches in the variance of the shock processes, when estimating different forms of targeting rule. The best-fitting model implies that US monetary policy is best described as being conducted under discretion, with an increase in central bank conservatism following the Volcker disinflation period, which is found to have occurred in 1982. More importantly, it identifies additional periods of policy change: the Fed relaxed policy temporarily in the aftermath of the 1987 stock market crash, and also lost conservatism following the 2000 dot-com crash, which it has never regained.

Finally, the counterfactual analysis using the best-fitting model suggests that the 'Great Moderation' in output and inflation volatility is due to both a reduction in shock variances and an increase in central bank anti-inflation conservatism. Decomposing the relative contribution of both effects implies that the far greater part of the 'Great Moderation' stems from the reduction in shock volatilities. More importantly, the counterfactuals show that inflation would never have breached 2% in the 1970s had the policy maker had access to a commitment technology. The potential gains from moving from discretion to commitment are substantial and dominate the gains from increasing central bank conservatism. Ensuring that the US Fed has access to commitment technologies and that they act to use such mechanisms is the 'good policy' that policymakers should focus on.

The plan of the paper is as follows. Section 2 outlines our model and the policy maker's preferences. The various descriptions of policy are discussed in Section 3. Section 4 considers data, priors and identification of the model, before presenting the estimation results in Section 5. Section 6 contrasts the results to those of Debortoli and Lakdawala (2016).

Section 7 then undertakes various counterfactual simulation exercises which facilitate an exploration of both the sources and welfare consequences of the 'Great Moderation', and also an assessment of the potential benefits of further improvements in the conduct of monetary policy. Section 8 concludes.

2.   The Model

The economy is comprised of households, a monopolistically competitive production sector, and the government. Full details of the underlying microfoundations of the model are given in the online Appendix A and only the linearized model is presented here.[3]

The household's optimization gives rise to the labor supply decision

$$\sigma \hat{X}_t + \varphi(\hat{y}_t - \hat{z}_t) = \hat{w}_t - \hat{\mu}_t, \tag{1}$$

and consumption Euler equation

$$\hat{X}_t = \mathbb{E}_t \hat{X}_{t+1} - \frac{1}{\sigma}\left(\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \mathbb{E}_t \hat{z}_{t+1}\right) - \hat{\xi}_t + \mathbb{E}_t \hat{\xi}_{t+1}, \tag{2}$$

where $\hat{X}_t$ is habits-adjusted consumption

$$\hat{X}_t = (1 - \theta)^{-1}(\hat{y}_t - \theta \hat{y}_{t-1}), \tag{3}$$

and $\hat{y}_t$ denotes output, $\hat{w}_t$ is real wages, $\hat{\pi}_t$ is inflation and $\hat{R}_t$ is the nominal interest rate. Here σ is the inverse of the intertemporal elasticity of substitution, φ is the inverse of the Frisch elasticity and θ is the habit persistence parameter. The process $\hat{\mu}_t = \tau \hat{\tau}_t / (1 - \tau)$ represents fluctuations in the labor income tax rate which serves as a cost-push shock, $\hat{z}_t$ is an innovation to non-stationary technology process which serves as a technology shock and $\hat{\xi}_t$ is a preference shock.

---

[3] An on-line Appendix contains information on the microfoundations of the model, solution algorithms, estimation and identification tests.

The firms' optimization decisions, in presence of both price and inflation inertia, give rise to a hybrid New Keynesian Phillips curve

$$\hat{\pi}_t = \chi_f \beta E_t \hat{\pi}_{t+1} + \chi_b \hat{\pi}_{t-1} + \kappa_c \hat{w}_t, \tag{4}$$

where the reduced form parameters are $\chi_f = \alpha/\Phi$, $\chi_b = \zeta/\Phi$, $\kappa_c = (1-\alpha)(1-\zeta)(1-\alpha\beta)/\Phi$, with $\Phi = \alpha(1+\beta\zeta) + (1-\alpha)\zeta$, where $1-\alpha$ is the Calvo (1983) probability of price change, β is the households' discount factor and ζ is the proportion of firms setting prices who follow a backward-looking rule of thumb, rather than setting prices optimally.

Hatted variables indicate that they have been linearized relative to their steady-states. The stationarity of the model's steady state is achieved by scaling by a non-stationary technology process discussed in Appendix A. The technology, cost-push and preference shocks follow AR(1) processes:

$$
\begin{aligned}
\hat{z}_t &= \rho^z \hat{z}_{t-1} + \sigma_z \varepsilon_t^z, & \varepsilon_t^z &\sim N(0,1), \\
\hat{\mu}_t &= \rho^\mu \hat{\mu}_{t-1} + \sigma_\mu \varepsilon_t^\mu, & \varepsilon_t^\mu &\sim N(0,1), \\
\hat{\xi}_t &= \rho^\xi \hat{\xi}_{t-1} + \sigma_\xi \varepsilon_t^\xi, & \varepsilon_t^\xi &\sim N(0,1).
\end{aligned}
$$

The model is then closed with one of the instrument or targeting rules considered in Section 3. The Fed's targeting rule can be inferred from their objectives.

In the empirical analysis it is assumed that the Fed's objective function takes the micro-founded form, although the coefficients on the quadratic terms are freely estimated. Specifically, the empirical loss function can be written as

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\omega_1 \left(\hat{X}_t + \hat{\xi}_t \right)^2 + \omega_2 \left(\hat{y}_t - \frac{\sigma}{\varphi} \hat{\xi}_t \right)^2 + \omega_\pi \hat{\pi}_t^2 + \omega_3 \left(\hat{\pi}_t - \hat{\pi}_{t-1} \right)^2 \right), \tag{5}$$

see Appendix B for its microfoundations. This allows us to flexibly capture Svensson's (2003) notion of a general targeting rule by allowing the central bank to define the relative importance of welfare-relevant terms. Strictly speaking, it should not be interpreted as a welfare function unless the estimated coefficients coincide with the microfounded weights.

Given that much of the literature on estimated instrument rules finds that there have been significant changes in the conduct of policy over time, targeting rules derived under an assumption of unchanging policy maker preferences may be too stylized to capture such changes. Therefore, the relative weight on inflation, ω_π, is allowed to be subject to regime switching between 1 and a value lower than 1 to capture policy regimes with lower conservatism. The estimation can therefore assess whether or not the Fed's attitudes to inflation targeting have varied over time. For example, has monetary policy been more conservative since the Volcker disinflation? Moreover, accounting for independent regime switching in the variances of shocks, σ_z, σ_μ, and σ_ζ helps to assess whether the lower interest rates observed during 2001-2007 were due to economic conditions, or the result of the Fed putting less emphasis on inflation targeting relative to its other objectives.

3. Policy

The four basic forms of policy considered are a simple instrument rule and three types of targeting rule: discretion, commitment and the intermediate case labelled 'quasi-commitment'. Across these alternative policies, the estimation permits changes in inflation conservatism by allowing Markov switching in instrument rule parameters, as well as in the relative weight given to inflation in the policy objective underpinning targeting rules, as detailed in this section.

3.1. Instrument Rules

The instrument rule is a generalized Taylor rule which, following An and Schorfheide (2007), is specified as

$$\widehat{R}_t = \rho^R \widehat{R}_{t-1} + (1 - \rho^R)[\psi_1 \hat{\pi}_t + \psi_2(\Delta \hat{y}_t + \hat{z}_t)] + \varepsilon_t^R,$$

where the Fed adjusts interest rates in response to movements in inflation and deviations of output growth from trend.[4]

Within the framework of a generalized Taylor rule, potential changes in US monetary policy are accounted for by allowing for either changes in the Fed's inflation target or rule parameters. In the former case, following Schorfheide (2005), the measure of excess inflation in the Taylor rule, $\hat{\pi}_t$, removes the inflation target from the data, where that target follows a two-state Markov-switching process. In the latter case, when the policy changes are described as shifts in rule parameters (ρ^R, ψ_1, ψ_2) between two regimes, the procedure developed by Farmer et al. (2011) is applied to solve the model.[5]

3.2. Targeting Rules

When implementing targeting rules, the central bank selects interest rates to minimize loss function (5) subject to the structural equations describing private sector behavior, equations (1)-(4), and the evolution of shocks. The targeting rules considered include the standard cases of discretion and full commitment, which are the two polar cases of how well the central bank can manage the expectations of the private sector. Under commitment the policy maker can make credible promises about the setting of the policy instrument in future periods, while under discretion they re-optimize and are expected to re-optimize in each period. This implies that under commitment there is a history-dependence in policy making arising from these past commitments, which is absent under discretion. The empirical implementation of commitment assumes that the targeting rule has been in place for a prolonged period, such that policy is considered to be timeless.[6]

---

[4]Rules of this form have not only been found to be empirically useful, but, when suitably parameterized, can often mimic optimal policy, see, for example, Schmitt-Grohe and Uribe (2007). Moreover, by allowing for an additional policy shock in the interest rate rule relative to the cases of optimal policy, we are further supporting the simple rule's ability to fit the data. As we shall see, despite this, discretionary policy is 'strongly' preferred by the data.

[5]The details of the solution algorithm are provided in Appendix C.

