How Optimal is US Monetary Policy?*

Xiaoshan Chen †
University of Stirling

Tatiana Kirsanova‡
University of Glasgow

Campbell Leith§
University of Glasgow

March 28, 2013

Abstract

Most of the literature estimating DSGE models for monetary policy analysis assume that policy follows a simple rule. In this paper we allow policy to be described by various forms of optimal policy - commitment, discretion and quasi-commitment. We find that, even after allowing for Markov switching in shock variances, the inflation target and/or rule parameters, the data preferred description of policy is that the US Fed operates under discretion with a marked increase in conservatism after the 1970s. Parameter estimates are similar to those obtained under simple rules, except that the degree of habits is significantly lower and the prevalence of cost-push shocks greater. Moreover, we find that the greatest welfare gains from the 'Great Moderation' arose from the reduction in the variances in shocks hitting the economy, rather than increased inflation aversion. However, much of the high inflation of the 1970s could have been avoided had policy makers been able to commit, even without adopting stronger anti-inflation objectives. More recently the Fed appears to have temporarily relaxed policy following the 1987 stock market crash, and has lost, without regaining, its post-Volcker conservatism following the bursting of the dot-com bubble in 2000.

Key Words: Bayesian Estimation, Interest Rate Rules, Optimal Monetary Policy, Great Moderation, Commitment, Discretion

JEL Reference Numbers: E58, E32, C11, C51, C52, C54

*We would like to thank Davide Debortoli, Eric Leeper, Charles Nolan, John Tsoukalas and seminar participants at the University of Surrey for helpful comments. All errors are ours.

†Address: Department of Economics, Stirling Management School, University of Stirling, Stirling, FK9 4LA; e-mail xiaoshan.chen@stirling.ac.uk

‡Address: Economics, Adam Smith Business School, Gilbert Scott Building, University of Glasgow, Glasgow G12 8QQ; e-mail tatiana.kirsanova@glasgow.ac.uk

§Address: Economics, Adam Smith Business School, Gilbert Scott Building, University of Glasgow, Glasgow G12 8QQ; e-mail campbell.leith@glasgow.ac.uk

1 Introduction

It is common practice when estimating DSGE models for use in monetary policy evaluation, to adopt a simple Taylor (1999)-type rule for monetary policy in order to be able to close the model for the purposes of estimation. However, such an approach raises a number of issues. In particular, partial equilibrium estimation of policy rules (see, for example, Clarida, Galí, and Gertler, 2000) suggests that their coefficients have changed over time, especially following the Volcker disinflation. While more multi-variate approaches also tend to require 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) or the nature of the underlying shock processes (Sims and Zha, 2006) to explain the evolution of inflation dynamics across time. Following Sims and Zha (2006) a large literature (see Section 8) then conducts counterfactual analyses which assess the extent to which the 'Great Moderation' in output and inflation volatility was good luck (a favorable shift in shock volatilities) or good policy (a desirable change in rule parameters and/or the implicit inflation target).

Moreover, although simple policy rules can often mimic optimal Ramsey policy, their estimated parameterization is often quite different from their optimized coefficients, (see, for example, Schmitt-Grohe and Uribe, 2007) which then begs the question why policy makers are implementing sub-optimal policy rules. This in turn, prompts the question as to whether or not policy makers actually have the ability to make the commitments assumed in standard Ramsey analyses of monetary policy, and whether they may, instead, be forced to pursue limited- or quasi-commitment policies, or even purely time-consistent discretionary policies with no commitment whatsoever. It is important to note that, just as failing to account for heteroscedasticity in the shock processes can bias parameter estimates (Sims and Zha, 2006), an inappropriate assumption about the degree of policy credibility can do so too (Erceg and Levin, 2003).

In this paper we examine all these issues by estimating a DSGE model of the US economy with, in contrast to the vast bulk of the literature, various descriptions of optimal policy rather than ad hoc simple rules. Specifically, we consider three basic forms of optimal policy: discretion, commitment and an intermediate case of imperfect commitment, variously termed, quasi- (Schaumburg and Tambalotti, 2007), loose- (Debertoli and Nunes, 2010) or limited- (Himmels and Kirsanova, 2013) commitment. Moreover, we allow for regime switches in the conservatism of central bank preferences, and switches in the volatility of shock processes hitting the economy.

We find that the description of policy most consistent with the data is that the Fed followed a (time-consistent) discretionary policy with a marked increase in conservatism following the Volcker disinflation. This dominates all other forms of optimal policy. This description of policy also dominates estimates based on simple rules with switches in shock variances and/or rule parameters or the inflation target. Taken together, this implies that the Fed is not making any credible policy commitments, either by following the Ramsey plan or following a simple rule. As with Sims and Zha (2006) we find that there is an increase in the anti-inflation stance of the US Fed around the time of the Volcker disinflation, and a reduction in the volatility of shocks in the early 1980s. However, we also find significant relaxation in the monetary policy stance in the aftermath of the stock market crash of October 1987, and the bursting of the dot-com bubble in 2000. Interestingly, while this policy stance was reversed early in 1990 as the Fed aggressively raised interest rates to stem the rise of inflation, our estimates suggest that monetary policy remained relatively accommodating following the dot-com crash all the way through to the financial crisis. That is, the Fed lost its conservatism following the dot-com crash and does not appear to have regained it.

Our estimates imply that, while most structural parameter estimates are robust to alternative descriptions of policy, some key parameter estimates obtained under simple rules are significantly different from those found under the various forms of optimal policy. This stems from the fact that optimal policy differentiates between shocks depending on their welfare implications, such that the inflationary consequences of taste and technology shocks are more aggressively offset than would be the case for cost-push shocks.[1] As a result, the first key difference between estimates based on optimal policy rather than simple rules, lies in the relative importance of cost-push shocks in creating meaningful policy trade-offs under all forms of optimal policy. However, although the estimated extent of cost-push shocks is significantly higher across all forms of optimal policy, the degree of habits varies significantly when considering discretionary policy. Therefore, the second key difference in parameter estimates is that a far smaller degree of habits are consistent with the policy response observed under discretion, relative to either simple rules or commitment. Essentially, under discretion the inability to precommit, means that offsetting the habits externality prompts an aggressive policy response as soon as a shock hits, which is not consistent with the data. In contrast, under commitment, policy is so effective that the habits externality is required in addition to pervasive cost-push shocks, to explain the volatility of inflation.

Finally, our counterfactual analysis suggests that the 'Great Moderation' in output and in-

[1] In fact, in the benchmark New Keynesian model without additional externalities such as those due to habits, the various forms of optimal policy would offset the inflation consequences of all shocks, other than cost-push shocks - see Leith, Moldovan, and Rossi (2012).

flation volatility is due to both a reduction in shock variances and an increase in central bank anti-inflation conservatism. However, decomposing the relative contribution of both effects implies that the far greater part of the welfare gains resulting from the Great Moderation stem from the reduction in shock volatilities. Additionally, despite the fact that Ramsey policies, which assume policy makers can improve key policy trade-offs by promising to behave in a particular way in the future, welfare-dominate all other descriptions of optimal policy and are typically the policy regime against which simple implementable rules are assessed, it appears that optimal policy *without* commitment is the data-preferred benchmark. We find that the potential gains from moving from discretion to commitment are substantial and dominate the gains from increasing central bank conservatism. In fact, our counterfactuals show that inflation would never have breached 2% in the 1970s had the policy maker had access to a commitment technology, *cet. par.* Accordingly, it appears that there is substantial scope to further improve monetary policy making by ensuring policy makers have access to commitment technologies and that they act to use such mechanisms.

The plan of the paper is as follows. Section 2 outlines our model, and the policy-maker's preferences. Our various descriptions of policy are discussed in Section 3. We then turn to consider the issues relating to the Bayesian estimation of our model in Section 4, and describe the data and priors in Section 5, before presenting our estimation results in Section 6. Section 7 then undertakes various counterfactual simulation exercises which enable us to explore both the sources and welfare consequences of the 'Great Moderation', but also assess the potential benefits of further improvements in the conduct of monetary policy. Section 8 relates our analyses to various strands of the literature on monetary policy. We then reach our conclusions in Section 9.

2 The Model

The economy is comprised of households, a monopolistically competitive production sector, and the government. There is a continuum of goods that enter the households' consumption basket. Households form external consumption habits at the level of the consumption basket as a whole - 'superficial' habits.[2] Furthermore, we assume the economy is subject to both price and inflation inertia. Both effects have been found to be important in capturing the hump-shaped responses of output and inflation to shocks evident in VAR-based studies, and are often employed in empirical

[2]For a comparison of the implications for optimal policy of alternative forms of habits see Leith et al. (2012).

applications of the New Keynesian model.[3]

2.1 Households

The economy is populated by a continuum of households, indexed by k and of measure 1. Households derive utility from consumption of a composite good, $C_t^k = \left(\int_0^1 \left(C_{it}^k\right)^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}$ where η is the elasticity of substitution between the goods in this basket and suffer disutility from hours spent working, N_t^k. Habits are both superficial and external implying that they are formed at the level of the aggregate consumption good, and that households fail to take account of the impact of their consumption decisions on the utility of others. To facilitate data-consistent detrending around a balanced growth path without restricting preferences to be logarithmic in form, we also follow Lubik and Schorfheide (2005) and An and Schorfheide (2007) in assuming that the consumption that enters the utility function is scaled by the economy wide technology trend, implying that household's consumption norms rise with technology as well as being affected by more familiar habits externalities. Accordingly, households derive utility from the habit-adjusted composite good,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{\left(C_t^k/A_t - \theta C_{t-1}/A_{t-1}\right)^{1-\sigma} (\xi_t)^{-\sigma}}{1-\sigma} - \frac{\left(N_t^k\right)^{1+\varphi} (\xi_t)^{-\sigma}}{1+\varphi} \right]$$

where $C_{t-1} \equiv \int_0^1 C_{t-1}^k dk$ is the cross-sectional average of consumption.[4] In other words households gain utility from consuming more than other households, and are disappointed if their consumption doesn't grow in line with technical progress and are subject to a time-preference or taste-shock, ξ_t. $\mathbb{E}_t$ is the mathematical expectation conditional on information available at time t, β is the discount factor $(0 < \beta < 1)$, and σ and φ are the inverses of the intertemporal elasticities of habit-adjusted consumption and work $(\sigma, \varphi > 0;\ \sigma \neq 1)$.

The process for technology is non-stationary,

$$\begin{aligned} \ln A_t &= \ln \gamma + \ln A_{t-1} + \ln z_t \\ \ln z_t &= \rho_z \ln z_{t-1} + \varepsilon_{z,t} \end{aligned}$$

[3] See for example Smets and Wouters (2003), Christiano, Eichenbaum, and Evans (2005) and Leith and Malley (2005).

[4] Note that this utility specification is slightly different from that in Lubik and Schorfheide (2005) who adopt the following specification, $(C_t - \theta\gamma C_{t-1})/A_t)^{1-\sigma} (\xi_t)^{-\sigma}/(1-\sigma)$. Their specification introduces a technology shock into the definition of habits adjusted consumption which then complicates the derivation of welfare. Therefore we adopt a specification which implies habits in detrended variables, which means that the only place the technology shock appears is in the consumption Euler equation.

Households decide the composition of the consumption basket to minimize expenditures, and the demand for individual good i is

$$C_{it}^k = \left(\frac{P_{it}}{P_t}\right)^{-\eta} C_t^k = \left(\frac{P_{it}}{P_t}\right)^{-\eta} \left(X_t^k + \theta C_{t-1}\right).$$

By aggregating across all households, we obtain the overall demand for good i as

$$C_{it} = \int_0^1 C_{it}^k dk = \left(\frac{P_{it}}{P_t}\right)^{-\eta} C_t. \tag{1}$$

Remainder of the Household's Problem The remainder of the household's problem is standard. Specifically, households choose the habit-adjusted consumption aggregate, $X_t^k = C_t^k/A_t - \theta C_{t-1}/A_{t-1}$, hours worked, N_t^k, and the portfolio allocation, D_{t+1}^k, to maximize expected lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{\left(X_t^k\right)^{1-\sigma} (\xi_t)^{-\sigma}}{1-\sigma} - \frac{\left(N_t^k\right)^{1+\varphi} (\xi_t)^{-\sigma}}{1+\varphi}\right]$$

subject to the budget constraint

$$\int_0^1 P_{it} C_{it}^k di + E_t Q_{t,t+1} D_{t+1}^k = W_t N_t^k (1-\tau_t) + D_t^k + \Phi_t + T_t$$

and the usual transversality condition. The household's period-t income includes: wage income from providing labor services to goods producing firms, $W_t N_t^k$, which is subject to a time-varying tax rate, τ_t , dividends from the monopolistically competitive firms, Φ_t, and payments on the portfolio of assets, D_t^k. Financial markets are complete and $Q_{t,t+1}$ is the one-period stochastic discount factor for nominal payoffs. Lump-sum transfers, T_t, are paid by the government. The tax rate, τ_t, will be used to finance lump-sum transfers, and can be designed to ensure that the long-run equilibrium is efficient in the presence of the habits and monopolistic competition externalities. However, we shall assume that the tax rate fluctuates around this efficient level such that it is responsible for generating an autocorrelated cost-push shock. Finally there is an autocorrelated preference shock, ξ_t.

In the maximization problem, households take as given the processes for C_{t-1}, W_t, Φ_t, and T_t, as well as the initial asset position D_{-1}^k. The first order conditions for labor and habit-adjusted consumption are

$$\frac{\left(N_t^k\right)^{\varphi}}{\left(X_t^k\right)^{-\sigma}} = \frac{W_t}{P_t A_t}(1-\tau_t)$$

and

$$Q_{t,t+1} = \beta \left(\frac{X_{t+s}^k \xi_{t+1}}{X_t^k \xi_t} \right)^{-\sigma} \frac{A_t}{A_{t+1}} \frac{P_t}{P_{t+1}}.$$

Taking expectations, the Euler equation for consumption can be written as

$$1 = \beta \mathbb{E}_t \left[\left(\frac{X_{t+1}^k \xi_{t+1}}{X_t^k \xi_t} \right)^{-\sigma} \frac{A_t}{A_{t+1}} \frac{P_t}{P_{t+1}} \right] R_t,$$

where $R_t^{-1} = \mathbb{E}_t \left[Q_{t,t+1} \right]$ denotes the inverse of the risk-free gross nominal interest rate between periods t and $t+1$.

2.2 Firms

We further assume that intermediate goods producers are subject to the constraints of Calvo (1983)-contracts such that, with fixed probability $(1-\alpha)$ in each period, a firm can reset its price and with probability α the firm retains the price of the previous period, but where, following Yun (1996) that price is indexed to the steady-state rate of inflation. When a firm can set the price, it can either do so in order to maximize the present discounted value of profits, $\mathbb{E}_t \sum_{s=0}^{\infty} \alpha^s Q_{t,t+s} \Phi_{it+s}$, or it can follow a simple rule of thumb as in (Galí and Gertler, 1999, or Leith and Malley, 2005). The constraints facing the forward looking profit maximizers are the demand for their own good (1) and the constraint that all demand be satisfied at the chosen price. Profits are discounted by the s-step ahead stochastic discount factor $Q_{t,t+s}$ and by the probability of not being able to set prices in future periods.

$$\begin{aligned}
& \max_{\{P_{it},\ Y_{it}\}} \mathbb{E}_t \sum_{s=0}^{\infty} \alpha^s Q_{t,t+s} \left[\left(P_{it} \pi^s - MC_{t+s} \right) Y_{it+s} \right] \\
s.t. Y_{it+s} &= \left(\frac{P_{it} \pi^s}{P_{t+s}} \right)^{-\eta} Y_{t+s} \\
\text{where } Q_{t,t+s} &= \beta^s \left(\frac{X_{t+1} \xi_{t+1}}{X_t \xi_t} \right)^{-\sigma} \frac{P_t}{P_{t+s}}.
\end{aligned}$$

The relative price set by firms able to reset prices optimally in a forward-looking manner, satisfies the following relationship

$$\frac{P_t^f}{P_t} = \frac{\eta}{\eta - 1} \frac{\mathbb{E}_t \sum_{s=0}^{\infty} (\alpha\beta)^s \left(X_{t+s} \xi_{t+s} \right)^{-\sigma} mc_{t+s} \left(\frac{P_{t+s} \pi^{-s}}{P_t} \right)^{\eta} \frac{Y_{t+s}}{A_{t+s}}}{\mathbb{E}_t \sum_{s=0}^{\infty} (\alpha\beta)^s \left(X_{t+s} \xi_{t+s} \right)^{-\sigma} \left(\frac{P_{t+s} \pi^{-s}}{P_t} \right)^{\eta-1} \frac{Y_{t+s}}{A_{t+s}}}, \tag{2}$$

where $mc_t = MC_t/P_t$ is the real marginal cost and P_t^f denotes the price set by all firms who are able to reset prices in period t and choose to do so in a profit maximizing way.