[6]See, for example, Svensson and Woodford (2005). To economize on terminology this is referred simply as commitment hereafter.

The remaining form of targeting rule is quasi-commitment, as developed in Schaumburg and Tambalotti (2007) and Debortoli and Nunes (2010). The policymaker may deviate from commitment-based plans with a fixed exogenous probability, known by all agents. The current policy maker forms a commitment plan to be followed until randomly ejected, with a given probability, from office. At which point a new policy maker will be appointed, and a new plan formulated until that policy maker is, in turn, removed. Therefore, the central bank can neither completely control the expectations of the private sector, nor perfectly coordinate the actions of all future policy makers. This implies that, in contrast to the cases of discretion and commitment, in each period there is a policy surprise resulting from the fact that expectations are formed as a probability-weighted average of policy with and without reneging, while actual policy will either renege or not. Such policy surprises imply that outcomes under quasi-commitment are not a probability-weighted average of those under discretion and commitment.

The procedure described by Svensson and Williams (2007) is used to solve for the equilibrium dynamics under discretion and commitment with Markov-switching in objectives.[7] In addition, this solution method is modified to incorporate the case of quasi-commitment, as Schaumburg and Tambalotti (2007) and Debortoli and Nunes (2010) do not allow for Markov switching in objectives. Appendix C presents the new algorithm.

4.   Data, Priors and Identification

The empirical analysis uses US data on output growth, inflation, and nominal interest rates from 1961Q1 up to 2008Q3, just before nominal interest rates were reduced to their

---

[7]The Svensson and Williams (2007) algorithm implies that although policy makers can anticipate any changes in their objectives, they do not attempt to tie the hands of their future selves by altering today's policy plan as part of a strategic game, instead they set today's policy cooperatively with their future selves. We consider that this algorithm is in line with the conduct of US Fed policy as there may be some evolution in the consensus surrounding the objectives of monetary policy. However, in other policy making environments, where interest rate decisions are made by partisan politicians who may alternate in office, this would be less defensible and the approach of Debortoli and Nunes (2010) would be applicable.

effective lower bound of 0.5% and the first round of quantitative easing was implemented. The data used in the estimation are plotted in Figure 3, alongside various counterfactual simulation results which will be discussed below. The estimation strategy is standard and is described in Appendix D.

The priors are presented in Table 1. These are set to be broadly consistent with the literature on the estimation of New Keynesian models, in particular for the structural model parameters we follow Smets and Wouters (2003). For the Markov-switching instrument rule parameters, in line with Bianchi (2013), the priors for the response to output growth and the smoothing term are set to be symmetric across regimes, while asymmetric priors are chosen for the response to inflation.[8] For targeting rules, the relative weights (i.e. $\omega_1, \omega_2,$ and ω_3) on the objective function are assumed to be distributed following beta distributions and ω_π is allowed to switch between 1 and a value lower than 1, where the beta distribution is used for the latter with a mean of 0.5. The prior for the probability of reneging on past promises under quasi-commitment policy, υ, follows Debortoli and Lakdawala (2016) with a uniform prior on the interval [0,1]. The parameters, γ^Q, π^A and r^A represent the values of output growth, inflation and interest rates, respectively, when the economy is in its steady state. The prior means of γ^Q, π^A and r^A are set to be broadly consistent with their data averages during this pre-sample period from 1950Q1 to 1960Q4. Parameter π^A is interpreted as an inflation target, and it is assumed to be constant for all models except the instrument rule model with Markov-switching inflation target, where the priors for π^A are set in line with Schorfheide (2005). The average real interest rate, r^A, determines the discount factor, $\beta = \left(1 + r^A/400\right)^{-1}$.

[Table 1 around here]

Finally, it is important to note that all model parameters are identifiable. To demonstrate

---

[8]This way of setting priors for the switching parameters is also discussed by Davig and Doh (2014), as a means of introducing a natural ordering of regime-dependent parameters in order to avoid the potential risk of 'label switching', as noted in Hamilton, Waggoner, and Zha (2007).

this, the identification tests of Komunjer and Ng (2011a) and Koop et al. (2013) were applied to the models which feature both policy and volatility switches. In all cases model parameters are identified, see details in Appendix E. This is in contrast to the identification of parameters in larger models, (see the application of these tests to the Smets and Wouters model in Iskrev (2010), Caglar et al. (2011) and Komunjer and Ng (2011b), respectively) which is one reason why we prefer to work with a simpler model.

5. Results

This section presents the results of the estimation. It begins by identifying which description of policy best fits the data. It then discusses the implications of this for inferences about structural parameters of the economy, which shocks drive the business cycle in the US, and whether the Fed's preferences have changed over time.

5.1. Policy, Structural Parameters and Shocks

The posterior means and the 90% confidence intervals are presented in Table 2 where each column corresponds to an alternative policy description, and these columns are ordered according to the log marginal likelihood values calculated using Geweke (1999) and Sims et al. (2008), respectively.

[Table 2 around here]

The first column of results in Table 2 is for the best-fitting model, which is discretionary policy. Following Kass and Raftery (1995) the evidence in favor of discretion relative to instrument rules with switches in rule parameters is identified as 'strong', and relative to commitment as 'decisive'. The probability of reneging on policy promises under the quasi-commitment policy is $v = 0.29$, which implies that the commitment plan is expected to last for just 10 months. These estimates suggest that the discretionary form of targeting rule

best fits the data, and there is no evidence of any commitment behavior on the part of the Fed.

The estimates obtained under the conventional instrument rule are broadly in line with other studies: an intertemporal elasticity of substitution, $\sigma = 2.9$, a measure of price stickiness, $\alpha = 0.77$, implying that price contracts typically last for one year; a relatively modest degree of price indexation, $\zeta = 0.09$, a sizeable estimate of the degree of habits, $\theta = 0.83$ and an inverse Frisch labor supply elasticity of $\varphi = 2.4$. Moving to the case of discretion, these deep parameter estimates remain largely the same, except that there is a significant decline in the degree of habits in the model, which falls to $\theta = 0.39$, and a modest increase in the degree of indexation in price setting to $\zeta = 0.16$. The quasi-commitment policy delivers similar values for these parameters. However, with a further increase in the degree of precommitment to the case of strict commitment, the degree of indexation rises to $\zeta = 0.26$, while the extent of habit persistence increases to a level closer to that observed under instrument rules, $\theta = 0.69$.

These differences in the estimated structural parameters across targeting rules reflect the need to ensure the policy maker faces a meaningful trade-off. In the benchmark New Keynesian model it is only cost-push shocks which present a trade-off between output and inflation stabilization for the policy maker. All other shocks would result in policy responses which perfectly stabilize inflation. Introducing a habits externality breaks this 'divine coincidence' and implies other shocks will matter to the policy maker. Therefore, in order to explain the observed volatility in inflation, the estimation under commitment retains the degree of habits relative to instrument rules. This increases the ability of shocks, other than the cost-push shock, to generate inflation volatility. In such an environment the degree of inflation indexation is also likely to affect these policy trade-offs.

The case of discretion is more subtle. The inability to commit to a small but sustained response to shocks implies that in the presence of the habits externality the policy maker

will react aggressively to such shocks, see Leith, Moldovan, and Rossi (2012). This would imply higher interest rate volatility than is observed in the data. Therefore, the estimation downplays the extent of habits under discretion, relative to commitment.

In addition to variations in the degree of habits and inflation indexation across the estimates obtained under targeting rules, the balance between different shocks also changes. Again, this reflects the need to generate meaningful policy trade-offs in order to explain the inflation volatility observed in the data. Therefore, we see a reduction in both the persistence and standard deviation of preference shocks under targeting rules relative to instrument rules. At the same time, the persistence and standard deviation of cost push shocks increase, dramatically so in the case of commitment. However, it is important to note that under discretion the unconditional variance of this shock is not dissimilar to those found in other studies employing instrument rules as their description of policy.[9]

To summarize, relative to conventional instrument rules, our preferred targeting rule adjusts structural and shock parameter estimates to create a meaningful trade-off for policy when explaining macroeconomic volatility. This includes a shift from preference to cost push shocks in explaining the US business cycle.

5.2. Inflation Conservatism

The results suggest that the Fed's stance on inflation targeting has varied over the sample period. Taking into account potential switches in shock volatilities, for each policy specification, the estimation identifies two distinct inflation targeting regimes with a different degree of conservatism. We label them 'more' and 'less' conservative regimes, depending on the size of the weight on inflation, ω_π, under targeting rules. Under all targeting rules, ω_π is

---

[9]It should be noted that the cost-push shock enters the Phillips curve with the reduced form coefficient κ_c, which lies in the range 0.036-0.065 across our estimates. Calculating the unconditional variance of the normalized cost-push process $\kappa_c \hat{\mu}_t$ for discretion implies that the variance of 0.002 and 0.017 in low and high volatility regimes, respectively, is lower than that estimated by Smets and Wouters, 2003 for a single volatility regime (0.0217). For the case of quasi-commitment the corresponding numbers are 0.0012 and 0.014. However, commitment requires substantial increases in the unconditional variance of the cost push shock to 0.16 and 0.65 for the low and high volatility regime, respectively.

more than halved in the less conservative regime from the default level of one in the more conservative regime.

As for instrument rules with either Markov-switching rule parameters or inflation targets, a 'less' conservative inflation regime can be also identified by observing a reduction in the size of the coefficient on excess inflation, ψ_1, or an increase in inflation target, π^A, respectively. In the former case, although policy satisfies Taylor principle across both regimes, ψ_1 falls from 2.124 to 1.219, while for the latter case, π^A rises from 3.34% to 4.33%.

We now explore when these less conservative inflation regimes were estimated to have occurred. Figure 1 plots the smoothed probabilities of being in the less conservative targeting regime, as well as being in the high volatility regime. In the case of quasi-commitment, the plot also shows the probability that the policy maker has reneged on previous commitments.

[Figure 1 around here]

The best-fitting model, discretion, provides more information than the instrument rule-based models on the conduct of monetary policy over recent years, as the smoothed probabilities show. The estimation finds the relaxation of monetary policy in the 1970s that is well documented in the existing literature following Clarida et al. (1998). However, unlike the vast majority of the literature our estimates date the Volcker disinflation as occurring in 1982 rather than 1979.[10] Additionally, the smoothed probabilities from this model also suggest that policy was relaxed briefly following the stock market crash of October 1987. More interestingly, a prolonged reduction in the Fed's weight on the inflation target is identified as occurring at the time of dot-com crash and persisting all the way through to the financial crisis. Such a pattern is not so apparent in the instrument rule-based models. Similarly, the less conservative policy episodes are largely confined to the mid to late 1970s under commitment. Quasi-commitment utilizes two mechanisms to capture a relaxation in the Fed's

---

[10]More recent papers also find that the date of the Volcker disinflation is later than previously thought. See, for example, Bianchi (2013), Schorfheide (2005).

anti-inflation stance. Specifically, we may observe a reduction in the weight attached to inflation stabilization in the objective function (lost conservatism) or periods of reneging on past policy commitments. Relative to discretion, quasi-commitment relies on extensive periods of lost conservatism to such an extent that it is easier to define when conservatism was not lost under this policy description – briefly in the early 1980s and a few years prior to the bursting of the dot-com bubble – and even then, not fully.[11] In addition, the quasi-commitment estimates imply that the Fed reneged on policy commitments relatively frequently in the 1970s, and was showing signs of having possibly done so in the lead up to the financial crisis too.

5.3. The Importance of Switches in Policy and Volatilities

Turning to explore how important accounting for both the switches in policy and shock volatilities are for our estimated results, Table 3 re-estimates our models without allowing for either form of switching.[12] In this case, the simple instrument rule is preferred by the data, but only marginally. This is because targeting rules are heavily penalized by being prevented from accounting for the less conservative policy in the 1970s. The ranking amongst targeting rules also changes: quasi-commitment is preferred to discretion with commitment struggling to fit the data. The apparent superiority of quasi-commitment relative to other forms of targeting rule is due to the presence of policy surprises. Without allowing for switches in shock volatilities these policy surprises, largely identified during the 1970s, serve as an additional shock to increase the ability of the model to fit the data. Once switches in shock volatilities are introduced in Table 2, quasi-commitment loses this advantage over instrument rules and discretion.

[Table 3 around here]

---

[11] Under quasi-commitment there are three mechanisms through which we can explain macroeconomic volatility: high volatility shocks, periods of lower conservatism, and episodes of reneging. The nature of this trade-off is revealed when we consider the role of switching shock volatilities and policy regimes in the next section.

[12] Here we only present selected parameters. The complete set of parameter estimates is given in Appendix G.

Introducing the possibility of policy switches, but not switches in the volatility of shocks, highlights several interesting features of the benchmark estimates that would be otherwise missed, see Table 4.

[Table 4 around here]

First, the ranking of policies changes again: the policy switches can account for the less conservative regime in the 1970s enabling discretion and quasi-commitment to dominate instrument rules.

Second, the differences between the less and more conservative regimes are greater than in the case in Table 2, where the switches in shock volatilities are present. Without volatility switching, as shown in Table 4, the instrument rule does not satisfy the Taylor principle in the less conservative regime. This mirrors the findings of Sims and Zha (2006) who warn of the biases that may be introduced by failing to account for heteroscedasticity in the error terms. For the instrument rule with switches in the inflation target, the differences in the targets are also widened. Similarly, for the targeting rules, the relative weight on inflation falls by more across all policy descriptions in the less conservative regime. These results support a generalization of the arguments in Sims and Zha (2006) that failure to account for shifts in shock volatility may overstate the apparent weakness in policy during certain periods.

Third, not including switching in shock volatilities also leads to a loss of nuance in the identification of periods with less conservative regime under discretion. Without volatility switches all policy descriptions pick up the high inflation in the 1970s as being the result of a less conservative targeting regime, and that this episode ends with the Volcker disinflation somewhere between 1979 and 1982, see Figure G1 in Appendix G. However, when we combine volatility shifts with policy shifts there are additional periods where the Fed appears to have lost conservatism.[13] These are often associated with well known periods of stock market

---

[13]There are less extensive periods of reduced conservatism under quasi-commitment when we do not allow for switches in shock volatilities. In essence, the less conservative regime under quasi-commitment allows

volatility, specifically in 1987 and following the bursting of the dot-com bubble.[14]

To summarize, with no switching in objectives the targeting rules find it more difficult to account for the inflation of 1970s than instrument rules. Adding switches in policy objectives results in discretion dominating all other forms of policy, see Table 4. Allowing for switches in shock volatilities, policy surprises generated by quasi-commitment policy become relatively less effective in explaining the data. As a result, quasi-commitment moves further down in the ranking of the data-preferred policies as shown in Table 2.

6.   Comparison with Debortoli and Lakdawala (2016)

Our estimates imply that discretion dominates all other descriptions of policy. This is in contrast to the conclusions of Debortoli and Lakdawala (2016) who argue that the data reject both discretion and commitment, preferring quasi-commitment. They reach this conclusion based on the fact that the estimated probability of reneging on past promises does not tend to either zero or one in estimation. This section seeks to explore the reasons underpinning the apparent disparity in conclusions.

The first thing to note is that our estimates of the probability of reneging on past policy commitments are not dissimilar to theirs. However, the fact that the estimates do not tend to the limiting case of discretion does not imply that quasi-commitment dominates discretion in terms of its ability to explain the data. Instead, the Bayes factor implies that discretion is decisively preferred to quasi-commitment. The reason for this is that the quasi-commitment model is not actually an intermediate case lying between the cases of commitment and discretion, as discussed before in Section 3.2. Instead, it introduces policy surprises – serving as a new kind of policy shock – which arise from the fact that economic

the estimation to accommodate higher shock volatilities without inducing an overly aggressive and therefore data-incoherent policy response during reneging periods.

[14]It is interesting to note that estimates involving targeting rules can identify these periods as being associated with heightened shock volatility, while the instrument rules cannot. Across all models we observe a reduction in shock volatilities in the mid 1980s, as commonly found in the literature.

agents form expectations based on the probability of experiencing a reneging regime in the next period. The realization or otherwise of the reneging regime is then always a shock relative to these expectations. When the probability of reneging is low, economic agents expect the policy maker to keep their promises so that reneging offers the policy maker the opportunity to exploit those expectations generating a sizeable policy shock. Conversely, when there is a high probability of reneging, the policy maker makes more extreme policy promises to retain a desirable influence over expectations which, in turn, imply a large policy shock whenever the policy maker keeps that promise (see Schaumburg and Tambalotti, 2007). The estimated probability of reneging needs to balance these two scenarios to produce policy shocks that match the volatility in the data. As discussed in Section 5.3, once switches in shock volatilities are allowed, there is less need to rely on such policy shocks to fit the data.[15]

Finally, we can check that our results are not driven by adopting an objective function which takes the form of the microfounded objective function (5) rather than the simpler specification used in Debortoli and Lakdawala (2016). We consider two forms of *ad hoc* objective function based on

$$L = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\hat{\pi}_t^2 + \omega_y \hat{y}_t^2 + \omega_R \left(\Delta \hat{R}_t \right)^2 \right).$$

Loss function type I excludes the term in the interest rate smoothing, $\omega_R = 0$, and only retains terms in inflation and the output gap. Loss function type II allows the interest rate smoothing term ω_R to be estimated.

[Table 5 is around here]

In Table 5 we compare four policies, all excluding Markov switching in policy objectives but including switching in shock volatilities, as in Debortoli and Lakdawala (2016). Discre-

---

[15]The quasi-commitment estimation also adds complexity to the model in the form of an additional estimated parameter, the need to estimate the probability that we have observed a reneging regime in each period and the scale of the state-space representation of the model relative to the discretionary case. This complexity is penalized in the construction of the Bayes factors.

tion with objectives in a microfounded form dominates the three cases of quasi-commitment, with type I and II ad hoc objectives and the case with objectives in a microfounded form (5) used throughout the paper.

Three clear messages emerge from this comparison. First, with objective function (5) discretion dominates quasi-commitment. Again, this confirms that when switches in shock volatilities are accounted for the policy-surprise shocks generated by quasi-commitment are less effective in fitting the data.