In addition to the familiar Calvo-type price setters, we also allow for inflation inertia. To do so we allow some firms to follow simple rules of thumb when setting prices. Specifically, when a firm is given the opportunity of posting a new price, we assume that rather than posting the profit-maximizing price (2), a proportion of those firms, ζ, follow a simple rule of thumb in resetting that price

$$P_t^b = P_{t-1}^* \pi_{t-1}, \tag{3}$$

such that they update there price in line with last period's rate of inflation rather than steady-state inflation, where P_{t-1}^* denotes an index of the reset prices given by

$$\ln P_{t-1}^* = (1-\zeta)\ln P_{t-1}^f + \zeta P_{t-1}^b.$$

P_t represents the price level at time t. With α of firms keeping last period's price (but indexed to steady-state inflation) and $(1-\alpha)$ of firms setting a new price, the law of motion of this price index is,

$$(P_t)^{1-\eta} = \alpha\left(P_{t-1}\pi\right)^{1-\eta} + (1-\alpha)\left(P_t^*\right)^{1-\eta}.$$

Denoting the fixed share of price-setters following the rule of thumb (3) by ζ, we can derive a price inflation Phillips curve, as detailed in Leith and Malley (2005). For this we combine the rule of thumb of price setters with the optimal price setting described above, leading to the price Phillips curve

$$\widehat{\pi}_t = \chi_f \beta \mathbb{E}_t \widehat{\pi}_{t+1} + \chi_b \widehat{\pi}_{t-1} + \kappa_c(\widehat{mc}_t),$$

where $\widehat{\pi}_t = \ln(P_t) - \ln(P_{t-1}) - \ln(\pi)$ is the deviation of inflation from its steady state value, $\widehat{mc}_t = \ln(W_t/P_t) - \ln A_t - \ln((\eta-1)/\eta)$, are log-linearized real marginal costs, and the reduced form parameter convolutions are defined as $\chi_f \equiv \alpha/\Phi$, $\chi_b \equiv \zeta/\Phi$, $\kappa_c \equiv (1-\alpha)(1-\zeta)(1-\alpha\beta)/\Phi$, with $\Phi \equiv \alpha(1+\beta\zeta) + (1-\alpha)\zeta$.

2.3 The Government

The government collects a distortionary tax on labor income which it rebates to households as a lump-sum transfer. The steady-state value of this distortionary tax will be set at a level which offsets the combined effect of the monopolistic competition distortion and the effects of the habits externality, as in Levine, McAdam, and Pearlman (2008), see Appendix B. However, shocks to the tax rate described by

$$\ln(1-\tau_t) = \rho^\mu \ln(1-\tau_{t-1}) + (1-\rho^\mu)\ln(1-\tau) - \varepsilon_t^\mu$$

serve as autocorrelated cost-push shocks to the NKPC. There is no government spending per se. The government budget constraint is given by

$$\tau_t W_t N_t = -T_t.$$

2.4 The Complete Model

The complete system of non-linear equations describing the equilibrium are given in Appendix A. Log-linearizing the equilibrium conditions (21) - (34) around the deterministic steady state detailed in the Appendix, gives the following set of equations:

$$\sigma \widehat{X}_t + \varphi \widehat{N}_t = \widehat{w}_t - \widehat{\mu}_t \text{ Labor Supply} \tag{4}$$

$$\widehat{X}_t = \mathbb{E}_t \widehat{X}_{t+1} - \frac{1}{\sigma}\left(\widehat{R}_t - \mathbb{E}_t \widehat{\pi}_{t+1} - \mathbb{E}_t \widehat{z}_{t+1}\right) - \widehat{\xi}_t + \mathbb{E}_t \widehat{\xi}_{t+1} \text{ Euler Equation} \tag{5}$$

$$\widehat{y}_t = \widehat{N}_t = \widehat{c}_t \text{ Resource Constraint} \tag{6}$$

$$\widehat{X}_t = (1-\theta)^{-1}(\widehat{c}_t - \theta \widehat{c}_{t-1}) \text{ Habits-Adjusted Consumption} \tag{7}$$

$$\widehat{\pi}_t = \chi_f \beta \mathbb{E}_t \widehat{\pi}_{t+1} + \chi_b \widehat{\pi}_{t-1} + \kappa_c(\widehat{w}_t), \text{ Hybrid NKPC} \tag{8}$$

$$\widehat{z}_t = \rho^z \widehat{z}_{t-1} + \varepsilon_{z,t} \text{ Technology Shock} \tag{9}$$

$$\widehat{\mu}_t = \rho^\mu \widehat{\mu}_{t-1} + \varepsilon_t^\mu \text{ Cost-Push Shock} \tag{10}$$

$$\widehat{\xi}_t = \rho^\xi \widehat{\xi}_{t-1} + \varepsilon_t^\xi \text{ Preference Shock} \tag{11}$$

where $\widehat{\mu}_t = \tau \widehat{\tau}_t / (1-\tau)$ represents autocorrelated fluctuations in the labor income tax rate which serves as a cost-push shock. The model is then closed through the addition of one of the descriptions of policy considered in Section 3.

2.5 Objective Function

Since we wish to assess the empirical implications of assuming policy is described by various forms of optimal policy rather than a simple rule we need to define the policy maker's objectives. Appendix C derives an objective function based on the utility of the households populating the economy as

$$L = -\frac{1}{2}\overline{N}^{1+\varphi}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \begin{array}{c} \frac{\sigma(1-\theta)}{1-\theta\beta}\left(\widehat{X}_t + \widehat{\xi}_t\right)^2 + \varphi\left(\widehat{y}_t - \frac{\sigma}{\varphi}\widehat{\xi}_t\right)^2 \\ + \frac{\alpha\eta}{(1-\beta\alpha)(1-\alpha)}\left(\widehat{\pi}_t^2 + \frac{\zeta\alpha^{-1}}{(1-\zeta)}\left[\widehat{\pi}_t - \widehat{\pi}_{t-1}\right]^2\right) \end{array} \right\} + tip + O[2] \tag{12}$$

which shall underpin the optimal policy estimation and analysis. Therefore, rather than adopt an ad hoc objective function defined in terms of output and inflation, we have an objective function which is fully consistent with the underlying model and which accounts for habits externalities, and both price level and inflation inertia. As a result the objective function contains dynamics in output and inflation.

When assuming optimal policy within the estimation, we shall assume that the policy maker possesses an objective function of this form, but where the weights on the various terms are freely estimated. This can capture the fact that the conservatism of the central bank differs from that of the representative household. Below, we shall contrast these estimated objective function weights with those of the representative household, given the estimated structural parameters of the model, in addition to assessing how the households' evaluation of the welfare implications of policy differs from that of the policy maker.

3 Policy

We consider four basic forms of policy, a simple rule and three types of optimal policy (discretion, commitment and quasi-commitment), to close our model when undertaking the estimation. We shall also allow for Markov switching in rule parameters, the inflation target, as well as the relative weight given to inflation under optimal policy.

3.1 Simple Rule Specification

When US monetary policy is described as a generalized Taylor rule, we specify this rule following An and Schorfheide (2007),

$$R_t = \rho^R R_{t-1} + (1-\rho^R)[\psi_1 \widehat{\pi}_t + \psi_2(\Delta\widehat{y}_t + \widehat{z}_t)] + \varepsilon_t^R \tag{13}$$

where the Fed adjusts interest rates in response to movements in inflation and deviations of output growth from trend.[5] When considering the simple rule without any switches in rule parameters we shall assume that $\psi_1, \psi_2 \geq 0$, such that our rule is determinate, and that the smoothing term in the rule is: $0 \leq \rho^R < 1$. Subsequently, equations (4)-(11) and (13) can be written as a linear rational expectation system of the form

$$\Gamma_0 X_t = \Gamma_1 X_{t-1} + \Psi Z_t + \Pi \eta_t \tag{14}$$

[5] It should be noted that 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, once we allow for Markov switching in either shocks or policy, discretionary policy is "decisively" preferred by the data.

where $X_t = [\widehat{z}_t, \widehat{\mu}_t, \widehat{\xi}_t, \widehat{y}_t, \widehat{\pi}_t, \widehat{R}_t, \mathbb{E}_t\widehat{y}_{t+1}, \mathbb{E}_t\widehat{\pi}_{t+1}]'$ is a vector of eight state variables, which includes six predetermined variables and two non-predetermined, or jump, variables. Vector Z_t stacks the exogenous shocks and η_t is composed of rational expectation forecast errors. Γ_0, Γ_1, Ψ and Π are matrices containing structural parameters. A standard solution technique, such as Sims (2002), can be used to solve the linear rational expectation system in equation (14). It returns a solution as a reduced AR(1) process.

$$X_t = \Phi_1 X_{t-1} + \Phi_2 Z_t, \;\; Z_t \sim NID(\mathbf{0}, \mathbf{\Sigma})$$

Within the framework of a generalized Taylor rule, we further account for potential changes in US monetary policy by allowing for either changes in the Fed's inflation target or rule parameters. In the former case the measure of excess inflation in the Taylor rule, $\widehat{\pi}_t$, involves removing the inflation target from the data, where, following Schorfheide (2005), we allow that inflation target to follow a two-state Markov-switching process. Modelling monetary policy changes as movements in the inflation target are not computationally demanding as Sims (2002) algorithm can still be employed to solve this model.[6]

However, when the policy changes are described as shifts in rule parameters (ρ^R, ψ_1, ψ_2) between two regimes, standard solution techniques are no longer applicable. Therefore, Svensson and Williams (2007), Davig and Leeper (2007) and Farmer, Waggoner, and Zha (2008, 2009, 2011) all provide algorithms to solve DSGE models with Markov-switches in structural parameters.[7] In this paper, we adopt the procedure developed by Farmer et al. (2008) to solve the model with Markov-switching in simple rule parameters. This model can be recast in the following system

$$\Gamma_0(S_t = j)X_t = \Gamma_1(S_t = j)X_{t-1} + \Psi(S_t = j)Z_t + \Pi(S_t = j)\eta_t. \tag{15}$$

Compared to the time-invariant interest rate rule contained in equation (14), Γ_0, Γ_1, Ψ and Π in (15) depend on an unobserved state variable, $S_t = j$, for $j \in \{1, 2\}$, that follows a two-state Markov process with transition probabilities

$$\Pr[S_t = 1|S_{t-1} = 1] = p_{11}, \Pr[S_t = 2|S_{t-1} = 2] = p_{22}.$$

Following Farmer et al. (2008), equation (15) can be rewritten as the following model with regime-invariant parameters

$$\overline{\Gamma}_0 X_t = \overline{\Gamma}_1 X_{t-1} + \overline{\Psi} Z_t + \overline{\Pi}\eta_t, \tag{16}$$

[6] The details of this model can be found in Schorfheide (2005).
[7] Chen and MacDonald (2012) provide a discussion of these algorithms.

where $\overline{\Gamma}_0, \overline{\Gamma}_1, \overline{\Psi}$ and $\overline{\Pi}$ are matrices that are functions of structural parameters and transition probabilities. Farmer et al. (2008) define a Minimum State Variable (MSV) solution to equation (16) and prove that it is also a solution to the original MSRE model specified in equation (15). Provided a unique solution exists, equation (16) can be written as an AR(1) process with Markov-switching parameters

$$X_t = \Phi_1(S_t = j)X_{t-1} + \Phi_2(S_t = j)Z_t, Z_t \sim NID(\mathbf{0}, \mathbf{\Sigma}).$$

It is important to note that the estimated rule in a particular state need not satisfy the Taylor principle, $\psi_1 > 1$, and that this need not imply indeterminacy provided the rule in alternative states is sufficiently responsive to inflation. This 'spillover' from one regime to another reflects the fact that economic agents are assumed to anticipate the Markov switching between different policy rules.

In addition to incorporating monetary policy changes, we also account for the 'good luck' factor that is normally modelled as a decrease in the volatility of shocks hitting the economy. Therefore, we allow for independent regime switching in the variances, Σ, of four shocks (i.e. $\sigma_z, \sigma_\mu, \sigma_\zeta$ and σ_R) that depends on the unobserved state variable, $s_t = i$, for $i \in \{1, 2\}$, and has the transition probabilities:

$$\Pr[s_t = 1|s_{t-1} = 1] = q_{11}, \Pr[s_t = 2|s_{t-1} = 2] = q_{22}$$

This results in a four-state transition matrix and $4^2 = 16$ states are carried at each iteration.

3.2 Optimal Monetary Policy

We now turn to describe our optimal monetary policy specifications. Relative to the number of models estimated with various simple rules, the empirical studies based on optimal policies are few, and tend to only focus on optimal policies under two polar extremes: full commitment or discretion. As with these studies (i.e. Givens, 2012; Le Roux and Kirsanova, 2013), we performed an empirical estimation based on optimal policies derived under either full commitment or discretion, but we also considered the intermediate case of quasi-commitment. To compute optimal policies, we recast the set of log-linearized equations in (4)-(11) the following state-space form

$$\begin{bmatrix} \boldsymbol{X}_{t+1} \\ \mathbb{E}_t \boldsymbol{x}_{t+1} \end{bmatrix} = \boldsymbol{A} \begin{bmatrix} \boldsymbol{X}_t \\ \boldsymbol{x}_t \end{bmatrix} + \boldsymbol{B} \boldsymbol{i}_t + \begin{bmatrix} \boldsymbol{C} \\ \mathbf{0} \end{bmatrix} \boldsymbol{\varepsilon}_{t+1}, \tag{17}$$

where $\boldsymbol{X}_t = [\widehat{z}_t, \widehat{\mu}_t, \widehat{\xi}_t, \widehat{y}_{t-1}, \widehat{\pi}_{t-1}, \widehat{R}_{t-1}]'$ is a vector of predetermined variables; $\boldsymbol{x}_t = [\widehat{y}_t, \widehat{\pi}_t]'$ is a vector of forward-looking variables; $\boldsymbol{i}_t = [\widehat{R}_t]$ is the control variable, and $\boldsymbol{\varepsilon}_t = [\varepsilon_t^z, \varepsilon_t^\mu, \varepsilon_t^\xi]$ contains a vector of zero mean *i.i.d.* shocks. Without loss of generality, the shocks are normalized so that the covariance matrix of $\boldsymbol{\varepsilon}_t$ is $\boldsymbol{I}$. Therefore, the covariance matrix of the shocks to $\boldsymbol{X}_{t+1}$ is $\mathbf{CC}'$. $\boldsymbol{A}$ and $\boldsymbol{B}$ are matrices containing the model's structural parameters. The central bank selects interest rates to maximize objective (12) subject to (17). We use the procedure described by Söderlind (1999) to solve for the equilibrium dynamics under both commitment and discretion.

However, estimation with micro-founded weights is problematic. Since the micro-founded weights are functions of structural parameters, they place very tight cross-equation restrictions on the model which are generally thought to be implausible. In particular, for standard estimates of the degree of price stickiness, the microfounded weight attached to inflation can be over 100 times that attached to the output terms (see Woodford, 2003, Ch.6). Optimal policies which were based on such a strong anti-inflation objective would clearly be inconsistent with observed inflation volatility. Therefore, for estimation, we adopt a form of the objective function which is consistent with the representative agents' utility, but allow the weights within that objective function in (12) to be freely estimated, and the resulting objective function is given by

$$\Gamma = -\overline{N}^{1+\varphi}\frac{1}{2}\mathbb{E}_0\sum_{t=0}^{\infty}\beta^t\left\{\omega_1\left(\widehat{X}_t+\widehat{\xi}_t\right)^2+\omega_2\left(\widehat{y}_t-\frac{\sigma}{\varphi}\widehat{\xi}_t\right)^2+\omega_\pi\widehat{\pi}_t^2+\omega_3\left[\widehat{\pi}_t-\widehat{\pi}_{t-1}\right]^2\right\}, \quad (18)$$

where the weight on inflation, ω_π, is normalized to 1. It is important to note that we do not augment our objective function with any ad hoc terms, such as a desire for interest rate smoothing, that are not implied by the underlying model. This facilitates an exploration of the policy implications of the estimated weights differing from the micro-founded weights.

Under full commitment, the central bank chooses a contingent interest rate plan for all future dates. When optimizing, the central bank internalizes the impact of its policies on the private sector's expectations. By being able to influence expectations through future policy commitments the policy maker can obtain a more favorable trade-off between the stabilization of inflation and output. The solution to the commitment problem is as follows

$$\begin{aligned}\begin{bmatrix}\boldsymbol{X}_{t+1}\\ \boldsymbol{\psi}_{t+1}\end{bmatrix} &= \boldsymbol{M}_c\begin{bmatrix}\boldsymbol{X}_t\\ \boldsymbol{\psi}_t\end{bmatrix}+\begin{bmatrix}\boldsymbol{C}\\ \mathbf{0}\end{bmatrix}\boldsymbol{\varepsilon}_{t+1}\\ \begin{bmatrix}\boldsymbol{x}_t\\ \boldsymbol{i}_t\end{bmatrix} &= \boldsymbol{G}_c\begin{bmatrix}\boldsymbol{X}_t\\ \boldsymbol{\psi}_t\end{bmatrix},\end{aligned}$$

where $\boldsymbol{\psi}_t$ is a vector of Lagrangian multipliers associated with forward-looking variables. The fact that the choice of interest rates depends on $\boldsymbol{\psi}_t$ implies that the central bank is assumed, under commitment, to honor past promises.

In contrast, under discretion, the central bank is not bound by its past promises. Therefore, in each period, it evaluates the current state of the economy and formulates the optimal policy. The policy outcome under discretion is only optimal in a constrained sense because the central bank can neither control the private sector's expectations by making promises about the policies that will be implemented in the future, nor coordinate with future policy makers.[8] Therefore, policy only depends on the current state

$$\begin{aligned} \boldsymbol{X}_{t+1} &= M_d \boldsymbol{X}_t + \boldsymbol{C}\varepsilon_{t+1} \\ \begin{bmatrix} \boldsymbol{x}_t \\ \boldsymbol{i}_t \end{bmatrix} &= \boldsymbol{G}_d \boldsymbol{X}_t. \end{aligned} \tag{19}$$

Given that much of the literature on estimated policy rules finds that there have been significant changes in the conduct of policy over time, we realize that both commitment and discretion policies derived under an assumption of unchanging policy maker preferences may be too stylized to capture such changes. Therefore, in our empirical analysis we attempt to relax the assumption of a time-invariant objective function that is used to derive optimal policies under both commitment and discretion. To do so, we adopt the algorithm developed by Svensson and Williams (2007) that solves optimal monetary policies in Markov jump-linear-quadratic systems. This algorithm can incorporate structural changes in both the model (17) and weights in the objective function (18).[9] However, in this paper, we only focus on potential changes in the US monetary policy objective on inflation targeting. Specifically, we allow the weight on inflation, ω_π, to be subject to regime shifting between 1 and a value lower than 1. By doing so, we can identify whether there are periods where the Fed has adopted different attitudes towards inflation at different points in time. For example, was there a more conservative monetary policy since the Volcker disinflation? Were the lower interest rates observed during 2001-2007 due to economic conditions, or were they the result of the Fed putting less emphasis on inflation targeting relative to its other objectives? Svensson and Williams (2007)'s 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

[8] In our model with endogenous state variables, due to habits formation and inflation inertia, current policies will influence future expectations through their impact on the states bequeathed to the future. However, crucially, under discretion the policy maker cannot make any additional commitments in the hope of favorably influencing expectations.

[9] The algorithm used to solve the Markov-jump linear quadratic system is described in Svensson and Williams (2007). We focus on the scenario where no learning occurs and the central bank and private agents can observe the different monetary policy regimes.

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 Debertoli and Nunes (2010) would be applicable.

Furthermore, we also consider an intermediate case of quasi-commitment. Schaumburg and Tambalotti (2007), Debertoli and Nunes (2010) and Himmels and Kirsanova (2013) all provide theoretical discussions of this description of policy. Under quasi-commitment, the policy maker deviates from full commitment-based plans with a fixed probability (which is known by the private sector). Effectively, the policy maker forms a commitment plan which they will adhere to until randomly ejected from office. At which point a new policy maker will be appointed, and a new plan formulated (based on the same objective function) until that policy maker is, in turn, removed. Therefore, the central bank can neither completely control the expectations of the private sector, nor can she perfectly coordinate the actions of all future policy makers. This framework incorporates elements of both discretion and commitment. Specifically, we follow Himmels and Kirsanova (2013) in recasting the quasi-commitment of Schaumburg and Tambalotti (2007) and Debertoli and Nunes (2010) in a general linear-quadratic form which can be solved using standard iterative techniques, such as Söderlind (1999). The optimization problem under quasi-commitment can be expressed by the following Lagrangian:

$$\min \mathbb{E}_0 \sum_{t=0}^{\infty} ((1-\upsilon)\beta)^t \left\{ \begin{array}{c} \left[\begin{array}{c} \omega_1 \left(\widehat{X}_t + \widehat{\xi}_t\right)^2 + \omega_2 \left(\widehat{y}_t - \frac{\sigma}{\varphi}\widehat{\xi}_t\right)^2 \\ +\omega_\pi \widehat{\pi}_t^2 + \omega_3 \left(\widehat{\pi}_t - \widehat{\pi}_{t-1}\right)^2 \end{array} \right] \\ +\upsilon\beta \boldsymbol{X}'_{t+1}\boldsymbol{S}\boldsymbol{X}_{t+1} \end{array} \right\}$$

subject to

$$\left[\begin{array}{c} \boldsymbol{X}_{t+1} \\ (1-\upsilon)\mathbb{E}_t \boldsymbol{x}_{t+1} + \upsilon \boldsymbol{H}\boldsymbol{X}_{t+1} \end{array} \right] = \boldsymbol{A} \left[\begin{array}{c} \boldsymbol{X}_t \\ \boldsymbol{x}_t \end{array} \right] + \boldsymbol{B}\boldsymbol{i}_t + \left[\begin{array}{c} \boldsymbol{C} \\ \boldsymbol{0} \end{array} \right] \varepsilon_{t+1}$$

where $0 \leq \upsilon \leq 1$ is the probability that the monetary authority reneges on the past policy promises at each period. The solution to the quasi-commitment problem is given by

$$\begin{aligned} \left[\begin{array}{c} \boldsymbol{X}_{t+1} \\ \boldsymbol{\psi}_{t+1} \end{array} \right] &= \boldsymbol{M}_{qc} \left[\begin{array}{c} \boldsymbol{X}_t \\ \boldsymbol{\psi}_t \end{array} \right] + \left[\begin{array}{c} \boldsymbol{C}\varepsilon_{t+1} \\ \boldsymbol{0} \end{array} \right] \\ \left[\begin{array}{c} \boldsymbol{x}_t \\ \boldsymbol{i}_t \end{array} \right] &= \boldsymbol{G}_{qc} \left[\begin{array}{c} \boldsymbol{X}_t \\ \boldsymbol{\psi}_t \end{array} \right], \end{aligned}$$

where $\boldsymbol{M}_{qc}$ and $\boldsymbol{G}_{qc}$ are partitioned with $\boldsymbol{X}_t$ and $\boldsymbol{\psi}_t$ as follows

$$\boldsymbol{M}_{qc} = \left[\begin{array}{cc} \boldsymbol{M}_{XX} & \boldsymbol{M}_{X\psi} \\ \boldsymbol{M}_{\psi X} & \boldsymbol{M}_{\psi\psi} \end{array} \right], \ \boldsymbol{G}_{qc} = \left[\begin{array}{cc} \boldsymbol{G}_{xX} & \boldsymbol{G}_{x\psi} \\ \boldsymbol{G}_{iX} & \boldsymbol{G}_{i\psi} \end{array} \right].$$

If $\upsilon = 1$, the monetary authority re-optimizes its policy every period and the resulting transition matrix is equivalent to the optimal policy under discretion as in (19). On the contrary, if $\upsilon = 0$, the monetary authority will keep its promises and this is essentially a policy problem with full commitment. In general, a value of υ close to 1 implies that the policy maker is forced to take inflationary expectations as under the discretionary policy problem, while υ close to zero means that she can make (partial) promises over future policy actions which have a beneficial impact on expectations, as under commitment.

The solutions of the quasi-commitment problem can be easily combined with a two-state Markov-switching model to identify the periods in which the policy maker reneges on previous plans before embarking on a new quasi-commitment policy, such that

$$\boldsymbol{M}_{qc} = \begin{bmatrix} \boldsymbol{M}_{XX} & \boldsymbol{M}_{X\psi}(S_t = j) \\ \boldsymbol{M}_{\psi X}(S_t = j) & \boldsymbol{M}_{\psi\psi}(S_t = j) \end{bmatrix}, \; \boldsymbol{G}_{qc} = \begin{bmatrix} \boldsymbol{G}_{xX} & \boldsymbol{G}_{x\psi}(S_t = j) \\ \boldsymbol{G}_{iX} & \boldsymbol{G}_{i\psi}(S_t = j) \end{bmatrix}$$

where the unobserved state variable, S_t, follows a two-state Markov process with transition probabilities

$$\Pr[S_t = 1 | S_{t-1} = 1] = (1 - \upsilon), \;\; \Pr[S_t = 2 | S_{t-1} = 2] = \upsilon.$$

If $S_t = 2$, elements in matrices $\boldsymbol{M}_{X\psi}, \boldsymbol{M}_{\psi X}, \boldsymbol{M}_{\psi\psi}$, $\boldsymbol{G}_{x\psi}$ and $\boldsymbol{G}_{i\psi}$ reflecting the Lagrange multipliers associated with forward-looking variables switch to zero indicating that the monetary policy authority breaks its promises.

Finally, as with the model with the simple rules, we allow for independent regime switching in variances of shocks under optimal policy, i.e. σ_z, σ_μ, and σ_ζ. This is to account for the 'good luck' factor and to obtain more reliable parameter estimates by avoiding the biases associated with the heteroscedastic errors that would emerge if such shifts in shock volatility were not accounted for.

Therefore, to summarize, we consider four basic forms of policy: simple rules, commitment, discretion and quasi-commitment. We also allow for Markov switches in the variances of the shock processes and, in the case of rules, switches in the inflation target or rule parameters, as well as changes in the degree of central bank conservatism under both optimal discretionary and commitment policies. We use three data series in estimation: output, inflation and interest rates. When considering optimal policy there are three shock processes for technology, preferences and cost-push shocks. While in the case of simple rules there is an additional shock to the interest rate rule.[10]

The next section will discuss our estimation strategy. However, before doing so it is important to note that all model parameters are identifiable. To demonstrate this, we used the Iskrev (2010)

[10] Adding an additional shock to the interest rate under optimal policy does not materially affect the results.

local identification test for our models based on a simple rule as well as optimal policy under both commitment and discretion.

4 Estimation Strategy

For estimation, the recursive equations derived under simple rule and optimal policy are linked to the observed variables through a measurement equation specified as:

$$\begin{bmatrix} \Delta GDP_t \\ INF_t \\ INT_t \end{bmatrix} = \begin{bmatrix} \gamma^Q + \Delta\widehat{y}_t + \widehat{z}_t \\ \pi^A + 4\widehat{\pi}_t \\ r^A + \pi^A + 4\gamma^Q + 4\widehat{R}_t \end{bmatrix}$$

The observed variables are quarterly output growth (ΔGDP_t), annualized domestic inflation (INF_t) and the nominal interest rate (INT_t). The parameters, γ^Q, π^A and r^A represent the values of output growth, inflation and interest rates when the economy is in its steady state. For the simple rule with a Markov-switching inflation target, π^A, shifts between a low and a high target, while for the other models, γ^Q, π^A and r^A remain time-invariant. These parameters will all be estimated as part of the model estimation.

We adopt the Bayesian approach in estimating all our models. For models with Markov-switching parameters, the posterior distribution is obtained through Bayes theorem

$$p(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T|\mathbf{Y}^T) = \frac{p(\mathbf{Y}^T|\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T)p\left(\boldsymbol{S}^T|\boldsymbol{\phi}\right)p\left(\boldsymbol{\phi}, \boldsymbol{\theta}\right)}{\displaystyle\int p(\mathbf{Y}^T|\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T)p\left(\boldsymbol{S}^T|\boldsymbol{\phi}\right)p\left(\boldsymbol{\phi}, \boldsymbol{\theta}\right)d\left(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T\right)} \tag{20}$$

where $p\left(\boldsymbol{\phi}, \boldsymbol{\theta}\right)$ is the prior for the structural parameters, $\boldsymbol{\theta}$, and the transition probabilities, $\boldsymbol{\phi}$**.** $p\left(\boldsymbol{S}^T|\boldsymbol{\phi}\right)$ is the prior for the unobserved states and $p(\mathbf{Y}^T|\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T)$ is the likelihood function. Since it is difficult to characterize the posterior distribution in equation (20), we follow Schorfheide (2005) to factorize the joint posterior as

$$p(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{S}^T|\mathbf{Y}^T) = \boldsymbol{p}\left(\boldsymbol{\theta}, \boldsymbol{\phi}|\mathbf{Y}^T\right)p\left(\boldsymbol{S}^T|\boldsymbol{\theta}, \boldsymbol{\phi}, \mathbf{Y}^T\right).$$

Due to the presence of Markov-switching parameters, the likelihood function is approximated using Kim (1994)'s filter, and then combined with the prior distribution to obtain the posterior distribution. Sims (2002) optimization routine CSMINWEL is used to find the posterior modes of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. The inverse Hessian is then calculated at these posterior modes and is used as the covariance matrix of the proposal distribution. It is scaled to yield a target acceptance rate of 25%-40%. We adopt Schorfheide (2005)'s strategy that employs a random walk Metropolis-Hastings algorithm to generate 500,000 draws from $\boldsymbol{p}\left(\boldsymbol{\theta}, \boldsymbol{\phi}|\mathbf{Y}^T\right)$, with the first 200,000 draws being discarded and save every 20th draw from the remaining draws.[11] Conditional on the saved draws of parameter vectors, $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, we then utilized Kim (1994) smoothing algorithm to generate draws from the history of unobserved states, $\boldsymbol{S}^T$. Posterior means are obtained by Monte-Carlo averaging.

Finally, we compute the log marginal likelihood values for each model to provide a coherent framework to compare models with different types of monetary policies. We first implement the commonly used modified harmonic mean estimator of Geweke (1999) for this task. We also utilize the approach of Sims, Waggoner, and Zha (2008) as a robustness check. The latter is designed for models with time-varying parameters, where the posterior density may be non-Gaussian.

5 Data and Model Priors

5.1 Data

Our empirical analysis uses the US data on output growth, inflation, nominal interest rates from 1961Q1 up to 2008Q3, just before nominal interest rates were reduced to their effective lower bound of 0.5% and the first round of quantitative easing was implemented. All data are seasonally adjusted and at quarterly frequencies. Output growth is the log difference of real GDP, multiplied by 100. Inflation is the log difference of the associated implict price deflator, scaled by 400. All data are taken from the FRED database.[12] The data used in the estimation are plotted in Figures 4-6, alongside various counterfactual simulation results which will be discussed below.

5.2 Priors

The priors are presented in Table 1. These are set to be broadly consistent with the literature on the estimation of New Keynesian models. For example, the mean of the Calvo parameter, α, is set so that average length of the contract is around one year. Following Smets and Wouters (2003), we choose the normal distribution for inverse of the Frisch labor supply elasticity, φ, and the inverse of the intertemporal elasticity of substitution, σ, with both priors having a mean of 2.5. Habits formation, indexation and the AR(1) parameters of the technology, cost-push, and taste shock processes are assumed to follow a beta distribution with a mean of 0.5 and a standard deviation of 0.15. It is important to note that the above priors are common to all model variants.

[11] Geweke (1992) convergence diagnostics indicate that convergence is achieved. These are available upon request.