Second, adding the interest rate smoothing term significantly raises the ability of quasi-commitment to fit the data. Quasi-commitment using the type II ad hoc objective achieves a better fit compared to the other two cases of quasi-commitment. However, the estimated weight on the smoothing term is implausibly large, $\omega_R = 1.5$ as the estimation seeks to limit the sharp movements in the policy instrument implied by the policy shocks described above. If we remove the interest rate smoothing term (type I ad hoc), then such policy results in the worst fit out of the four cases considered in Table 5.

Third, quasi-commitment policies with ad hoc welfare objectives identify similar probabilities of reneging and periods of high volatility as the quasi-commitment policy presented in Table 2, see Figure G2 in Appendix G.

7. Counterfactuals

The best-fitting model is obtained under discretionary policy with Markov switching in the weight on inflation target in the policy maker's objectives, as well as switches in the volatility of shocks hitting the economy. This allows us to undertake various counterfactual exercises. For example, exploring what the outcomes would have been if shock volatilities had not declined in the 1980s, or what would have happened had the Fed adopted a tougher anti-inflation stance in the 1970s. Moreover, this section explores how much further economic outcomes would have improved had the policy maker not only adopted tougher anti-inflation

policies in the 1980s, but also been able to act under commitment.

7.1. Good Luck

The series of counterfactuals begins by analyzing the role of good luck in stabilizing US output and inflation. To do so the pattern of switches in policy regimes is fixed as estimated, but the counterfactual sets the volatility of shocks at their high or low values. The estimated shocks are therefore re-scaled by the relative standard deviations from the high and low volatility regimes. Panel A of Figure 2 plots the actual and counterfactual series for inflation, interest rates and output growth. We can see that the high volatility of shocks plays a significant role in raising inflation during the 1970s. In the absence of these high volatility shocks, inflation would never have risen above 5%. In addition, it is apparent that output growth fluctuations could have been dampened if policy makers had had the 'good luck' of experiencing the low shock volatility regime during the 1970s and early 1980s. Moreover, it is also notable that under the policy regimes estimated in the post-Volcker period, inflation and output fluctuations would not have changed too dramatically regardless of the magnitude of shocks. This may be an indication that tougher anti-inflation policies in the 1980s helped stabilize the US economy.

[Figure 2 around here]

7.2. Conservative Monetary Policy

The second set of counterfactual analyses assesses the impact that increased conservatism would have had on US inflation and output, especially during the 1970s. To simulate the set of counterfactual variables we subject the economy to the sequence of estimated shocks, but set the weight on inflation in the policy maker's objective function, ω_π, to either its default value of one in the more conservative regime, or to 0.436 in the less conservative regime, throughout the sample period. The first two pictures in Panel B of Figure 2 plot the actual

and counterfactual series for inflation and interest rates. The third picture plots the output loss, which is the difference between model implied output with estimated objective function weights and the counterfactual output when the policy maker is more conservative.

Panel B of Figure 2 shows that even if the Fed had adopted a tougher anti-inflation stance in the 1970s, it would not have been able to completely avoid higher inflation, but observed inflation would have been significantly lowered at a cost of higher output losses. Similarly, the two periods of rising inflation that occurred following the stock market crash of 1987 and the bursting of the dot-com bubble could also have been mitigated if the Fed had maintained its stance on inflation targeting. The counterfactual paths for interest rates largely reflect the tightness or slackness of policy implied by the alternative scenarios. However, since the effective stance of monetary policy is reflected in the real interest rate, the path for nominal interest rates under the less conservative policy are above those implied by the more conservative policy, reflecting the latter's success in controlling inflation.

7.3. The Value of Commitment

Finally, Panel C of Figure 2 assesses the implications of moving from discretion to commitment. Both the shock volatility and policy switches follow their estimated realizations, but we change whether or not the policy maker has access to a commitment technology. The results are striking. If the Fed had been able to make credible policy commitments in the 1970s, even although it was subject to high volatility shocks and had a reduced weight on the inflation target in that period, inflation would have remained below 2% throughout the sample period. Although it appears that there would have been non-trivial losses in output with a peak loss of around 1% by the mid 1970s, the welfare analysis in the next section suggests that these losses are more than compensated for by the reduction in inflation volatility.

7.4.   Welfare Analysis

In addition to providing the counterfactual figures above, it is insightful to compute the unconditional variances of key variables and the value of unconditional welfare (using both the estimated policy objective (5) and the fully microfounded objectives where the weights are microfounded functions of the estimated structural parameters of the model) under alternative counterfactuals.

As a benchmark case we consider the worst case scenario where the economy is permanently in the high shock volatility regime and adopt a less conservative policy with $\omega_\pi = 0.436$ under discretion. We can then consider the extent to which 'good policy' or 'good luck' alone would be able to stabilize inflation, output and interest rates and improve welfare.

Table 6 presents variances of output, inflation and interest rate under different conservatism – volatility scenarios. The degree of conservatism ranges from that estimated under the less conservative through the more conservative regimes, both of which are using estimated policy objective weights, to the extreme level implied by the fully microfounded welfare function. Two welfare metrics are used to measure losses, one with estimated weights and the other with microfounded weights.

[Table 6 around here]

Panel A in Table 6 shows that under discretion either implementing the 'more' conservative regime, or enjoying a reduction in shock volatility alone, would reduce by more than half the volatility in inflation and interest rates implied by the worst case scenario. However, it is the 'good luck' that would lead to significant output stabilization and, therefore, achieve bigger gains as measured by either the central bank's estimated or the microfounded welfare metrics.

If the policy maker further increases the level of conservatism to the levels implied by the microfounded objectives, there is a striking reduction in inflation volatility to negligible levels.

However, it significantly worsens output volatility in the high volatility regime. Clearly, the Fed has not implemented monetary policy with a degree of inflation conservatism anywhere near that implied by microfounded objectives.

Turning to Panel B of Table 6 we consider the same experiment, but now assume that policy is conducted under commitment. In the absence of 'good luck', being able to act with commitment allows the central bank to almost completely stabilize inflation volatility, but at the cost of moderate increases in output fluctuations. It is also important to note that welfare is clearly improved regardless of the degree of central bank conservatism. This result suggests that the reduction in inflation volatility achieved by being able to act under commitment is such that the issue of conservatism becomes of second-order importance. Therefore, the dimension of 'good policy' policymakers should be concerned with is not the weight given to inflation stabilization in the policy maker's objective function, i.e. the conservatism of the central bank, but rather that they have the tools and credibility to effectively pursue a commitment policy and to make time-inconsistent promises which they will keep. Finally, under commitment we again see substantial decreases in output volatility when there is good luck.

8.   Conclusions

A time consistent targeting rule – discretionary policy – provides the best fit to the data, outperforming conventional instrument rules and the other forms of optimal policy with different degrees of precommitment. Bayes factors suggest that there is 'strong' evidence in favor of this description of policy relative to simple instrument rules, and 'decisive' evidence relative to targeting rules formed under either commitment or quasi-commitment. However, the ranking of policies in terms of fitting the data crucially depends on whether or not we account for potential changes in the Fed's degree of inflation conservatism and in shock volatilities. A failure to take into account policy switches hinders the ability of targeting

rules to account for the monetary policy response to the high inflation of the 1970s relative to instrument rules. The absence of variation in shock volatilities exaggerates the fit of quasi-commitment because it can rely on policy surprises as a source of volatility. We demonstrate how inferences about shock processes, habit persistence and inflation indexation change across different policy specifications.

The preferred model implies that there was an increase in central bank conservatism following the Volcker disinflation period, which is estimated to occur in 1982. This description of policy also finds that the Fed relaxed policy temporarily in the aftermath of the 1987 stock market crash, and also lost conservatism following the 2000 dot-com crash, which it has never regained.

Based on estimates from the best-fit model, a range of counterfactual simulations are undertaken which throw light on various aspects of policy. First, there have been significant welfare gains to the conservatism in policy making that was adopted following the Volcker disinflation. However, these gains are small compared to those attained from the estimated reduction in shock volatilities. Relative to the average rate of inflation of 6.51% in the 1970s, a policy maker acting under discretion, but with the higher degree of conservatism observed later on in the sample, would have reduced average inflation to 4.71%. In contrast, inflation would have been expected to be 3.39% in the same period had the economy been lucky enough to have been in the low volatility regime. Second, had the US Fed been able to commit, rather than acting under discretion, then in the 1970s the average rate of inflation would have been below 2%, regardless of the level of conservatism. Taken together, this suggests that attempts to improve monetary policy outcomes should concentrate on ensuring that the Fed is able to make and communicate credible promises concerning future policy, and that this is of more importance than altering the preferences of the central banker.

References

Adolfson, M., Laseen, S., Linde, J., Svensson, L. E., 2011. Optimal Monetary Policy in an Operational Medium-Sized DSGE Model. Journal of Money, Credit and Banking 43 (7), 1287–1331.

An, S., Schorfheide, F., 2007. Bayesian Analysis of DSGE Models. Econometric Reviews 26 (2-4), 113–172.