[12] The specific data series used are the Effective Federal Funds Rate - FEDFUNDS, Gross Domestic Product in United States-USARGDPQDSNAQ and the Gross Domestic Product: Implicit Price Deflator-GDPDEF.

In addition, variances of shocks are chosen to be highly dispersed inverted Gamma distributions to generate realistic volatilities for the endogenous variables. In models that allow for Markov-switching shock processes, the priors for shock variances are set to be symmetric across regimes.

Furthermore, for models featuring a simple rule, we use comparatively loose priors for the policy rule parameters that are consistent with An and Schorfheide (2007) in the case of time-invariant simple rule. As for Markov-switching rule parameters, in line with Bianchi (2012), 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.[13] For optimal policy, the relative weights (i.e. ω_1, ω_2, ω_3) on the objective function are assumed to be distributed following beta distributions and ω_π is normalized to 1 in a time-invariant objective function. In the case where ω_π is allowed to switch between 1 and a value lower than 1, 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 a beta distribution with a mean of 0.3 and standard deviation of 0.02, implying a prior belief that the frequency of policy re-optimizations lies between 8 months and one year. Loosening this prior tends to push the estimated parameter closer to one, such that the quasi-commitment policy reduces to that of discretion. Maintaining a tight prior enables us to explore the implications of describing policy as being a form of quasi-commitment and facilitates a comparison of episodes of re-optimization with other policy switches.

The prior means γ^Q, π^A and r^A are set to be broadly consistent with the average output growth rate and inflation rate during this pre-sample period from 1950Q1 to 1960Q4. For the model with Markov-switching inflation target, the priors for the inflation targets are set in line with Schorfheide (2005). Finally, the average real interest rate, r^A, is linked to the discount factor, β, such that $\beta = \left(1 + r^A/400\right)^{-1}$.

6 Results

In this section we present the results of our estimation. Since the majority of the literature estimating DSGE models relies on a simple rule without Markov switching in either rule parameters or the model's shock processes, we begin by contrasting such an approach to one where we allow

[13] This way of setting priors for the switching parameters is also discussed by Davig and Doh (2009), to introduce a natural ordering of regime-dependent parameters and to avoid the potential risk of 'label switching' as noted in Hamilton, Waggoner, and Zha (2007).

policy to be described by one of the forms of optimal policy outlined above - namely discretion, commitment and quasi-commitment. However, as noted by Clarida et al. (2000) and others, there is evidence that there have been changes in the estimated rule parameters, which may correspond to shifts in monetary policy regime, such as that begun with the Volcker disinflation. Accordingly, our second set of estimates allows for such policy regime shifts, both in terms of changes in the inflation target or rule parameters, as well as the variation in the weight given to the inflation target under our various forms of optimal policy. Finally, Sims and Zha (2006) argue that estimates of shifts in the parameters of simple rules may be biased in the presence of heteroscedastic errors and that much of the 'Great Moderation' can be assigned to a favorable shift in the variance of shocks (i.e. 'good luck'), rather than a systematic change/improvement in the conduct of monetary policy ('good policy'). Therefore, we conclude our set of estimated results by allowing for both changes in policy and the volatility of shocks processes.

6.1 Simple Rules and Optimal Policy

In this subsection we contrast results when monetary policy is described by an inertial Taylor rule for interest rates, with those obtained when policy is based on one of the notions of optimality, discretion, commitment or quasi-commitment.[14] The posterior means and the 90% confidence intervals obtained from estimating time-invariant models are presented in Table 2 where each column corresponds to an alternative policy description, and these columns are ordered according to log marginal likelihood values calculated using Geweke (1999) and Sims et al. (2008), respectively.[15] The first column of results in Table 2 is for the best-fit model, which in this case utilizes the simple interest rate rule, followed by time-consistent discretionary policy then quasi-commitment, and finally, commitment. In terms of marginal data densities, the simple rule, discretion and quasi-commitment are relatively similar, while describing policy with optimal commitment leads to a clear deterioration in explanatory power. Table 2 also reports the Bayes

[14]We have also estimated these policy variants after allowing for Markov switching in the volatility of the shock processes. This does not affect the ranking or qualitative results in Table 2. These results are available upon request.

[15]It is important to note that across all estimations the posterior distributions differ from the prior. The possible exception to this is the estimate of the persistence of the cost-push shock, when policy is described by one of the variants of the simple rule. Nevertheless, the application of the Iskrev (2010) local identification test to the model based on simple rules (and optimal policy under discretion and commitment), is supportive of identification at the central parameter estimates also in this case. In applying the Iskrev (2010) test we examine the rank of the Jacobian of a vector of model parameters across 10,000 draws from the prior distribution, as well as at the prior and posterior means. Plots contrasting prior and posterior distributions are available upon request (and are contained in Figure 1 for our data-preferred estimates based on discretionary policy with Markov switching in conservatism and shock variances).

Factors for each model relative to the first model in the Table. In this case, using Kass and Raftery (1995) adaption of Jeffreys (2007) criteria for quantifying the evidence in favor of one model rather than another, the evidence in favor of simple rules over discretion is only "substantial", while evidence in favor of discretion relative to commitment is "decisive".[16] The probability of reneging on policy promises under the quasi-commitment policy is identified, as its posterior mean, $\upsilon = 0.31$, is slightly shifted up from its prior, but the posterior distribution is slightly more dispersed around its mean. Given a quarterly data period, this estimate implies that the quasi-commitment plan is expected to implemented for about one year. Looser priors for this parameter raise the estimates of υ as the data seek to reduce the quasi-commitment policy to its special case of discretion. Therefore, this preliminary set of results does not suggest that there is a significant degree of commitment within US monetary policy.

This then begs the question why are the data apparently inconsistent with policy under commitment, when the rhetoric of central banks would suggest that making credible promises is their *raison d'etre*? We can provide an explanation by exploring the differences in estimates of both the structural model parameters and the shock process hitting our estimated economies as we vary the description of policy. If we consider estimates obtained under the conventional inertial interest rate rule, then our results are broadly in line with other studies: an intertemporal elasticity of substitution, σ, of 2.8; a measure of price stickiness, $\alpha = 0.78$, implying that price contracts typically last for one year; a relatively modest degree of price indexation, $\zeta = 0.1$, a sizeable estimate of the degree of habits, $\theta = 0.82$ and an inverse Frisch labor supply elasticity of $\varphi = 2.4$. Moving from these estimates obtained under a conventional interest rate rule to the case of optimal policy under discretion, these deep parameter estimates remain largely the same, except that there is a sizeable decline in the degree of habits in the model, which falls to $\theta = 0.48$, and a modest increase in the degree of indexation in price setting to $\zeta = 0.16$. At the same time, the simple rule relies on taste shocks (both in terms of size and persistence) to explain the volatility in the data, while policy under discretion significantly raises the estimated persistence of cost-push shocks in order to fit the data. These differences in estimates across models where policy is described by optimal (but time-consistent) policy rather than an ad hoc rule, reflects the nature of the optimal policy problem. In the absence of inflation inertia and habits, the model would reduce to the benchmark New Keynesian model considered by Woodford (2003) where only cost-push shocks generate a meaningful trade-off for the monetary policy maker, as

[16] Following Jeffreys (2007), Kass and Raftery (1995) argue that values of the Bayes Factor associated with two models lying between 0 and 3.2 constitutes evidence which is "not worth more than a bare mention", between 3.2 and 10 is "substantial" evidence, between 10 and 100 is "strong" evidence and above 100 is "decisive" evidence.

monetary policy can optimally respond to technology and taste shocks without generating any inflation. Adding inflation indexation to pricing contracts creates further policy trade-offs (see Steinsson, 2003), as does the externality associated with habits formation (see Leith et al., 2012), both of which would imply that the inflation consequences of taste and technology shocks are no-longer perfectly offset by monetary policy. Accordingly, in order to replicate the observed fluctuations in inflation, optimal policy must be faced with meaningful trade-offs which prevent it from perfectly stabilizing inflation, and the estimated degree of habits, price indexation and cost-push shocks provide the most data coherent means of doing so.

As we increase the level of commitment the policy maker can achieve, then this alters the trade-offs facing the policy maker further. Although we have constructed an economy that doesn't experience any inflationary bias problem, any inability to fully commit implies a stabilization bias (Svensson, 1997). This additional bias arises as the policy maker acting without commitment cannot make credible promises as to how she will behave in the future which improves the policy trade-offs she faces today. When the policy maker can commit, she will make promises which enable her to stabilize inflation at a lower cost in terms of fluctuations in real variables. Therefore, in order for commitment policy to be consistent with the observed fluctuations in inflation, the parameter estimates significantly raise both the degree of inflation indexation ($\zeta = 0.59$), the extent of habits, $\theta = 0.64$ as well as the variance ($\sigma_\mu^2 = 3.628$) and persistence ($\rho^\mu = 0.968$) of cost-push shocks.[17] Therefore, although commitment introduces an inertia to policy which may have been thought useful in explaining the data, it is, in fact, simply too effective in stabilizing inflation to be the data-preferred description of policy.[18]

6.2 Changes in Policy Regime

Log marginal likelihood values for the above time-invariant models suggest that a simple rule could explain the data as well as, but not much better than, optimal discretionary policy. In addition, within alternative descriptions of optimal policy there was "decisive" evidence against the hypothesis that policy makers could not precommit. In this subsection we relax the assumption that a single description of policy was in place at all points in time. Instead, we shall allow the simple rule to vary either through changes in the inflation target over time, or through changes in rule parameters. We shall then contrast this policy with those under Svensson and Williams

[17] 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 0.1-0.3 across our estimates.

[18] We shall turn to evaluate the gains to both optimal policy in general, and commitment in particular in Section 7.

(2007)-type commitment and discretion when the weight attached to inflation stabilization may also vary over time. These switches in inflation targets, rule parameters and policy objective weights are intended to capture the changes in policy regime at the US Fed discussed by Clarida et al. (2000) and others. Again we present the estimates across these alternative descriptions of policy, ordered according to decreasing goodness of fit, see Table 3.

There is a clear improvement in fit when monetary policy is allowed to switch between regimes. Moreover, the ranking of policy descriptions has altered, with the optimal policy under discretion (with Markov switching in the weight attached to inflation in the objective function) dominating all other descriptions of policy, including those rule-based estimations which allow for shifts in rule parameters or the inflation target or which assume commitment with switches in policy objective weights. It is also important to note that the Bayes Factors associated with these estimates now imply that the evidence in favor of the estimates being based on discretionary policy rather than rules or commitment is "decisive" using the interpretation of Kass and Raftery (1995).

Under optimal policy when the weight attached to inflation stabilization is allowed to switch across regimes, the posterior mean of ω_π in the second regime reduces to 0.347 under discretion and 0.302 under commitment. As for simple rules, for the model with Markov-switching rule parameters, the posterior mean of the rule coefficient on inflation, ψ_1, switches between 2.075 and 0.909, with no overlap in the confidence intervals across the two regimes. In contrast, the differences between other policy parameters (ψ_2, ρ^R) over the two regimes are less apparent. The estimates of the inflation coefficient, ψ_1, across the two regimes correspond to 'active' and 'passive' monetary rules (Leeper, 1991), where the former satisfies the Taylor principle by raising nominal interest rates by more than the increase in excess inflation, *cet. par.*, while the latter fails to do so. It should be noted that this model remains determinate as economic agents anticipate moving between the two policy regimes, despite the mild passivity of the rule in the second regime.[19] Similarly, when we consider monetary policy changes as shifts in the inflation target, $\widehat{\pi}_t$ incorporates the jumps in the estimated inflation target which varies between 4.23% and 6.06%, respectively and the rule remains active. All these results suggest that the Fed's stance on inflation targeting has varied over the studied sample period.

Figure 2 plots the smoothed probabilities of being in the 'weak inflation targeting regime' within our four policy variants, as well as the probabilities that a policy maker acting under quasi-commitment has reneged on his past policy promises, but retained the same policy objectives.

[19] In fact, the coefficient on inflation in the passive rule could fall as low as $\psi_1 = 0.1$ and the model would still remain determinate, *cet. par.*.

By 'weak inflation targeting regime' we mean those regimes where the inflation target is higher, the degree of conservatism lower or the interest rate rule more passive/less active. Although there are some slight differences in timings of changes in regime, there is a clear prevalence of the weak inflation targeting regime in the 1970s. Accordingly, what Clarida et al. (2000) find as a passive monetary policy rule in the 1970s is also clear from our results. However, it can also be captured by an increase in the inflation target in the rule, particularly around the time of the two oil price shocks, or by a decline in the weight attached to inflation stabilization under either commitment or discretion at similar points in time. Under quasi-commitment there is also a clear rise in the extent of estimated policy re-optimizations in the 1970s which is broadly consistent with the timings of the changes in regime suggested by the other modelling devices.[20] We shall see below that the Markov switching in policy will be more informative in respect of the conduct of monetary policy in recent years when combined with switches in shock volatility.

Consistent with previous results obtained from estimating time-invariant model variants, structural parameters are similar across discretionary policy and both forms of simple rule, although with a lower habits and slightly higher inflation indexation under discretion relative to simple rules. Additionally, the estimates under commitment, tend to partially recover the higher estimates of habits and the degree of indexation, as well as increasing the size and persistence of the cost-push shocks for the reasons discussed above. Therefore, allowing for policy regime changes generates results consistent with policy makers engaging in optimal policy, but without possessing (or, at least, utilizing) any commitment technologies.

6.3 The Great Moderation

Much of the debate on the observed moderation in the volatility of both output and inflation since the mid 1980s has focused on whether or not this was due to 'good luck' or 'good policy'. We now contribute to this debate by widening the description of policy from a simple rule to optimal policy under either commitment or discretion. The results, again ordered by decreasing log marginal likelihood, are presented in Table 4. Allowing for separate Markov switching processes for policy and shock volatility further improves the fit of all model variants. Describing policy as being

[20] Debortoli and Lakdawala (2013) also find evidence that policy re-optimizations occur predominantly in 1970s. However, our results differ slightly from those of Debortoli and Lakdawala (2013) in that they find policy re-optimizations to be relatively infrequent. This may reflect the larger scale model they consider, different priors, alternative detrending techniques or, the importance assigned to interest rate smoothing in their estimated objective function.

optimal, but time consistent, remains clearly preferred by the data[21] - with the evidence in favor of discretion relative to simple rules with time varying parameters being classed as “strong”, and relative to commitment as “decisive”. Once more, the data simply does not like the implications of commitment policy, and seeks to raise the degree of habits, inflation indexation and the magnitude and persistence of cost-push shocks above conventional estimates to ensure that the model can fit the data, particularly in terms of inflation volatility. This inability of commitment policies to explain the data as effectively as policy under discretion suggests that there may be significant gains to commitment - an issue we explore in the next section.

It is important to note that once we account for time-varying volatility of shocks, differences across monetary policy regimes are narrowed in all policy variants. In particular, the mean posterior of ω_π in the second regime rises to 0.436 under discretion and 0.373 under commitment. While, for the simple rule with Markov-switching rule parameters, ψ_1 in the second regime does not actually turn passive.[22] In addition, for the simple rule with a Markov-switching inflation target, the difference between inflation targets across regimes is also reduced. These results may imply that at least part of the ‘Great Moderation’ was due to ‘good luck’ - an issue that we consider below. It also suggests that estimates of policy changes which fail to account for heteroscedasticity in the error processes may contain significant biases which overstate the extent to which policy did not appear to successfully target inflation in the 1970s, see Sims and Zha (2006).

Examining the periods when either the high volatility or relatively accommodating monetary policies are estimated to be in place is also more informative than when only policy was allowed to switch. Figure 3 plots the smoothed probabilities of being in the weak inflation targeting regime, as well as being in the high volatility regime. It also includes a plot for the case where the policy maker acts under quasi-commitment, but faces Markov switching in shock volatilities. Compared to the smoothed probabilities obtained from rule-based models, these obtained from our best-fit model based on discretionary policy with switches in the weight attached to inflation in the objective function, provide additional information on the conduct of monetary policy over recent years. Aside from the usual relaxation of monetary policy in the 1970s, the smoothed probabilities from this model also suggest that policy was relaxed briefly following the stock market crash of

[21] Plots of the prior and posterior distributions in Figure 1 indicate that almost all parameters are well identified, with the possible exception of the inverse of the Frisch elasticity of labor supply, φ, where the posterior distribution has not changed significantly relative to its prior.

[22] This is in line with Sims and Zha (2006) who also find that policy was not passive in the 1970s. However, they argue that this was because their model includes money and that policy in the 1970s is better described by a money growth rule than an interest rate rule.

October 1987. More interestingly, a prolonged reduction in the Fed's weight attached to inflation stabilization is identified from the dot-com crash all the way through to the financial crisis. Such a pattern is not so apparent in the rule based models.

7 Counterfactuals

Our best-fit model is obtained under discretionary policy with Markov switching in the weight on inflation stabilization 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 US Fed adopted a tougher anti-inflation stance in the 1970s. Moreover, we can explore how much further welfare 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.

We begin our counterfactuals by analyzing the role of good luck in stabilizing US output and inflation. To do so we fixed the pattern of switches in policy regimes to those estimated from the data, but consider the counterfactual where the volatility of shocks is either at its high or low value. We take the estimated shocks and re-scale them by the relative standard deviations from the high and low volatility regimes, so that similar kinds of shock are imposed, but they are scaled to mimic the standard deviations observed under the two volatility regimes. Figure 4 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 the low shock volatility regime during the 1970s and early 1980s. Moreover, we should also note that under the policy regimes estimated post-Volcker, 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 in stabilizing the US economy.

Therefore, in the second set of counterfactual analysis, we assess the impact that increased conservatism would have had on US inflation and output, especially during the 1970s. To do so, we reinsert the estimated shock processes back in the model and fix the weight on inflation in the objective function, ω_π, to either 1 or 0.436 throughout the sample period. Figure 5 plots the actual and counterfactual series for inflation and interest rates, as well as output losses with

$\omega_\pi = 1$ for the entire sample period. Output losses are the difference between model implied output with estimated objective function weights and the counterfactual output when policy maker is more conservative, $\omega_\pi = 1$. The first panel of Figure 5 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 the output losses as shown in Panel 3 of Figure 5. 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 real measure of the monetary policy stance is captured 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.

Finally, in Figure 6 we keep the shock volatility and policy switches fixed at their estimated values, and vary whether or not the policy-maker has access to a commitment technology. That is, we assess the implications of moving from discretion to commitment, *cet. par.* The results are striking. If the US Fed had been able to make credible policy commitments in the 1970s, even although it was subject to high volatility shocks and reduced the weight attached to inflation stabilization 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. However, our welfare analysis below, suggests that these losses are more than compensated for by the reduction in inflation volatility.

In addition to providing counterfactual figures, we also compute the unconditional variances of key variables, as well as the value of unconditional welfare (both using the policy maker's estimated weights and the measure of utility that would be consistent with the estimated structural parameters) under alternative counterfactuals. Following Bianchi (2012), we use the unconditional variances of key variables (and the associated welfare losses) computed under the worst case scenario as the benchmark case for the 'good luck' versus 'good policy' debate. That is our benchmark adopts the high shock volatility regime in conjunction with discretionary policy with the lower level of estimated conservatism, $\omega_\pi = 0.436$. We can then consider the extent to which 'good policy' or 'good luck' alone would be able to stabilize inflation, output and interest rates. Tables 5 and 6 show that under discretion either an increase in central bank conservatism or reduction in shock volatility alone would reduce more than half of the volatility in inflation and interest rates implied by the worst case scenario. However, it is the ‘good luck’ that could lead to significant output stabilization and therefore achieve bigger gains in welfare.

Turning to the second half of Tables 5 and 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 can allow central banks 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 whether the estimated increase in central bank conservatism took place. 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’ we 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 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 Links to Existing Monetary Policy Literature

Early estimates of structural DSGE models utilized a variety of econometric techniques, including minimum distance estimators (Rotemberg and Woodford, 1997, and Christiano et al., 2005), maximum likelihood (Fuhrer and Moore, 1995, and Ireland, 2001), and GMM (Leith and Malley, 2005; Clarida et al., 2000). However, following Smets and Wouters (2003)’s successful estimation of a DSGE model for the Euro-area, it has become increasingly common to estimate such models using Bayesian techniques.[23] There is now a vast literature on the subject and we have to be very selective in our discussion of the existing literature. Our research has links to the literature on simple rules with Markov switching, on optimal monetary policy under discretion and commitment, on switching policy objectives, and on the implications for optimal policy of alternative sources of model inertia. We therefore, restrict our discussion to these four topics.

8.1 Estimating DSGE Models with Simple Policy Rules

Typically the literature estimating DSGE models for monetary policy analysis adopts a simple interest rate rule as a description of policy in order to close the model and achieve determinacy prior to estimation. However, the partial equilibrium analysis of estimated policy rules has often

[23] For an early survey see An and Schorfheide (2007), and more recently Schorfheide (2011) and Guerron-Quintana and Nason (2012). Textbook treatments include, Canova (2007) and DeJong and Dave (2007).

found there to be significant breaks in estimated policy rules with the rules often switching from a passive to an active parameterization around the time of the Volcker disinflation (see, for example, the estimates in Clarida et al., 2000).[24] Such breaks have also been found in rules estimated within a DSGE framework, whether, following, Lubik and Schorfheide (2005) they assume a particular equilibrium selection device for periods of indeterminacy, or allow for a long-run analysis of determinacy where economic agents anticipate random switches between policy regimes (see Davig and Leeper, 2007 and Farmer et al., 2008, 2009, 2011).[25]

Sims and Zha (2006), however, point out that the original Clarida et al. (2000) analysis and other estimates of changes in policy rules are potentially subject to bias to the extent that errors are heteroscedastic, and found that allowing for switches in the volatility of shock processes significantly reduced the importance of changes in policy rule parameters in explaining inflation dynamics. These results led Sims and Zha (2006) to conclude that the 'Great Moderation' in output and inflation volatility was largely due to good luck rather than a desirable change in policy. There have been many subsequent papers attempting to explore this issue but all of these have relied on simple policy rules as their benchmark description of policy. Within this literature we can find papers which utilize structural VARs (Primiceri, 2005; Benati and Surico, 2009; Cogley et al., 2010) and others DSGE models. Within the latter category, there are papers which allow for breaks in policy rules (Davig and Doh, 2009; Bianchi, 2012) or inflation targets (for example, Schorfheide, 2005; Liu et al., 2011) often with breaks in the volatility of shock processes, and occasionally changes in the estimated degree of nominal inertia (Eo, 2009, and for the UK see Liu and Mumtaz, 2011, and Chen and MacDonald, 2012).

8.2 Estimating DSGE Models with Optimal Monetary Policy

Our paper describes monetary policy as optimal, alongside possible Markov switching in shock volatilities and central bank conservatism. Here the literature is not nearly as extensive as in the case of a rule-based description of policy, although there have been some papers attempting to estimate models using a description of optimal policy rather than a simple rule. For example,

[24] The Taylor principle requires monetary policy to adjust nominal interest rates such that real interest rates rise in response to an increase in inflation relative to its target. An active interest rate rule satisfies this requirement, while a passive one does not. An active monetary policy rule is typically seen as a sufficient condition for determinacy in the benchmark New Keynesian model (see Woodford, 2003, Ch.2).

[25] Castelnuovo (2012) uses an alternative equilibrium selection device to that of Lubik and Schorfheide (2005) assumption that impulse responses be continuous when moving from the boundary of the determinate to indeterminate policy space and finds that this raises the uncertainty in assigning the Great Moderation to either of the usual explanations. While Eo (2009) contrasts the the Lubik and Schorfheide (2005) and Davig and Leeper (2007) approaches.

Ozlale (2003) and Favero and Rovelli (2003) both estimate the Rudebusch and Svensson (1997) model and find structural breaks in the central bank's preference parameters around the time of Volcker's appointment. However, as this model is purely backward-looking this does not offer an opportunity to discriminate between various forms of optimal policy, particularly, the question as to whether or not the policy maker is able to precommit. Dennis (2004) implicitly assumes that policy makers cannot precommit, by considering the case of discretionary (time-consistent) policy and also finds a break in policy-maker preferences. In contrast, using data for Sweden Adolfson, Laseen, Linde, and Svensson (2011) find that policy under commitment is more able to fit the data than a simple rule. Ilbas (2010) estimates the Euro-area model of Smets and Wouters (2003) under the assumption that policy is optimal and policy makers can commit. It is notable that these papers only consider one form of optimal policy. Papers which do consider both discretion and commitment include Givens (2012) for the US and Le Roux and Kirsanova (2013) for the UK. Both find that discretion fits the data better than commitment, but the difference in marginal data density is not very large. Coroneo et al. (2012) explore the explanatory power of discretion and commitment using moment inequalities, which reject the model under commitment but not discretion.

In contrast to these papers we consider not only the cases of commitment and discretion, but allow for the possibility of an intermediate case of imperfect commitment given by the case of 'quasi-commitment', as described in Schaumburg and Tambalotti (2007) and Debertoli and Nunes (2010). Independently of the current paper, Debortoli and Lakdawala (2013) consider the cases of commitment and quasi-commitment, finding that the majority of re-optimizations occured in the 1970s. Moreover, since shifts in shock volatilities and/or simple rule parameters have been found to be important in the literature exploring the 'Great Moderation' we allow for Markov switching in shock volatilities and the degree of central bank conservatism when considering the alternative forms of optimal policy.[26]

8.3 Markov Switching in Policy Objectives

While much of the analysis of regime switches in policy has focussed on switches in estimated rule parameters and/or inflation targets, there is some research exploring changes in policy maker preferences. Debertoli and Nunes (2013) point out that since breaks in policy rule parameters

[26] In addition to the Markov switching models considered above, some papers also focus on learning in explaining inflation and output dynamics surrounding monetary policy shifts. However, Schorfheide (2005) finds that Bayesian posterior odds favour the full-information variant of their model. Therefore, we did not explore this issue further in the current paper.

may simply imply that policy is moving towards (or away from) the efficient frontier, treating such estimated breaks as reflecting a break in policy maker preferences may be inappropriate. Accordingly, Debertoli and Nunes (2013) allow for a policy maker to be randomly ejected from office, before being replaced with another policy maker with different preferences. In such a situation the policy maker commits policy as far into the future as they reasonably can given their expected termination from office, but also seeks to anticipate the policies pursued by their replacement. Effectively, they are playing a game against the future policy maker. While such a description of changing policy maker preferences may be appropriate when considering monetary policy under the control of partisan politicians who alternate in office,[27] it may not reflect the policy implications of changes in Fed policy objectives which may be thought more to follow the evolution of monetary theory and practice. Therefore, we followed Svensson and Williams (2007) who develop tools for analyzing optimal monetary policies in Markov jump-linear-quadratic systems with the aim of allowing policy to target distributional forecasts rather than point estimates of key macro variables. Within their general framework, which applies the recursive saddlepath method of Marcet and Marimon (2011), they similarly can allow for breaks in policy maker preferences, but where, implicitly, current policymakers anticipate such breaks and do not undertake current policies with the aim of thwarting their future selves. Another paper exploring breaks in policy maker preferences, although only for the case of discretion, is Blake and Zampolli (2011).

8.4 The Role of Inertia

While models utilizing simple rules often rely on devices such as habits and inflation inertia to achieve the hump-shaped response to monetary policy shocks found in the data, the impact of such devices on the conduct of optimal policies depends on their associated externalities. Amato and Laubach (2004) and Leith et al. (2012) both consider the optimal policy response to habits, where the former paper assumes habits to be internal (i.e. households anticipate the impact of current consumption on their future utility via habits effects) and the latter external (households fail to internalize the impact of their consumption decisions on the future utility of themselves and others). As a result, in the former case the presence of habits does not impose any additional policy trade-offs for the policy maker, while the latter does. Moreover, the policy response to this externality differs depending on whether or not the policy maker can commit or not. Similarly the presence of inflation inertia will affect the desirability of price level control when the policy maker can commit. As shown in Woodford (2003) in the context of the benchmark

[27] As in, for example, the UK prior to the Bank of England being granted independence in 1997.

New Keynesian model, the promise to not only stabilize inflation under commitment, but the price level itself significantly improves the trade-off between output and inflation stabilization. Such a commitment is less desirable in the presence of inflation inertia, see Steinsson (2003). Finally, the optimal policy response to shocks will depend on the nature of the shocks: in the simple benchmark New Keynesian model, the inflationary consequences of technology and taste shocks will be aggressively offset by monetary policy, such that it is only cost-push shocks which generate meaningful policy trade-offs. In the presence of habits externalities this is no longer the case (see Leith et al., 2012).

Accordingly, the differing responses to various shocks across the different descriptions of optimal policy in the presence of the additional frictions caused by habits and inflation inertia, were exploited by the estimation procedure in explaining the data. This ability to differentiate between habits and cost-push shocks as the source of policy trade-offs in estimation would not be so apparent when policy is described by an estimated rule whose parameters are not necessarily in any way optimal.

9 Conclusions

In this paper we explored the implications of describing policy using various notions of optimal policy, namely discretion, commitment and quasi-commitment, when estimating a DSGE model of the US economy. In contrast to conventional estimates which assume that policy can be described by a simple, ad hoc, rule, estimates under optimal policy tend to imply a different mixture of shocks, habits formation and inflation inertia in explaining the data. This reflects the fact that optimal policy will aggressively stabilize inflation unless it faces a meaningful policy trade-off either because of externalities associated with features such as consumption habits or inefficient cost-push shocks. Although we permitted the policy rules to experience switches in the implicit inflation target or parameters of the rule, and allowed for switches in the volatility of the shock processes hitting the economy, our best-fit model implies that policy was best described as being conducted under discretion, with an increase in central bank conservatism following the Volcker disinflation period. This description of policy, also implied 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. Analysis of Bayes Factors suggests that there is “strong” evidence in favor of this description of policy relative to simple rules, and “decisive” evidence relative to optimal policy under commitment.

Based on estimates from our best-fit model, we undertake a range of counterfactual simulations which throw light on various aspects of policy. Firstly, we find that 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. Secondly, we were able to explore the implications of being able to commit, rather than acting under discretion. Here we found that the gains to commitment were significant and dominate the degree of central bank conservatism. In the 1970s the average rate of inflation would have been below 2% had the Fed had the ability to commit regardless of the level of conservatism. Taken together, this suggests that attempts to improve monetary policy outcomes should concentrate on ensuring central banks are 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., S. Laseen, J. Linde, and L. E. Svensson (2011). Optimal Monetary Policy in an Operational Medium-Sized DSGE Model. *Journal of Money, Credit and Banking 43*(7), 1287–1331.

Amato, J. D. and T. Laubach (2004). Implications of habit formation for optimal monetary policy. *Journal of Monetary Economics 51*(2), 305–325.

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

Benati, L. and P. Surico (2009). VAR Analysis and the Great Moderation. *American Economic Review 99*(4), 1636–1652.

Bianchi, F. (2012). Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics. *Review of Economic Studies*. Forthcoming.

Blake, A. P. and F. Zampolli (2011). Optimal Policy in Markov Switching Rational Expectations Models. *Journal of Economic Dynamics and Control 35*(10), 1626–1651.

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

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

Canova, F. (2007). *Methods for Applied Macroeconomic Research.* Oxford: Princeton University Press.

Castelnuovo, E. (2012). Policy Switch and the Great Moderation: The Role of Equilibrium Selection. *Macroeconomic Dynamics 16*(3), 449–471.

Chen, X. and R. MacDonald (2012). Realised and Optimal Monetary Policy Rules in an Estimated Markov-Swicthing DSGE Model of the United Kingdom. *Journal of Money, Credit and Banking 44*(6), 1091–1116.

Christiano, L., M. Eichenbaum, and G. Evans (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy 113*(1), 1–45.

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

Cogley, T., G. E. Primiceri, and T. J. Sargent (2010). Inflation-Gap Persistence in the US. *American Economic Journal: Macroeconomics 2*(1), 43–69.

Coroneo, L., V. Corradi, and P. Santos Monteiro (2012). Testing the degree of commitment via moment inequalities. The Warwick Economics Research Paper Series (TWERPS) No. 985.

Davig, T. and T. Doh (2009). Monetary policy regime shifts and inflation persistence. Federal reserve Bank of Kansas City, RWP 08-16.

Davig, T. and E. Leeper (2007). Generalizing the Taylor principle. *American Economic Review 97*(3), 607–35.