Bianchi, F., 2013. Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics. Review of Economic Studies 80 (2), 491–511.

Boivin, J., Giannoni, M. P., 2006. Has Monetary Policy Become More Effective? Review of Economics and Statistics 88 (3), 445–462.

Caglar, E., Chadha, J. S., Shibayama, K., 2011. Bayesian Estimation of DSGE models: Is the Workhorse Model Identified?, working Paper No 1125, University of Kent.

Calvo, G., 1983. Staggered Prices in a Utility-Maximising Framework. Journal of Monetary Economics 12, 383–398.

Clarida, R. H., Galí, J., Gertler, M., 1998. Monetary Policy Rules in Practice. Some International Evidence. European Economic Review 42, 1033–1067.

Davig, T., Doh, T., 2014. Monetary policy regime shifts and inflation persistence. The Review of Economics and Statistics 96 (5), 862–875.

Debortoli, D., Lakdawala, A., 2016. How Credible is the Federal Reserve? A Structural Estimation of Policy Re-optimizations. AEJ: Macroeconomics 8 (3), 42–76.

Debortoli, D., Nunes, R., 2010. Fiscal Policy under Loose Commitment. Journal of Economic Theory 145 (3), 1005–1032.

Dennis, R., 2004. Inferring Policy Objectives from Economic Outcomes. Oxford Bulletin of Economics and Statistics 66 (S1), 735–64.

Erceg, C. J., Levin, A. T., 2003. Imperfect Credibility and Inflation Persistence. Journal of Monetary Economics 50, 915–944.

Farmer, R. E. A., Waggoner, D. F., Zha, T., 2011. Minimal State Variable Solutions to Markov-switching Rational Expectations Models. Journal of Economic Dynamics and Control 35 (12), 2150–2166.

Favero, C. A., Rovelli, R., 2003. Macroeconomic stability and the preferences of the Fed. A formal analysis, 1961-98. Journal of Money, Credit and Banking 35 (4), 545–556.

Geweke, J., 1999. Using simulation methods for Bayesian econometric models: inference, development, and communication. Econometric Reviews 18 (1), 1–73.

Givens, G. E., 2012. Estimating Central Bank Preferences under Commitment and Discretion. Journal of Money, Credit and Banking 44 (6), 1033–1061.

Hamilton, J. D., Waggoner, D. F., Zha, T., 2007. Normalisation in Econometrics. Econometric Reviews 26, 221–252.

Ilbas, P., 2010. Estimation of monetary policy preferences in a forward-looking model: a Bayesian approach. International Journal of Central Banking 6 (3), 169–209.

Ireland, P. N., 2007. Changes in the Federal Reserve's Inflation Target: Causes and Consequences. Journal of Money, Credit, and Banking 39 (8), 1851–1882.

Iskrev, N., 2010. Local identification in DSGE models. Journal of Monetary Economics 52 (7), 189–202.

Kass, R. E., Raftery, A. E., 1995. Bayes Factors. Journal of the American Statistical Association 90 (430), 773–795.

Komunjer, I., Ng, S., 2011a. Dynamic Identification of Dynamic Stochastic General Equilibrium Models. Econometrica 79 (6), 1995–2032.

Komunjer, I., Ng, S., 2011b. Supplementary Material to 'Dynamic Identification of Dynamic Stochastic General Equilibrium Models'. Econometrica 79 (6), 1995–2032.

Koop, G., Pesaran, M. H., Smith, R. P., 2013. On Identification of Bayesian DSGE Models. Journal of Business and Economic Statistics 31 (3), 300–314.

Le Roux, S., Kirsanova, T., 2013. Commitment vs. Discretion in the UK: An Empirical Investigation of the Monetary and Fiscal Policy Regime. International Journal of Central Banking 9 (4), 99–151.

Leith, C., Moldovan, I., Rossi, R., 2012. Optimal Monetary Policy in a New Keynesian Model with Habits in Consumption. Review of Economic Dynamics 15 (3), 416–435.

Lubik, T., Schorfheide, F., 2005. A Bayesian Look at the New Open Economy Macroeconomics. In: NBER Macroeconomics Annual 2005. MIT Press, pp. 313–379.

Schaumburg, E., Tambalotti, A., 2007. An Investigation of the Gains from Commitment in Monetary Policy. Journal of Monetary Economics 54 (2), 302–324.

Schmitt-Grohe, S., Uribe, M., 2007. Optimal Simple and Implementable Monetary and Fiscal Rules. Journal of Monetary Economics 54 (6), 1702–1725.

Schorfheide, F., 2005. Learning and Monetary Policy Shifts. Review of Economic Dynamics 8, 392–419.

Sims, C. A., Waggoner, D. F., Zha, T., 2008. Methods for inference in large multiple-equation Markov-switching models. Journal of Econometrics 146, 113–144.

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

Smets, F., Wouters, R., 2003. An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. Journal of the European Economic Association 1 (1), 1123–75.

Svensson, L., 2003. What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules. Journal of Economic Literature 41, 426–477.

Svensson, L., Williams, N., 2007. Monetary Policy with Model Uncertainty: Distribution Forecast Targeting, cEPR Discussion Papers 6331, C.E.P.R. Discussion Papers.

Svensson, L. E., Woodford, M., 2005. Implementing Optimal Policy through Inflation-Forecast Targeting. In: Bernanke, B. S., Woodford, M. (Eds.), The Inflation-Targeting Debate. University of Chicago Press, pp. 19–92.

Taylor, J. B., 1999. Staggered Price and Wage Setting in Macroeconomics. In: Taylor, J. B., Woodford, M. (Eds.), Handbook of Macroeconomics. North Holland, pp. 1009–1050.

Table 1: Distribution of Priors

Parameters		Range	Density	Mean	Std Dev
inv. of intertemp. elas. of subst.	σ	$\mathbb{R}$	Normal	2.50	0.25
Calvo parameter	α	$[0,1)$	Beta	0.75	0.02
inflation inertia	ζ	$[0,1)$	Beta	0.50	0.15
habit persistence	θ	$[0,1)$	Beta	0.50	0.15
inverse of Frisch elasticity	φ	$\mathbb{R}$	Normal	2.50	0.25
AR coeff., taste shock	ρ^ξ	$[0,1)$	Beta	0.50	0.15
AR coeff., cost-push shock	ρ^μ	$[0,1)$	Beta	0.50	0.15
AR coeff., productivity shock	ρ^z	$[0,1)$	Beta	0.50	0.15
steady state interest rate	r^A	$\mathbb{R}^+$	Gamma	3.5	2
inflation target	π^A	$\mathbb{R}^+$	Gamma	3.5	2
steady state growth rate	γ^Q	$\mathbb{R}$	Nomal	0.52	1
probability of reneging	υ	$[0,1]$	Uniform	0.5	0.25
Markov Switching s.d. of shocks					
preference shocks	$\sigma_{\xi(S=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
cost-push shocks	$\sigma_{\mu(S=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
technology shocks	$\sigma_{z(S=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
policy shocks	$\sigma_{R(S=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
Markov switching rule parameters					
interest rate smoothing	$\rho^R_{(s=1=2)}$	$[0,1)$	Beta	0.50	0.25
inflation (more conservative)	$\psi_{1(s=1)}$	$\mathbb{R}^+$	Gamma	1.50	0.50
inflation (less conservative)	$\psi_{1(s=2)}$	$\mathbb{R}^+$	Gamma	1.0	0.50
output	$\psi_{2(s=1=2)}$	$\mathbb{R}^+$	Gamma	0.50	0.25
Weights on Objectives					
gap term, $\hat{X}_t - \hat{\xi}_t$	ω_1	$[0,1)$	Beta	0.50	0.15
gap term, $\hat{y}_t - \frac{\sigma}{\varphi}\hat{\xi}_t$	ω_2	$[0,1)$	Beta	0.50	0.15
change in inflation, $\hat{\pi}_t - \hat{\pi}_{t-1}$	ω_3	$[0,1)$	Beta	0.50	0.15
inflation, $\hat{\pi}_t$	$\omega_{\pi(s=2)}$	$[0,1)$	Beta	0.50	0.15
Markov switching in Inflation Target					
inflation target $(s=1)$	$\pi^A_{(s=1)}$	$\mathbb{R}^+$	Gamma	6	2
inflation target $(s=2)$	$\pi^A_{(s=2)}$	$\mathbb{R}^+$	Gamma	3	2
Transition Probabilities					
policy: remains more conservative	p_{11}	$[0,1)$	Beta	0.90	0.05
policy: remains less conservative	p_{22}	$[0,1)$	Beta	0.90	0.05
volatility: remains with low volatility	q_{11}	$[0,1)$	Beta	0.90	0.05
volatility: remains with high volatility	q_{22}	$[0,1)$	Beta	0.90	0.05

Notes: For policy switches $s=1$ is more conservative regime and $s=2$ is less conservative regime. For volatility switches $S=1$ is less volatile regime and $S=2$ is more volatile regime.