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

Debertoli, D. and R. Nunes (2013). Monetary regime switches and unstable objectives. Board of Govenors of the Federal Reserve System, International Finance Discussion Paper No. 1036.

Debortoli, D. and A. K. Lakdawala (2013). How Credible is the Federal Reserve? A Structural Estimation of Policy Re-Optimizations. University of San Diego, mimeo.

DeJong, D. N. and C. Dave (2007). *Strcutural Macroeconometrics*. Oxford: Princeton University Press.

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

Eo, Y. (2009). Bayesian Analysis of DSGE Models with Regime Switching. Mimeo, Washington University in St. Louis.

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

Eser, F., C. Leith, and S. Wren-Lewis (2009). When is monetary policy all we need? University of Oxford Working Paper No.430.

Farmer, R. E. A., D. F. Waggoner, and T. Zha (2008). Minimal State Variable Solutions to Markov-switching Rational Expectations Models. Bank of Atlanta Working Paper 2008-23.

Farmer, R. E. A., D. F. Waggoner, and T. Zha (2009). Understanding Markov-switching Rational Rxpectations Models. *Journal of Economic Theory 144*, 1849–1867.

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

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

Fuhrer, J. C. and G. R. Moore (1995). Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output. *American Economic Review 85*(1), 219–239.

Galí, J. and M. Gertler (1999). Inflation Dynamics: A Structural Econometric Analysis. *Journal of Monetary Economics 44*, 195–222.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In J. O. Berger, J. M. Bernardo, A. P. Dawid, and A. F. M. Smith (Eds.), *Proceedings* *of the Fourth Valencia International Meetings on Bayesian Statistics*, pp. 169–194. Oxford: Oxford University Press.

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.

Guerron-Quintana, P. A. and J. A. Nason (2012). Bayesian Estimation of DSGE Models. Federal Reserve Bank of Phildelphia, Working Paper No. 12-4.

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

Himmels, C. and T. Kirsanova (2013). Escaping Expectation Traps: How Much Commitment is Required? *Journal of Economic Dynamics and Control 37*(3), 649–665.

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. (2001). Sticky-Price Models of the Business Cycle: Specification and Stability. *Journal of Monetary Economics 46*, 3–18.

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.

Jeffreys, H. (2007). *Theory of Probability* (3 ed.). Oxford: Oxford University Press.

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

Kim, C.-J. (1994). Dynamic linear models with Markov-switching. *Journal of Econometrics 60*, 1–22.

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

Leeper, E. M. (1991). Equilibria Under 'Active' and 'Passive' Monetary and Fiscal Policies. *Journal of Monetary Economics 27*, 129–147.

Leith, C. and J. Malley (2005). Estimated General Equilibrium Models for the Analysis of Monetary Policy in the US and Europe. *European Economic Review 49*(8), 2137–2159.

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

Levine, P., P. McAdam, and J. Pearlman (2008). Quantifying and sustaining welfare gains from monetary commitment. *Journal of Monetary Economics 55*, 1253–1276.

Liu, P. and H. Mumtaz (2011). Evolving Macroeconomic Dynamics in a Small Open Economy: An Estimated Markov Switching DSGE Model for the UK. *Journal of Money, Credit and Banking 43*(7), 1443–1474.

Liu, Z., D. F. Waggoner, and T. Zha (2011). Sources of Macroeconomic Fluctuations: A Regime-Switching DSGE Approach. *Quantitative Economics 2*(2), 251–301.

Lubik, T. and F. Schorfheide (2005). A Bayesian Look at the New Open Economy Macroeconomics. In *NBER Macroeconomics Annual 2005*, pp. 313–379.

Marcet, A. and R. Marimon (2011). Recursive Contracts. CEP Discussion Papers DP1055, Centre for Economic Performance, LSE.

Ozlale, U. (2003). Price stability vs. output stability: tales of Federal Reserve administrations. *Journal of Economic Dynamics and Control 27*(9), 1595–1610.

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

Rotemberg, J. J. and M. Woodford (1997). An Optimization-based Econometric Framework for the Evaluation of Monetary Policy. In *NBER Macroeconomics Annual*, pp. 297–344.

Rudebusch, G. and L. E. Svensson (1997). Policy Rules for Inflation Targeting. In *Monetary Policy Rules*, pp. 203–262.

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

Schmitt-Grohe, S. and M. Uribe (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.

Schorfheide, F. (2011). Estimation and Evaluation of DSGE Models: Progress and Challenges. NBER Working Paper No. 16781.

Sims, C. A. (2002). Solving Linear Rational Expectations Models. *Computational Economics 20*(1-2), 1–20.

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

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

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

Söderlind, P. (1999). Solution and Estimation of RE Macromodels with Optimal Policy. *European Economic Review 43*, 813–823.

Steinsson, J. (2003). Optimal Monetary Policy in an Economy with Inflation Presistence. *Journal of Monetary Economics 50*, 1425–1456.

Svensson, L. (1997). Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts. *The American Economic Review 87*(1), 98–114.

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

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

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton, NJ.: Princeton University Press.

Yun, T. (1996). Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics 37*(2-3), 345–370.

A The Complete Model

The complete system of non-linear equations describing the equilibrium are given by

$$
\begin{aligned}
N_t^{\varphi}\left(\frac{X_t}{A_t}\right)^{\sigma} &= \frac{W_t}{A_t P_t}(1-\tau_t) \equiv w_t(1-\tau_t) && (21)\\
\left(\frac{X_t}{A_t}\right)^{-\sigma}\xi_t^{-\sigma} &= \beta \mathbb{E}_t\left[\left(\frac{X_{t+1}}{A_{t+1}}\right)^{-\sigma}\frac{A_t}{A_{t+1}}\ \xi_{t+1}^{-\sigma} R_t \pi_{t+1}^{-1}\right] && (22)\\
N_t &= \frac{Y_t}{A_t}\int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\eta} di && (23)\\
X_t &= C_t - \theta C_{t-1} && (24)\\
Y_t &= C_t && (25)\\
\tau_t W_t N_t &= -T_t && (26)\\
\frac{P_t^f}{P_t} &= \frac{\eta}{\eta-1}\frac{\mathbb{E}_t \sum_{s=0}^{\infty}(\alpha\beta)^s\left(\frac{X_{t+s}\xi_{t+s}}{A_{t+s}}\right)^{-\sigma} mc_{t+s}\left(\frac{P_{t+s}\pi^{-s}}{P_t}\right)^{\eta}\frac{Y_{t+s}}{A_{t+s}}}{\mathbb{E}_t \sum_{s=0}^{\infty}(\alpha\beta)^s\left(\frac{X_{t+s}\xi_{t+s}}{A_{t+s}}\right)^{-\sigma}\left(\frac{P_{t+s}\pi^{-s}}{P_t}\right)^{\eta-1}\frac{Y_{t+s}}{A_{t+s}}} && (27)\\
mc_t &= \frac{W_t}{A_t P_t} && (28)\\
P_t^b &= P_{t-1}^{*}\pi_{t-1} && (29)\\
\ln P_{t-1}^{*} &= (1-\zeta)\ln P_{t-1}^{f} + \zeta P_{t-1}^{b} && (30)\\
P_t^{1-\eta} &= \alpha\left(\pi P_{t-1}\right)^{1-\eta} + (1-\alpha)\left(P_t^{*}\right)^{1-\eta} && (31)\\
\ln A_t &= \ln\gamma + \ln A_{t-1} + \ln z_t && (32)\\
\ln z_t &= \rho_z \ln z_{t-1} + \varepsilon_{z,t} && (33)\\
\ln(1-\tau_t) &= \rho^{\mu}\ln(1-\tau_{t-1}) + (1-\rho^{\mu})\ln(1-\tau) - \varepsilon_t^{\mu} && (34)
\end{aligned}
$$

with an associated equation describing the evolution of price dispersion, $\Delta_t = \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\eta} di$, which is not needed to tie down the equilibrium upon log-linearization. The model is then closed with the addition of a description of monetary policy, which will either be rule based, or derived from various forms of optimal policy discussed in the main text.

In order to render this model stationary we need to scale certain variables by the non-stationary level of technology, A_t such that $k_t = K_t/A_t$ where $K_t = \{Y_t, C_t, W_t/P_t\}$. All other real variables are naturally stationary. Applying this scaling, the steady-state equilibrium conditions reduce to:

$$
\begin{aligned}
N^{\varphi}X^{\sigma} &= w(1-\tau) \\
1 &= \beta R\pi^{-1}/\gamma = \beta r/\gamma \\
y &= N = c \\
X &= c(1-\theta) \\
\frac{\eta}{\eta-1} &= \frac{1}{w}.
\end{aligned}
$$

This system yields

$$
N^{\sigma+\varphi}(1-\theta)^{\sigma} = w(1-\tau). \tag{35}
$$

which can be solved for N. Note that this expression depends on the real wage w, which can be obtained from the steady-state pricing decision of our monopolistically competitive firms. In Appendix B we contrast this with the labor allocation that would be chosen by a social planner in order to fix the steady-state tax rate required to offset the net distortion implied by monopolistic competition and the consumption habits externality.

B The Social Planner's Problem

The subsidy level that ensures an efficient long-run equilibrium is obtained by comparing the steady state solution of the social planner's problem with the steady state obtained in the decentralized equilibrium. The social planner ignores the nominal inertia and all other inefficiencies and chooses real allocations that maximize the representative consumer's utility subject to the aggregate resource constraint, the aggregate production function, and the law of motion for habit-adjusted consumption:

$$
\max_{\{X_t^*, C_t^*, N_t^*\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(X_t^*, N_t^*, \xi_t, A_t\right)
$$

$$
\begin{aligned}
s.t.\ Y_t^* &= C_t^* \\
Y_t^* &= A_t N_t^* \\
X_t^* &= C_t^*/A_t - \theta C_{t-1}^*/A_{t-1}
\end{aligned}
$$

The optimal choice implies the following relationship between the marginal rate of substitution between labor and habit-adjusted consumption and the intertemporal marginal rate of substitution in habit-adjusted consumption

$$
\chi\left(N_t^*\right)^{\varphi}\left(X_t^*\right)^{\sigma} = (1-\theta\beta)\,\mathbb{E}_t\left(\frac{X_{t+1}^*\xi_{t+1}}{X_t^*\xi_t}\right)^{-\sigma}.
$$

The steady state equivalent of this expression can be written as

$$\chi\left(N^{*}\right)^{\varphi+\sigma}(1-\theta)^{\sigma}=(1-\theta\beta).$$

If we contrast this with the allocation achieved in the steady-state of our decentralized equilibrium (35) we can see that the two will be identical whenever the tax rate is set optimally to be

$$\tau^{*}\equiv 1-\frac{\eta}{\eta-1}(1-\theta\beta).$$

Notice that in the absence of habits the optimal tax rate would be negative, such that it is effectively a subsidy which offsets the monopolistic competition distortion. However, for the estimated values of the habits parameter the optimal tax rate is positive as the policy maker wishes to prevent households from overconsuming.

C Derivation of Objective Function

Individual utility in period t is

$$\Gamma_0=\mathbb{E}_0\sum_{t=0}^{\infty}\beta^{t}\left(\frac{X_t^{1-\sigma}\xi_t^{-\sigma}}{1-\sigma}-\frac{N_t^{1+\varphi}\xi_t^{-\sigma}}{1+\varphi}\right)$$

where $X_t=c_t-\theta c_{t-1}$ is habit-adjusted aggregate consumption after adjusting consumption for the level of productivity, $c_t=C_t/A_t$.

Linearization up to second order yields

$$\begin{aligned}\Gamma_0 \;\; = \;\; & \mathbb{E}_0\sum_{t=0}^{\infty}\beta^{t}\left(\overline{X}^{1-\sigma}\left\{\frac{1-\theta\beta}{1-\theta}\left(\widehat{c}_t+\frac{1}{2}\widehat{c}_t^{2}\right)-\frac{1}{2}\sigma\widehat{X}_t^{2}-\sigma\widehat{X}_t\widehat{\xi}_t\right\}\right.\\ & \left.-\overline{N}^{1+\varphi}\left\{\widehat{N}_t+\frac{1}{2}(1+\varphi)\widehat{N}_t^{2}-\sigma\widehat{N}_t\widehat{\xi}_t\right\}\right)+tip(3).\end{aligned}$$

where where $tip(3)$ includes terms independent of policy of third order and higher and for every variable Z_t with steady state value Z we denote $\widehat{Z}_t=\log(Z_t/Z)$.

The second order approximation to the production function yields the exact relationship $\hat{N}_t=\hat{\Delta}_t+\hat{y}_t$, where $y_t=Y_t/A_t$ and $\Delta_t=\int_0^1\left(\frac{P_t(i)}{P_t}\right)^{-\eta}di$. We substitute $\hat{N}_t$ out and follow Eser et al. (2009) in using

$$\sum_{t=0}^{\infty}\beta^{t}\hat{\Delta}_t=\frac{\alpha}{1-\alpha\beta}\hat{\Delta}_{-1}+\frac{1}{2}\sum_{t=0}^{\infty}\beta^{t}\frac{\alpha\eta}{(1-\beta\alpha)(1-\alpha)}\left(\widehat{\pi}_t^{2}+\frac{\zeta\alpha^{-1}}{(1-\zeta)}\left[\widehat{\pi}_t-\widehat{\pi}_{t-1}\right]^{2}\right)$$

to yield

$$\Gamma_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\overline{X}^{1-\sigma} \left\{ \frac{1-\theta\beta}{1-\theta} \left(\widehat{c}_t + \frac{1}{2}\widehat{c}_t^2 \right) - \frac{1}{2}\sigma \widehat{X}_t^2 - \sigma \widehat{X}_t \widehat{\xi}_t \right\} \right.$$
$$\left. -\overline{N}^{1+\varphi} \left(\begin{array}{c} \widehat{y}_t + \frac{1}{2}\frac{\alpha\eta}{(1-\beta\alpha)(1-\alpha)} \left(\widehat{\pi}_t^2 + \frac{\zeta\alpha^{-1}}{(1-\zeta)} \left[\widehat{\pi}_t - \widehat{\pi}_{t-1} \right]^2 \right) \\ +\frac{1}{2}(1+\varphi)\widehat{y}_t^2 - \sigma \widehat{y}_t \widehat{\xi}_t \end{array} \right) \right) + tip(3).$$

The second order approximation to the national income identity yields

$$\widehat{c}_t + \frac{1}{2}\widehat{c}_t^2 = \widehat{y}_t + \frac{1}{2}\widehat{y}_t^2 + tip\,(3)\,.$$

Finally, we use that in the efficient steady-state $\overline{X}^{1-\sigma}(1-\theta\beta) = (1-\theta)\overline{N}^{1+\varphi}$ and collect terms to arrive at

$$\Gamma_0 = -\frac{1}{2}\overline{N}^{1+\varphi}\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{\sigma(1-\theta)}{1-\theta\beta} \left(\widehat{X}_t + \widehat{\xi}_t \right)^2 + \varphi \left(\widehat{y}_t - \frac{\sigma}{\varphi}\widehat{\xi}_t \right)^2 \right.$$
$$\left. + \frac{\alpha\eta}{(1-\beta\alpha)(1-\alpha)} \left(\widehat{\pi}_t^2 + \frac{\zeta\alpha^{-1}}{(1-\zeta)} \left[\widehat{\pi}_t - \widehat{\pi}_{t-1} \right]^2 \right) \right\} + tip\,(3)\,.$$

Figure 1: Prior and Posterior Distributions of Parameters

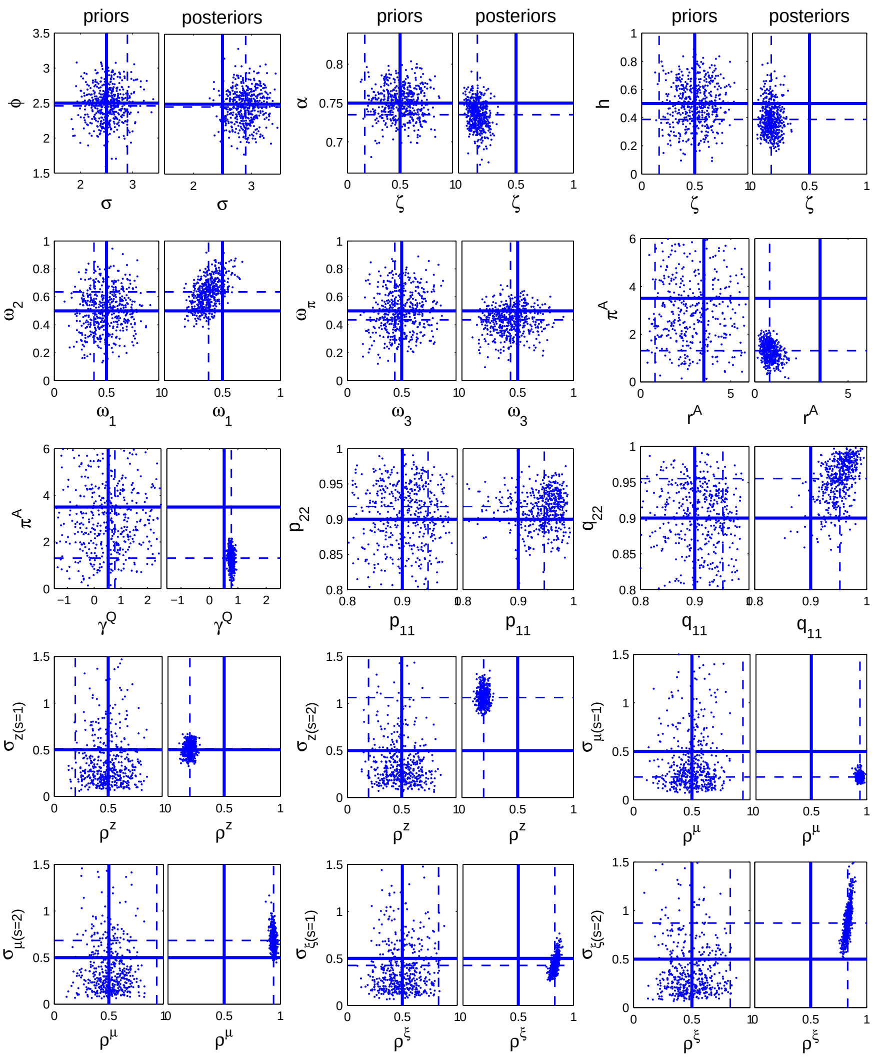

Notes: The panels depict 500 draws from prior and posterior distributions from the estimates in the first column of Table 4. The draws are plotted for pairs of estimated parameters and the intersections of lines signify prior (solid) and posterior (dashed) means, respectively.

Figure 2: Markov Switching Probabilities - Policy Switches Only

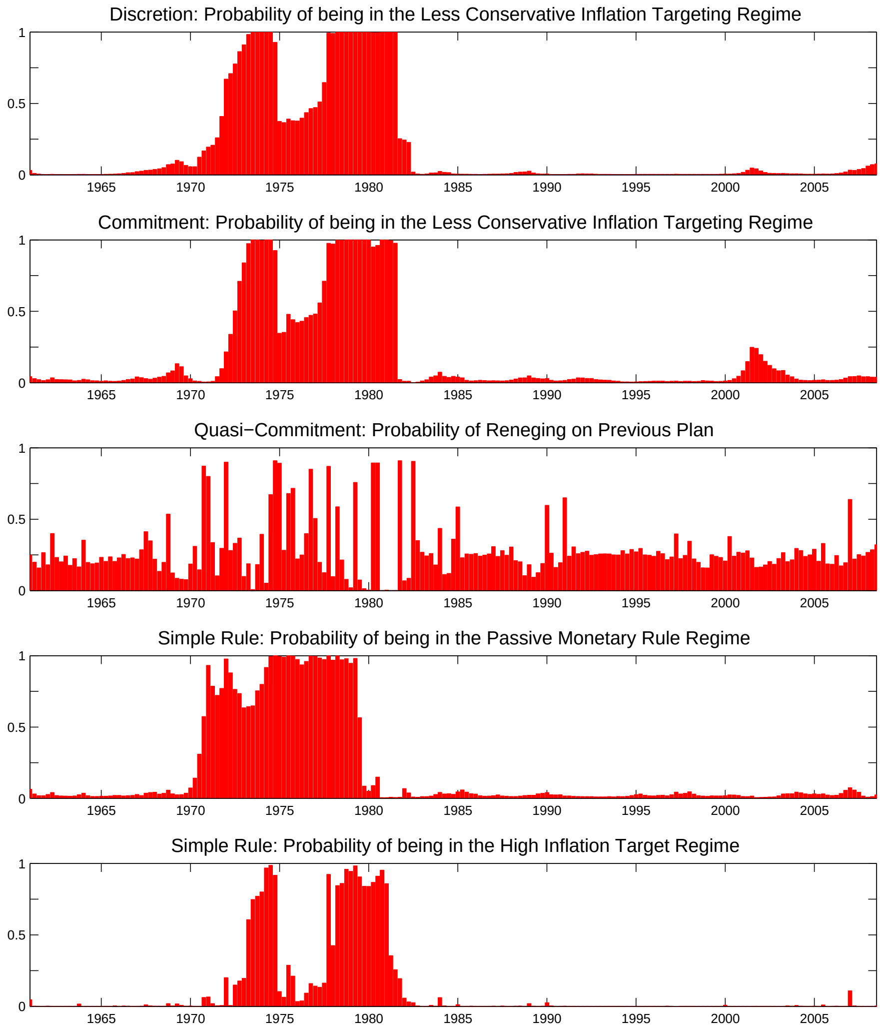

Figure 3: Markov Switching Probabilities - Policy and Volatility Switches

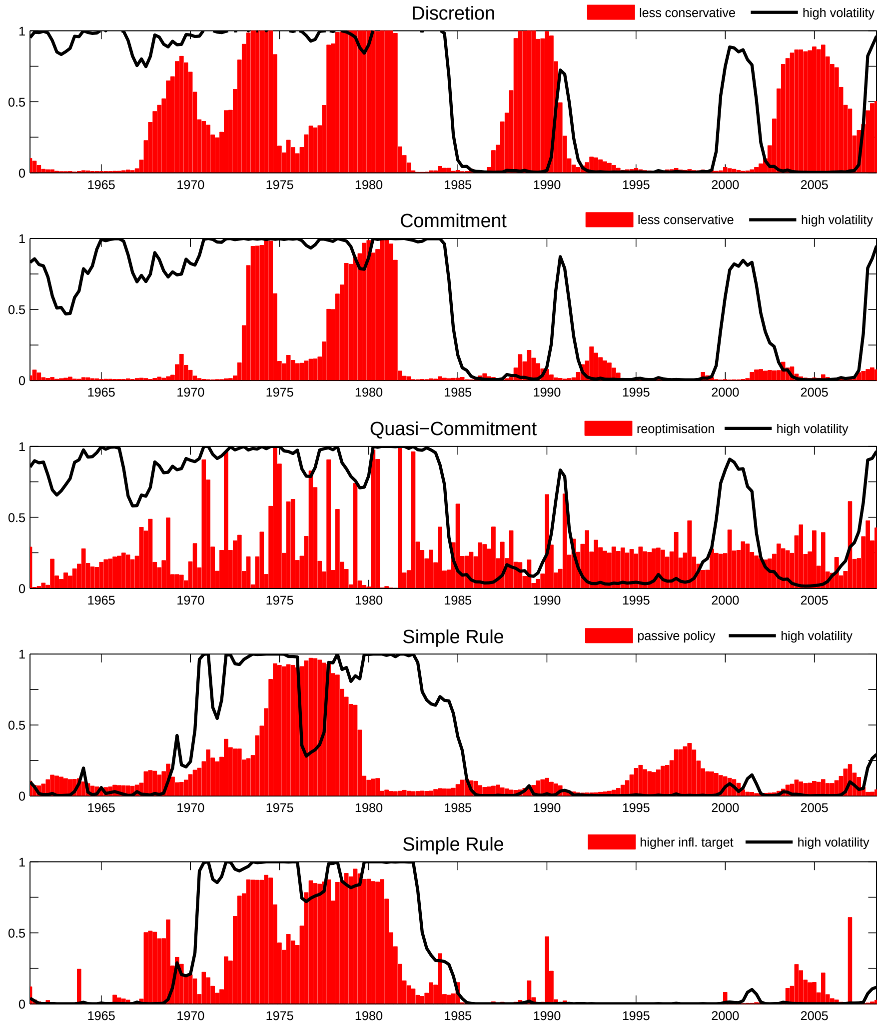

Figure 4: Counterfactuals under Different Volatility Regimes

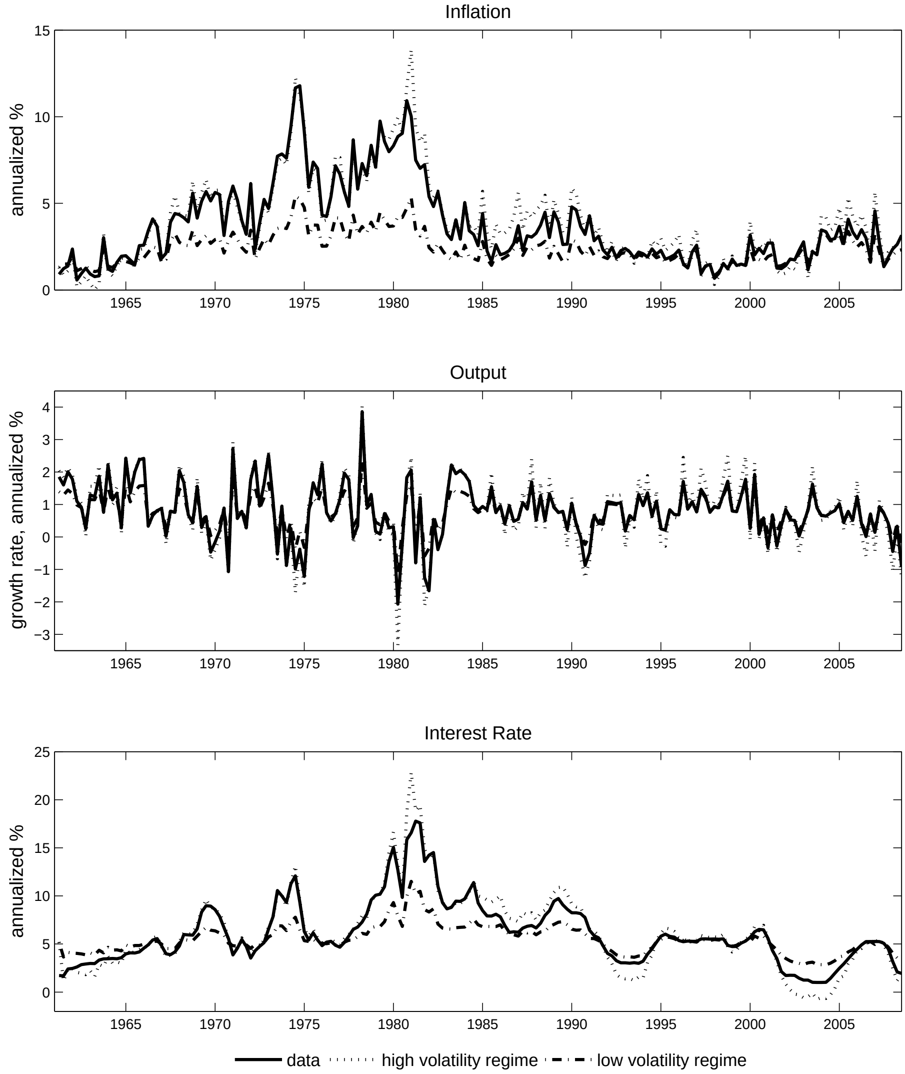

Figure 5: Counterfactuals under Different Levels of Conservatism

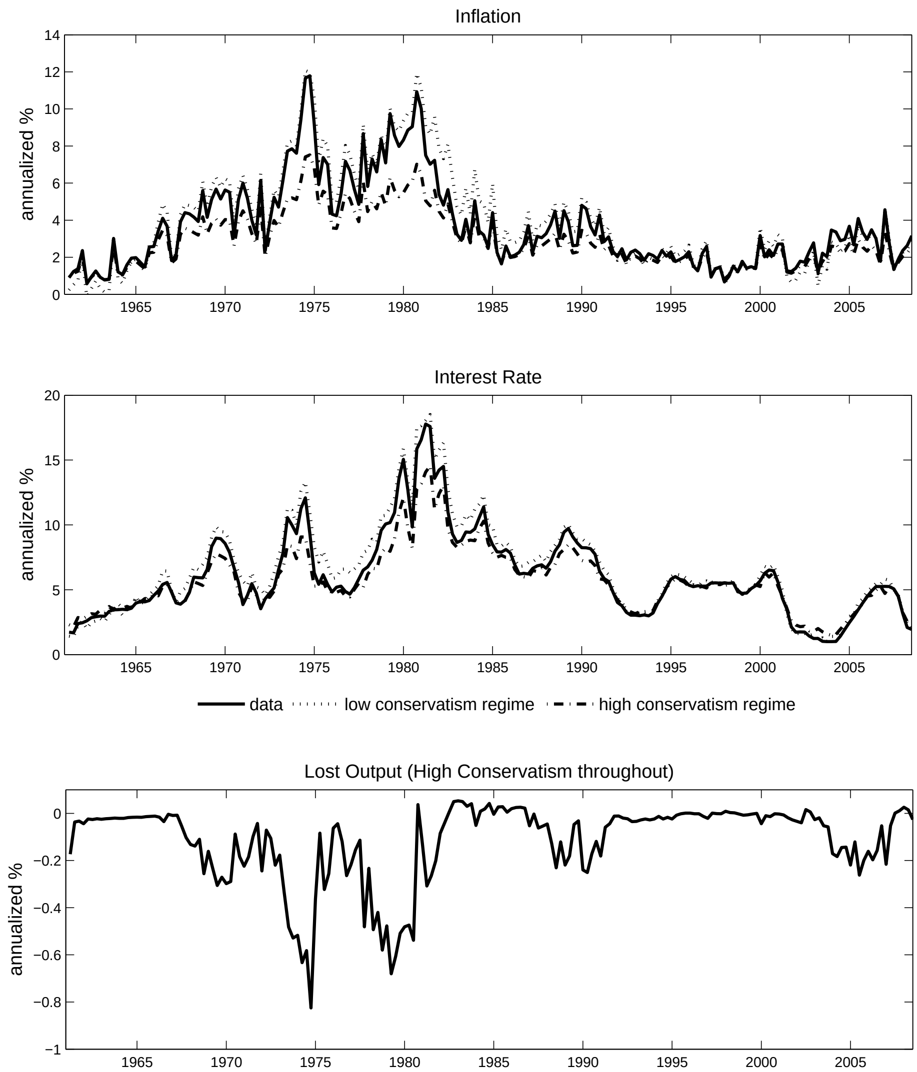

Notes: Lower panel plots the difference between output observed given the model account of regime switches and output attained if only the conservatism regime is realized.