Table 2: Estimation Results

Parameters	Discretion	Rule - Parameters	Rule - Target	Quasi-Commitment	Commitment
			Model Parameters		
σ	2.901 [2.526,3.244]	2.937 [2.564,3.309]	2.934 [2.556,3.301]	2.692 [2.356,3.038]	2.912 [2.480,3.338]
α	0.735 [0.708,0.763]	0.770 [0.742,0.799]	0.775 [0.746,0.804]	0.754 [0.732,0.776]	0.775 [0.748,0.803]
ζ	0.165 [0.069,0.254]	0.088 [0.031,0.142]	0.084 [0.030,0.138]	0.182 [0.096,0.270]	0.262 [0.114,0.419]
θ	0.387 [0.206,0.560]	0.827 [0.702,0.956]	0.790 [0.631,0.950]	0.372 [0.201,0.544]	0.694 [0.304,0.953]
φ	2.459 [2.060,2.844]	2.442 [2.030,2.855]	2.424 [2.004,2.838]	2.286 [1.889,2.672]	2.199 [1.782,2.638]
			Shock Processes		
ρ^{ξ}	0.830 [0.791,0.870]	0.890 [0.853,0.927]	0.901 [0.866,0.938]	0.893 [0.869,0.919]	0.919 [0.898,0.941]
ρ^{μ}	0.939 [0.914,0.963]	0.504 [0.262,0.759]	0.502 [0.252,0.751]	0.923 [0.900,0.948]	0.992 [0.986,0.998]
ρ^{z}	0.195 [0.141,0.248]	0.329 [0.228,0.427]	0.359 [0.257,0.462]	0.186 [0.134,0.238]	0.162 [0.106,0.218]
$\sigma_{\xi(S=1)}$	0.425 [0.297,0.546]	0.682 [0.527,0.837]	0.545 [0.390,0.690]	0.495 [0.334,0.649]	0.404 [0.249,0.555]
$\sigma_{\xi(S=2)}$	0.873 [0.599,1.139]	1.467 [1.040,1.888]	1.346 [0.958,1.721]	0.909 [0.652,1.167]	1.224 [0.720,1.757]
$\sigma_{\mu(S=1)}$	0.236 [0.182,0.292]	0.277 [0.169,0.381]	0.276 [0.169,0.383]	0.251 [0.188,0.315]	1.329 [0.737,1.905]
$\sigma_{\mu(S=2)}$	0.684 [0.527,0.840]	0.546 [0.343,0.751]	0.545 [0.390,0.690]	0.864 [0.658,1.065]	2.806 [1.697,3.913]
$\sigma_{z(S=1)}$	0.512 [0.391,0.622]	0.601 [0.540,0.660]	0.603 [0.542,0.664]	0.433 [0.352,0.515]	0.452 [0.372,0.526]
$\sigma_{z(S=2)}$	1.064 [0.932,1.193]	1.184 [0.981,1.380]	1.156 [0.977,1.329]	1.034 [0.918,1.148]	0.989 [0.870,1.103]
$\sigma_{R(S=1)}$	−	0.140 [0.124,0.156]	0.146 [0.129,0.162]	−	−
$\sigma_{R(S=2)}$	−	0.412 [0.332,0.489]	0.455 [0.379,0.529]	−	−
			Data Means		
r^{A}	0.802 [0.294,1.282]	0.541 [0.189,0.873]	0.509 [0.165,0.828]	0.803 [0.330,1.266]	0.722 [0.257,1.184]
$\pi^{A}_{(s=1)}$	1.305 [0.629,1.943]	3.558 [2.986,4.122]	3.336 [2.745,3.948]	1.962 [1.588,2.326]	2.755 [2.303,3.189]
$\pi^{A}_{(s=2)}$	−	−	4.329 [3.662,5.001]	−	−
γ^{Q}	0.773 [0.669,0.897]	0.713 [0.592,0.832]	0.700 [0.566,0.829]	0.790 [0.697,0.885]	0.828 [0.721,0.931]

Table 2: Estimation Results – continued

Parameters	Discretion	Rule - Parameters	Rule - Target	Quasi-Commitment	Commitment
			Policy Parameters		
υ	–	–	–	0.290 [0.227,0.355]	–
$\rho^R_{(s=1)}$	–	0.825 [0.793,0.858]	0.821 [0.793,0.851]	–	–
$\rho^R_{(s=2)}$	–	0.868 [0.779,0.946]	–	–	–
$\psi_{1(s=1)}$	–	2.124 [1.798,2.447]	2.014 [1.655,2.370]	–	–
$\psi_{1(s=2)}$	–	1.219 [0.809,1.635]	–	–	–
$\psi_{2(s=1)}$	–	0.511 [0.327,0.692]	0.587 [0.381,0.784]	–	–
$\psi_{2(s=2)}$	–	0.274 [0.102,0.438]	–	–	–
ω_1	0.380 [0.232,0.534]	–	–	0.624 [0.476,0.777]	0.503 [0.320,0.690]
ω_2	0.635 [0.468,0.800]	–	–	0.749 [0.618,0.884]	0.559 [0.280,0.843]
ω_3	0.436 [0.200,0.667]	–	–	0.369 [0.141,0.586]	0.454 [0.195,0.695]
$\omega_{\pi(s=1)}$	1	–	–	1	1
$\omega_{\pi(s=2)}$	0.436 [0.279,0.589]	–	–	0.301 [0.204,0.395]	0.373 [0.216,0.527]
			Markov Transition Probabilities		
p_{11}	0.947 [0.903,0.989]	0.964 [0.942,0.988]	0.902 [0.840,0.964]	0.798 [0.715,0.882]	0.978 [0.959,0.997]
p_{22}	0.918 [0.876,0.962]	0.846 [0.812,0.880]	0.812 [0.740,0.889]	0.914 [0.865,0.966]	0.798 [0.722,0.877]
q_{11}	0.952 [0.919,0.986]	0.956 [0.928,0.985]	0.979 [0.960,0.998]	0.907 [0.852,0.962]	0.958 [0.931,0.986]
q_{22}	0.955 [0.910,0.997]	0.843 [0.779,0.910]	0.946 [0.902,0.992]	0.941 [0.905,0.977]	0.933 [0.887,0.976]
			Log Marginal Data Densities and Bayes Factors		
Geweke	−759.78 (1.00)	−764.16 (80.29)	−765.83 (425.76)	−770.29 (3.67e+4)	−793.62 (4.98e+14)
Sims et.al.	−759.91 (1.00)	−764.21 (74.08)	−765.95 (422.76)	−770.34 (3.40e+4)	−793.95 (6.12e+14)

Notes: Here and in Tables 3-5 for each parameter the posterior distribution is described by its mean and 90% confidence interval in square brackets. Bayes Factors for marginal data densities are in parentheses.

Table 3: Selected Parameter Estimates - No Switching

Parameters	Simple Rule	Quasi-Commitment	Discretion	Commitment
		Selected Model Parameters		
ζ	0.103 [0.039,0.166]	0.170 [0.087,0.252]	0.156 [0.066,0.241]	0.594 [0.489,0.737]
θ	0.823 [0.685,0.964]	0.421 [0.210,0.627]	0.476 [0.267,0.680]	0.643 [0.444,0.782]
		Policy Parameters		
υ	−	0.556 [0.329,0.816]	−	−
ρ^R	0.791 [0.756,0.826]	−	−	−
ψ_1	1.716 [1.455,1.972]	−	−	−
ψ_2	0.492 [0.290,0.697]	−	−	−
ω_1	−	0.703 [0.552,0.861]	0.458 [0.287,0.627]	0.627 [0.490,0.808]
ω_2	−	0.828 [0.727,0.935]	0.758 [0.628,0.901]	0.446 [0.316,0.620]
ω_3	−	0.390 [0.163,0.619]	0.451 [0.213,0.692]	0.489 [0.268,0.712]
		Data Means		
r^A	0.706 [0.246,1.139]	0.759 [0.143,1.330]	0.966 [0.352,1.569]	1.088 [0.459,1.540]
π^A	4.746 [3.800,5.677]	2.586 [1.899,3.095]	2.656 [1.008,4.221]	4.050 [3.642,4.674]
γ^Q	0.688 [0.547,0.826]	0.737 [0.613,0.861]	0.716 [0.593,0.835]	0.726 [0.594,0.797]
		Log Marginal Data Densities and Bayes Factors		
Geweke	−841.01 (1.00)	−841.67 (1.94)	−842.49 (4.41)	−855.43 (1.84e+6)
Sims et.al	−841.09 (1.00)	−841.54 (1.57)	−842.69 (4.96)	−858.26 (2.85e+7)