Figure 6: Counterfactuals: Commitment versus Discretion

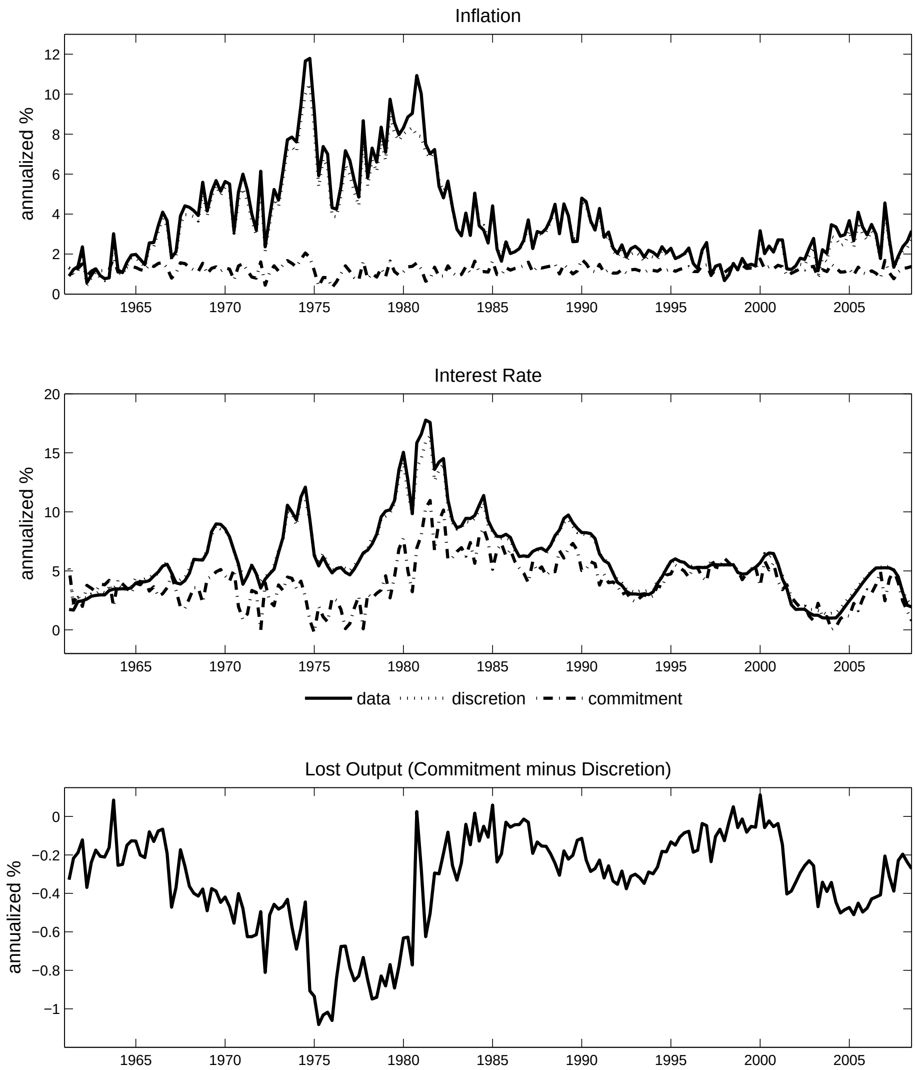

Notes: Lower panel plots the difference between output observed given the model account of regime switches, assuming discretionary policymaking, and output attained if the policy maker is able to commit *cet. par.*

Table 1: Distribution of Priors

Parameters		Range	Density	Mean	Std Dev
Inv. of intertemp. elas. of subst.	σ	$\mathbb{R}$	Nomal	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	$\mathbb{R}$	Nomal	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	$[0,1)$	Beta	0.3	0.02
Markov Switching s.d. of shocks					
taste shocks	$\sigma^2_{\xi(s=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
cost-push shocks	$\sigma^2_{\mu(s=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
productivity shocks	$\sigma^2_{z(s=1=2)}$	$\mathbb{R}^+$	Inv. Gamma	0.50	5
policy shocks	$\sigma^2_{R(s=1=2)}$	$\mathbb{R}^+$	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 (strong inflation targeting)	$\psi_{1(S=1)}$	$\mathbb{R}^+$	Gamma	1.50	0.50
inflation (weak inflation targeting)	$\psi_{1(S=2)}$	$\mathbb{R}^+$	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	3	2
inflation target $(S=2)$	$\pi^A_{(S=2)}$	$\mathbb{R}^+$	Gamma	6	2
Transition Probabilities					
policy: remaining with strong infl. targeting	p_{11}	$[0,1)$	Beta	0.90	0.05
policy: remaining with weak infl. targeting	p_{22}	$[0,1)$	Beta	0.90	0.05
volatility: remaining with low volatility	q_{11}	$[0,1)$	Beta	0.90	0.05
volatility: remaining with high volatility	q_{22}	$[0,1)$	Beta	0.90	0.05

Table 2: Estimation Results - No Switching

Parameters	Simple Rule	Discretion	Quasi-Commitment	Commitment
		Model Parameters		
σ	2.802 [2.407,3.188]	2.722 [2.338,3.101]	2.412 [2.066,2.750]	2.832 [2.477,3.062]
α	0.779 [0.751,0.807]	0.760 [0.734,0.786]	0.793 [0.772,0.814]	0.768 [0.734,0.789]
ζ	0.103 [0.039,0.166]	0.156 [0.066,0.241]	0.169 [0.083,0.252]	0.594 [0.489,0.737]
θ	0.823 [0.685,0.964]	0.476 [0.267,0.680]	0.531 [0.306,0.779]	0.643 [0.444,0.782]
φ	2.417 [2.005,2.833]	2.387 [2.146,2.627]	1.832 [1.439,2.244]	2.312 [2.046,2.749]
		Shock Processes		
ρ^{ξ}	0.899 [0.859 0.941]	0.845 [0.806 0.885]	0.902 [0.873 0.933]	0.862 [0.793 0.903]
ρ^{μ}	0.500 [0.246 0.747]	0.947 [0.922 0.974]	0.947 [0.920 0.975]	0.968 [0.948 0.989]
ρ^{z}	0.320 [0.219 0.418]	0.223 [0.172 0.275]	0.213 [0.164 0.262]	0.199 [0.132 0.241]
σ^2_{ξ}	0.987 [0.715 1.247]	0.763 [0.486 1.039]	0.938 [0.568 1.300]	0.983 [0.706 1.165]
σ^2_{μ}	0.567 [0.341 0.788]	0.489 [0.356 0.613]	0.722 [0.506 0.919]	3.628 [2.816 4.916]
σ^2_{z}	0.797 [0.731 0.863]	0.827 [0.756 0.896]	0.844 [0.769 0.920]	0.770 [0.701 0.815]
σ^2_{R}	0.251 [0.227 0.273]	–	–	–

continued on the next page

Table 2: Estimation Results - No Switching – continued

Parameters	Simple Rule	Discretion	Quasi-Commitment	Commitment
Data Means				
r^A	0.706 [0.246,1.139]	0.966 [0.352,1.569]	1.073 [0.379,1.712]	1.088 [0.459,1.540]
π^A	4.746 [3.800,5.677]	2.656 [1.008,4.221]	2.353 [1.883,2.842]	4.050 [3.642,4.674]
γ^Q	0.688 [0.547,0.826]	0.716 [0.593,0.835]	0.732 [0.611,0.863]	0.726 [0.594,0.797]
Policy Parameters				
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	–	0.458 [0.287,0.627]	0.710 [0.563,0.858]	0.627 [0.490,0.808]
ω_2	–	0.758 [0.628,0.901]	0.807 [0.694,0.927]	0.446 [0.316,0.620]
ω_3	–	0.451 [0.213,0.692]	0.379 [0.150,0.609]	0.489 [0.268,0.712]
υ	–	–	0.308 [0.276,0.340]	–
Log Marginal Data Densities and Bayes Factors				
Geweke (1999)	−841.01 (1.00)	−842.494 (4.41)	−843.88 (17.6)	−855.43 (1.84e+6)
Sims et al. (2008)	−841.09 (1.00)	−842.691 (4.96)	−843.91 (16.8)	−858.26 (2.85e+7)

Notes: For each parameter the posterior distribution is described by mean and 90% confidence interval in square brackets. Bayes Factors for marginal data densities are in parentheses. Computation of the q_L statistic of Sims et al. (2008), which assesses the overlap between the weighting matrix and the posterior density, indicates that the calculated marginal log likelihoods are reliable in every case.

Table 3: Estimation Results - Switches in Policy Only

Parameters	Discretion	Rule - Parameters	Rule - Target	Commitment
		Model Parameters		
σ	2.896 [2.500,3.288]	2.621 [2.382,2.861]	2.791 [2.403,3.187]	2.921 [2.560,3.277]
α	0.731 [0.706,0.758]	0.775 [0.747,0.803]	0.779 [0.750,0.807]	0.770 [0.744,0.796]
ζ	0.155 [0.069,0.239]	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]	0.825 [0.698,0.954]	0.810 [0.658,0.961]	0.606 [0.388,0.843]
φ	2.331 [1.916,2.757]	2.425 [2.025,2.848]	2.410 [2.003,2.846]	2.271 [1.872,2.679]
		Shock Processes		
ρ^{ξ}	0.805 [0.766,0.844]	0.887 [0.850,0.927]	0.898 [0.858,0.941]	0.904 [0.877,0.933]
ρ^{μ}	0.957 [0.937,0.978]	0.501 [0.250,0.748]	0.499 [0.250,0.751]	0.986 [0.978,0.995]
ρ^{z}	0.213 [0.164,0.261]	0.307 [0.208,0.403]	0.317 [0.218,0.417]	0.210 [0.154,0.268]
σ^2_{ξ}	0.515 [0.289,0.719]	0.981 [0.755,1.199]	0.848 [0.609,1.090]	0.797 [0.511,1.069]
σ^2_{μ}	0.444 [0.327,0.554]	0.275 [0.169,0.382]	0.569 [0.340,0.795]	2.325 [1.697,2.947]
σ^2_{z}	0.829 [0.755,0.896]	0.797 [0.169,0.382]	0.795 [0.727,0.861]	0.779 [0.711,0.846]
σ^2_{R}	–	0.235 [0.213,0.256]	0.252 [0.229,0.275]	–
		Data Means		
r^{A}	0.766 [0.303,1.213]	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]	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]	0.677 [0.540,0.808]	0.681 [0.544,0.822]	0.741 [0.619,0.862]

continued on the next page

Table 3: Estimation Results - Switches in Policy Only – continued

Parameters	Discretion	Rule - Parameters	Rule - Target	Commitment
		Policy Parameters		
$\rho^R_{(S=1)}$	–	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]	–	–	0.502 [0.331,0.666]
ω_2	0.650 [0.460,0.847]	–	–	0.523 [0.295,0.732]
ω_3	0.442 [0.164,0.698]	–	–	0.460 [0.205,0.710]
$\omega_{\pi(S=1)}$	1	–	–	1
$\omega_{\pi(S=2)}$	0.347 [0.219,0.477]	–	–	0.302 [0.194,0.414]
p_{11}	0.978 [0.962,0.994]	0.962 [0.939,0.989]	0.956 [0.930,0.984]	0.979 [0.962,0.996]
p_{22}	0.940 [0.900,0.981]	0.802 [0.734,0.870]	0.796 [0.722,0.876]	0.816 [0.735,0.901]
		Log Marginal Data Densities and Bayes Factors		
Geweke (1999)	-810.98 (1.00)	-825.33 ($1.72e{+}6$)	-831.74 ($1.04e{+}9$)	-832.85 ($3.14e{+}9$)
Sims et al. (2008)	-811.24 (1.00)	-825.44 ($1.46e{+}6$)	-831.81 ($8.52e{+}8$)	-832.98 ($2.75e{+}9$)

Notes: For each parameter the posterior distribution is described by mean and 90% confidence interval in square brackets. Bayes Factors for marginal data densities are in parentheses. Computation of the q_L statistic of Sims et al. (2008), which assesses the overlap between the weighting matrix and the posterior density, indicates that the calculated marginal log likelihoods are reliable in every case.

Table 4: Estimation Results - Switches in Policy and Volatility

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]	2.323 [2.035,2.639]	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]	0.785 [0.763,0.807]	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]	0.157 [0.083,0.231]	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]	0.466 [0.375,0.554]	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]	1.776 [1.425,2.099]	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]	0.895 [0.871,0.919]	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]	0.930 [0.904,0.958]	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]	0.193 [0.142,0.245]	0.162 [0.106,0.218]
$\sigma^2_{\xi(s=1)}$	0.425 [0.297,0.546]	0.682 [0.527,0.837]	0.545 [0.390,0.690]	0.526 [0.340,0.710]	0.404 [0.249,0.555]
$\sigma^2_{\xi(s=2)}$	0.873 [0.599,1.139]	1.467 [1.040,1.888]	1.346 [0.958,1.721]	1.104 [0.757,1.444]	1.224 [0.720,1.757]
$\sigma^2_{\mu(s=1)}$	0.236 [0.182,0.292]	0.277 [0.169,0.381]	0.276 [0.169,0.383]	0.296 [0.215,0.372]	1.329 [0.737,1.905]
$\sigma^2_{\mu(s=2)}$	0.684 [0.527,0.840]	0.546 [0.343,0.751]	0.545 [0.390,0.690]	1.000 [0.747,1.237]	2.806 [1.697,3.913]
$\sigma^2_{z(s=1)}$	0.512 [0.391,0.622]	0.601 [0.540,0.660]	0.603 [0.542,0.664]	0.453 [0.359,0.547]	0.452 [0.372,0.526]
$\sigma^2_{z(s=2)}$	1.064 [0.932,1.193]	1.184 [0.981,1.380]	1.156 [0.977,1.329]	1.064 [0.939,1.187]	0.989 [0.870,1.103]
$\sigma^2_{R(s=1)}$	–	0.140 [0.124,0.156]	0.146 [0.129,0.162]	–	–
$\sigma^2_{R(s=2)}$	–	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]	0.637 [0.238,1.029]	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]	2.108 [1.628,2.575]	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]	0.771 [0.675,0.871]	0.828 [0.721,0.931]

continued on the next page

Table 4: Estimation Results - Switches in Policy and Volatility – continued

Parameters	Discretion	Rule - Parameters	Rule - Target	Quasi-Commitment	Commitment
			Policy Parameters		
υ	–	–	–	0.292 [0.262,0.323]	–
$\rho^R_{(S=1)}$	–	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]	–	–	0.746 [0.608,0.879]	0.503 [0.320,0.690]
ω_2	0.635 [0.468,0.800]	–	–	0.830 [0.730,0.934]	0.559 [0.280,0.843]
ω_3	0.436 [0.200,0.667]	–	–	0.404 [0.167,0.636]	0.454 [0.195,0.695]
$\omega_{\pi(S=1)}$	1	–	–	1	1
$\omega_{\pi(S=2)}$	0.436 [0.279,0.589]	–	–	–	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]	–	0.978 [0.959,0.997]
p_{22}	0.918 [0.876,0.962]	0.846 [0.812,0.880]	0.812 [0.740,0.889]	–	0.798 [0.722,0.877]
q_{11}	0.952 [0.919,0.986]	0.956 [0.928,0.985]	0.979 [0.960,0.998]	0.904 [0.841,0.968]	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 [0.907,0.976]	0.933 [0.887,0.976]
			Log Marginal Data Densities and Bayes Factors		
Geweke (1999)	−759.78 (1.00)	−764.16 (80.29)	−765.83 (425.76)	−791.10 (4.00e+13)	−793.62 (4.98e+14)
Sims et al. (2008)	−759.91 (1.00)	−764.21 (74.08)	−765.95 (422.76)	−791.10 (3.51e+13)	−793.95 (6.12e+14)

Notes: For each parameter the posterior distribution is described by mean and 90% confidence interval in square brackets. Bayes Factors for marginal data densities are in parentheses. Computation of the q_L statistic of Sims et al. (2008), which assesses the overlap between the weighting matrix and the posterior density, indicates that the calculated marginal log likelihoods are reliable in every case.

Table 5: Unconditional Variances and Welfare under Alternative Policies and Volatilities

Regime: (conservatism, volatility)	Output	Inflation	Interest Rate	Welfare Cost (est. weights)	Welfare Cost (micro. weights)
			Discretion		
(low, 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%]
(high, 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%]
(low, 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%]
			Commitment		
(low, 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%]
(high, 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%]
(low, 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%]
(high, 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%]

Notes: The figures in the first three columns measure the unconditional variances of output, inflation and interest rates for estimated parameters in regime (conservatism, volatility). The welfare cost using estimated weights is computed using equation (18). The welfare costs using micro-founded weights is based on equation (12), but is 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.
* denotes the Benchmark Case for Table 6.

Table 6: Percentage Reduction (-) or Increase (+) in Variances and Welfare Relative to the Benchmark Case in Table 5.

Regime: (conservatism, volatility)	Output	Inflation	Interest Rate	Welfare Cost (est. weights)	Welfare Cost (micro. weights)
			Discretion		
(high, high)	+2.79%	-65.83%	-59.16%	-3.81%	-59.21%
(low, low)	-59.58%	-60.98%	-64.95%	-78.24%	-82.77%
(high, low)	-61.62%	-86.26%	-84.61%	-78.71%	-92.35%
			Commitment		
(low, high)	+12.91%	-97.39%	-48.58%	-19.98%	-87.13%
(high, high)	+14.20%	-99.10%	-51.97%	-19.23%	-89.86%
(low, low)	-57.61%	-98.86%	-74.89%	-81.53%	-97.15%
(high, low)	-58.49%	-99.61%	-76.53%	-81.37%	-97.70%