Table 4: Selected Parameter Estimates - Switches in Policy Only

Parameters	Discretion	Quasi-Commitment	Rule - Parameters	Rule - Target	Commitment
			Selected Model Parameters		
ζ	0.155 [0.069,0.239]	0.182 [0.091,0.274]	0.102 [0.038,0.163]	0.123 [0.054,0.195]	0.229 [0.078,0.366]
θ	0.479 [0.286,0.835]	0.371 [0.192,0.543]	0.825 [0.698,0.954]	0.810 [0.658,0.961]	0.606 [0.388,0.843]
			Policy Parameters		
υ	−	0.325 [0.239,0.411]	−	−	−
$\rho^R_{(s=1)}$	−	−	0.746 [0.708,0.786]	0.797 [0.762,0.831]	−
$\rho^R_{(s=2)}$	−	−	0.845 [0.794,0.900]	−	−
$\psi_{1(s=1)}$	−	−	2.075 [1.824,2.315]	1.805 [1.507,2.097]	−
$\psi_{1(s=2)}$	−	−	0.909 [0.621,1.189]	−	−
$\psi_{2(s=1)}$	−	−	0.483 [0.309,0.645]	0.498 [0.285,0.714]	−
$\psi_{2(s=2)}$	−	−	0.245 [0.098,0.393]	−	−
ω_1	0.259 [0.035,0.414]	0.633 [0.480,0.785]	−	−	0.502 [0.331,0.666]
ω_2	0.650 [0.460,0.847]	0.759 [0.631,0.893]	−	−	0.523 [0.295,0.732]
ω_3	0.442 [0.164,0.698]	0.349 [0.126,0.559]	−	−	0.460 [0.205,0.710]
$\omega_{\pi(s=1)}$	1	1	−	−	1
$\omega_{\pi(s=2)}$	0.347 [0.219,0.477]	0.348 [0.254,0.440]	−	−	0.302 [0.194,0.414]
			Data Means		
r^A	0.766 [0.303,1.213]	0.997 [0.377,1.591]	0.695 [0.276,1.105]	0.662 [0.239,1.054]	0.975 [0.358,1.561]
$\pi^A_{(s=1)}$	2.683 [1.275,4.022]	2.097 [1.770,2.431]	3.736 [3.183,4.299]	4.234 [3.470,4.995]	3.064 [2.733,3.411]
$\pi^A_{(s=2)}$	−	−	−	6.058 [5.217,6.862]	−
γ^Q	0.683 [0.567,0.800]	0.722 [0.598,0.842]	0.677 [0.540,0.808]	0.681 [0.544,0.822]	0.741 [0.619,0.862]
			Log Marginal Data Densities and Bayes Factors		
Geweke	−810.98 (1.00)	−814.83 (47.0)	−825.33 (1.72e+6)	−831.74 (1.04e+9)	−832.85 (3.14e+9)
Sims et.al.	−811.24 (1.00)	−814.30 (21.21)	−825.44 (1.46e+6)	−831.81 (8.52e+8)	−832.98 (2.75e+9)

Table 5: Estimation Results - MS Shocks only

Para-meters	Discretion micro-founded objective	Quasi-commitment Ad Hoc objective type II	Quasi-commitment micro-founded objective	Quasi-commitment Ad Hoc objective type I
Model Parameters				
σ	2.866 [2.503,3.227]	2.588 [2.220,2.944]	2.375 [2.054,2.688]	2.186 [1.851,2.509]
α	0.751 [0.724,0.779]	0.808 [0.788,0.827]	0.787 [0.765,0.808]	0.811 [0.793,0.828]
ζ	0.173 [0.075,0.261]	0.173 [0.089,0.256]	0.194 [0.142,0.243]	0.153 [0.080,0.224]
θ	0.459 [0.220,0.715]	0.495 [0.409,0.580]	0.478 [0.383,0.570]	0.293 [0.233,0.349]
φ	2.274 [1.872,2.675]	2.020 [1.577,2.439]	1.793 [1.453,2.137]	2.031 [1.614,2.436]
Shock Processes				
ρ^{ξ}	0.843 [0.810,0.877]	0.822 [0.758,0.891]	0.898 [0.875,0.922]	0.875 [0.842,0.910]
ρ^{μ}	0.936 [0.911,0.961]	0.926 [0.891,0.963]	0.930 [0.903,0.957]	0.936 [0.907,0.968]
ρ^{z}	0.183 [0.132,0.239]	0.300 [0.215,0.386]	0.194 [0.142,0.243]	0.201 [0.142,0.258]
$\sigma_{\xi(S=1)}$	0.443 [0.311,0.575]	0.510 [0.315,0.709]	0.510 [0.332,0.681]	0.480 [0.340,0.616]
$\sigma_{\xi(S=2)}$	0.898 [0.622,1.171]	1.905 [1.187,2.657]	1.082 [0.747,1.404]	1.186 [0.846,1.509]
$\sigma_{\mu(S=1)}$	0.234 [0.178,0.286]	0.829 [0.433,1.260]	0.317 [0.219,0.411]	0.583 [0.372,0.781]
$\sigma_{\mu(S=2)}$	0.769 [0.579,0.951]	2.247 [1.557,2.910]	1.094 [0.773,1.431]	1.801 [1.345,2.239]
$\sigma_{z(S=1)}$	0.476 [0.380,0.569]	0.526 [0.441,0.610]	0.450 [0.358,0.542]	0.438 [0.347,0.526]
$\sigma_{z(S=2)}$	1.064 [0.361,1.189]	0.962 [0.794,1.111]	1.061 [0.937,1.184]	1.024 [0.893,1.148]
Data Means				
r^{A}	0.763 [0.277,1.213]	0.732 [0.249,1.218]	0.666 [0.245,1.082]	0.662 [0.241,1.071]
π^{A}	1.706 [0.693,2.643]	2.276 [1.879,2.678]	2.150 [1.674,2.636]	2.481 [2.178,2.793]
γ^{Q}	0.789 [0.692,0.885]	0.761 [0.645,0.882]	0.783 [0.682,0.883]	0.786 [0.684,0.887]

Table 5: Estimation Results - MS Shocks only – continued

Parameters	Discretion micro-founded objective	Quasi-commitment Ad Hoc objective type II	Quasi-commitment micro-founded objective	Quasi-commitment Ad Hoc objective type I
Policy Parameters				
υ	–	0.194 [0.127,0.261]	0.260 [0.187,0.333]	0.144 [0.110,0.177]
ρ^R	–	–	—	–
ψ_1	–	–	—	–
ψ_2	–	–	—	–
ω_1	0.454 [0.275,0.642]	–	0.746 [0.609,0.882]	–
ω_2	0.715 [0.569,0.867]	–	0.819 [0.714,0.927]	–
ω_3	0.444 [0.198,0.676]	–	0.402 [0.167,0.633]	–
ω_π	1	1	1	1
ω_y	–	0.819 [0.711,0.933]	–	0.866 [0.781,0.952]
ω_R	–	1.533 [0.734,2.349]	–	–
Markov Transition Probabilities				
p_{11}	0.916 [0.866,0.968]	0.948 [0.902,0.997]	0.900 [0.840,0.964]	0.879 [0.813,0.944]
p_{22}	0.892 [0.849,0.934]	0.959 [0.931,0.986]	0.939 [0.904,0.973]	0.940 [0.903,0.978]
Log Marginal Data Densities and Bayes Factors				
Geweke	−776.22 (1.0)	−782.97 (854.06)	−792.73 (1.48e+7)	−837.80 (5.54e+26)
Sims et.al	−776.23 (1.0)	−782.81 (718.38)	−792.74 (1.49e+07)	−837.64 (4.67e+26)

Note: The prior for ω_R is Gamma (1,1) and for ω_y it is Beta (0.5,0.15).

Table 6: Unconditional Variances and Welfare under Alternative Policies and Volatilities

Regime: (conservatism, volatility)	Output	Inflation	Interest Rate	Welfare Cost (est. weights)	Welfare Cost (micro. weights)
A: Discretion					
(less, high)*	0.147 [0.092,0.228]	2.044 [1.413,3.157]	1.452 [0.936,2.459]	3.726 [2.250,6.554]	1.05% [0.69%,1.54%]
(more, high)	0.151 [0.100,0.234]	0.698 [0.467,1.00]	0.593 [0.449,0.844]	3.584 [2.126,6.397]	0.41% [0.30%,0.60%]
(micro, high)	0.177 [0.127,0.259]	0.002 [0.001,0.003]	0.480 [0.403,0.566]	–	0.08% [0.05%,0.15%]
(less, low)	0.060 [0.036,0.093]	0.798 [0.541,1.231]	0.509 [0.311,0.893]	0.811 [0.485,1.451]	0.17% [0.11%,0.26%]
(high, low)	0.057 [0.035,0.089]	0.281 [0.179,0.407]	0.223 [0.166,0.322]	0.793 [0.470,1.435]	0.07% [0.05%,0.115%]
(micro, low)	0.061 [0.042,0.094]	0.001 [0.000,0.001]	0.232 [0.193,0.276]	–	0.02% [0.01%,0.03%]
B: Commitment					
(less, high)	0.166 [0.112,0.250]	0.053 [0.037,0.081]	0.746 [0.624,0.893]	2.982 [1.588,5.720]	0.13% [0.09%,0.20%]
(more, high)	0.168 [0.117,0.251]	0.018 [0.012,0.026]	0.697 [0.0.586,0.829]	3.009 [1.616,5.753]	0.10% [0.07%,0.17%]
(micro, high)	0.179 [0.129,0.261]	0.000 [0.000,0.000]	0.463 [0.387,0.547]	–	0.08% [0.05%,0.15%]
(less, low)	0.062 [0.040,0.095]	0.023 [0.015,0.033]	0.364 [0.296,0.446]	0.688 [0.377,1.319]	0.03% [0.02%,0.04%]
(more, low)	0.061 [0.040,0.094]	0.008 [0.005,0.012]	0.341 [0.279,0.414]	0.694 [0.383,1.326]	0.02% [0.02%,0.04%]
(micro, low)	0.062 [0.042,0.095]	0.000 [0.000,0.000]	0.225 [0.187,0.268]	–	0.02% [0.01%,0.03%]

Notes: The welfare costs are computed using equation (5) where weights are either estimated or microfounded functions of estimated structural parameters. The microfounded welfare costs are expressed as a percentage of steady-state consumption. For both commitment and discretionary policy we compute social welfare using regimes and regime parameters identified for discretionary policy.

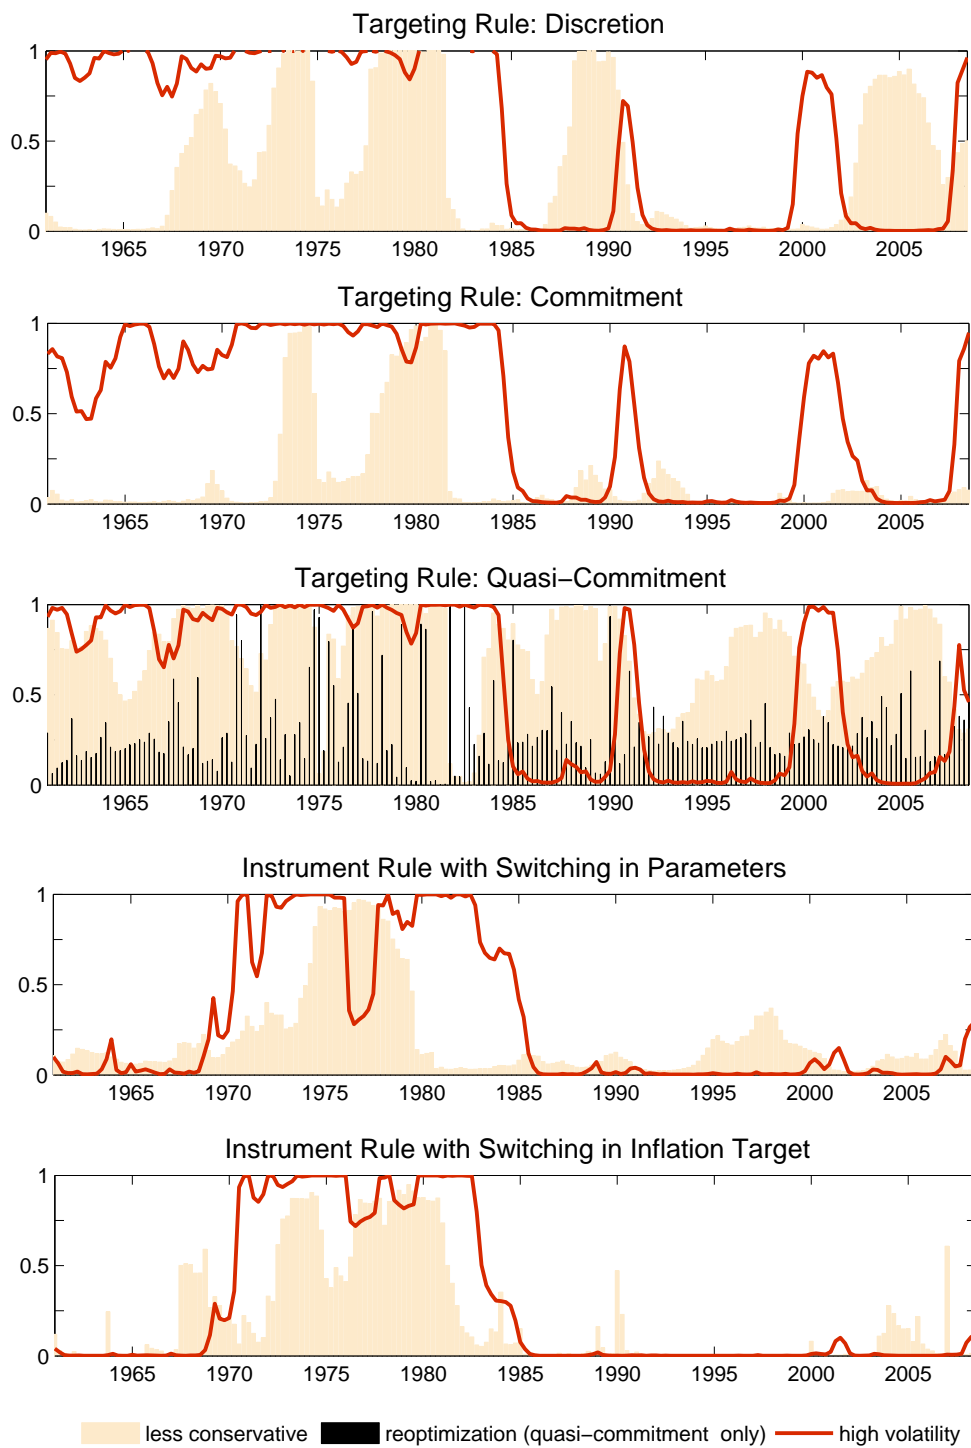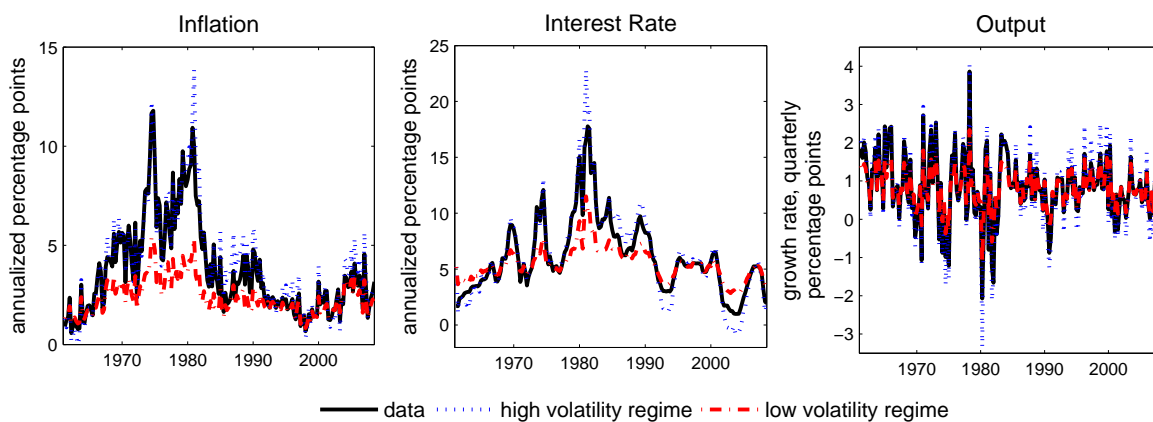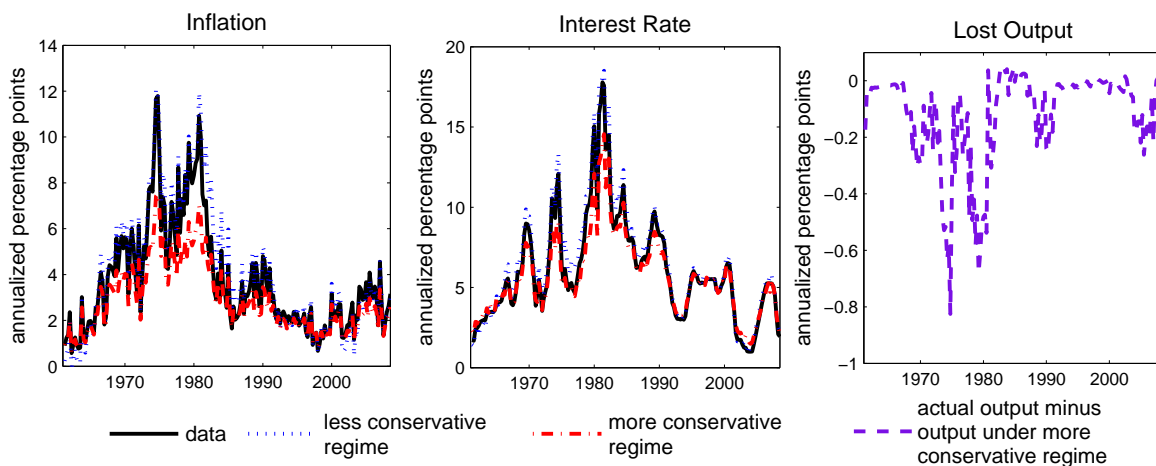Figure 1: Markov Switching Probabilities: Policy and Volatility Switches

Figure 2: Counterfactuals

A: Good Luck

Inflation	Interest Rate	Output

data · · · · · high volatility regime · – · – low volatility regime

B: Conservative Monetary Policy

Inflation	Interest Rate	Lost Output

data · · · · · less conservative regime · – · – more conservative regime

– – actual output minus output under more conservative regime

C: The Value of Commitment

Inflation	Interest Rate	Lost Output

data · · · · · discretion · – · – commitment

– – actual output minus output under commitment regime