Changes in Federal Reserve Preferences *

Aeimit Lakdawala †

Department of Economics

Michigan State University

May 2016

Abstract

Using a model of optimizing central bank behavior, I estimate the dynamic behavior of preferences, which are captured by the relative weight put on stabilizing inflation versus minimizing the output gap. Unlike previous work, I let this parameter vary continuously over time. There is a drastic but steady rise in the weight on inflation around the appointment of Paul Volcker; however, I find variation in preferences throughout the sample period. The results suggest that preference changes have been more complex than typically assumed in the literature. The estimated preference series is used to perform counterfactual experiments and to construct a new measure of monetary shocks.

Keywords: monetary policy, Federal Reserve preferences, monetary shocks, time-varying parameter, extended Kalman filter

JEL classification: E430, E520, E580

*I would like to thank James Hamilton, Davide Debortoli, Giacomo Rondina, Valerie Ramey, Marjorie Flavin and seminar participants at UCSD, Federal Reserve Bank of St. Louis, Federal Reserve Bank of Boston, University of Georgia, Emory and Midwest Macro for helpful comments and suggestions.

†aeimit@msu.edu

1 Introduction

The Federal Reserve has a dual mandate of full employment and price stability, and quite often there is a short-run tradeoff between these two objectives. For example, although the Federal Reserve tried to combat the high unemployment rate following the financial crisis of 2007-08 by using both traditional and unconventional measures, it was cautious in implementing more aggressive policy actions. This was in part due to the fear of causing a future rise in inflation. Thus at each point in time, the Fed is forced to choose which it dislikes more: high unemployment or high inflation. The weight that it places on inflation relative to unemployment can be thought of as a measure of Federal Reserve preferences and is a crucial input into monetary policy decisions. The primary goal of this paper is to estimate how these preferences have evolved over time in order to gain a deeper understanding of the motivations behind Federal Reserve actions.

Most of the existing work in the literature has used some version of a Taylor rule to model monetary policy, where the interest rate responds to output and inflation. Consequently, attempts to investigate changes in the behavior of the Federal Reserve have typically involved looking for instability in the Taylor rule.[1] While there is evidence of changes in these parameters, these changes cannot be used as direct evidence that preferences have changed. This is because Taylor rules are only reduced form representations of monetary policy behavior. The estimated coefficients of the Taylor rule capture the combined effect of underlying structural parameters of the economy and preferences of the central bank. In this paper I use the restrictions from a simple model of optimizing central bank behavior to identify and estimate a time-varying series of Federal Reserve preferences. The central bank's goal is simultaneously to stabilize inflation around an inflation target and minimize the output gap. I use the relative importance that the central bank puts on inflation stabilization versus output gap minimization as a measure of central bank preferences. Using the central bank's optimization problem, I derive an optimal policy rule where the coefficients are functions of both the preference parameter and the underlying structure of the economy.

[1] Typically, the parameters are allowed to be different across split samples (Clarida et al. (2000), Lubik and Schorfheide (2004) and Boivin and Giannoni (2006)), allowed to switch between regimes (Fernandez-Villaverde et al. (2010)) or vary continuously over time (Boivin (2006), Kim and Nelson (2006), Cogley and Sargent (2005)).

Not only is it important to structurally model the preference parameter, but there is reason to believe that this parameter is not constant over time. There are two main broad reasons why the preference parameter can change over time. First, the composition of the Federal Open Market Committee (FOMC)[2] changes over time. This includes changes in the chairman, with the most famous example being the appointment of Paul Volcker in 1979; Meltzer (2006) states that the biggest difference with the appointment of Paul Volcker was the changing of the weight the Fed put on inflation relative to unemployment. While the Fed chairman has changed only a handful of times in the last few decades, the composition of the FOMC changes more often with rotating voting rights for Presidents of four of the regional Federal Reserve Banks and staggered changes in members of the Board of Governors. Additionally even if the composition of the FOMC committee is the same over two periods, there can be changes in the preferences of the committee members.[3] Second, Fed preferences can change due to political pressure on the Fed. There are accounts of Presidents Johnson and Nixon putting pressure on Fed chairmen Martin and Burns to refrain from monetary tightening (Meltzer (2011)). More recently with the financial crisis, Di Maggio (2010) provides evidence of monetary policy being influenced by Congress. In some instances the Fed has been more influenced by this pressure (see Burns (1987)) while in others the Fed has stood its ground (see Greenspan (2008)).

But most papers in the literature that estimate Fed preferences have assumed constant preferences. When efforts have been made to introduce time variation into Fed Preferences, they have been done in restrictive ways. The common approach is estimating two values for the the preference parameter by splitting the sample at the appointment of Paul Volcker (see Dennis (2006), Favero and Rovelli (2003), Ozlale (2003), Salemi (2006), Ilbas (2012), Givens (2012) and Best (ming)). Given the potential for gradual changes in the preference parameter due to continual changes in FOMC composition or political pressure, it is important to allow for a more flexible form of estimation that can capture this. In this paper I use a time-varying parameter approach to

[2]The FOMC is the Fed's main monetary policy making arm.

[3]For example the Minneapolis Fed President Naryana Kocherlakota famously switched from being an "inflation hawk" to a "dove" in 2012, see for example http://blogs.wsj.com/economics/2012/09/28/another-reason-kocherlakota-changed-his-mind/

estimate a preference parameter that can change continuously over time. The combination of an optimal policy setup with the time variation in the preferences introduces non-linearities that can not be addressed using the linear estimation techniques that are standard in this literature. I deal with these non-linearities by using the Extended Kalman Filter which is conveniently embedded in a Bayesian Markov Chain Monte Carlo algorithm. The only other paper in my knowledge to allow a more general form of time variation is Owyang and Ramey (2004), where Fed preferences are modeled as a Markov-switching process. They find multiple switches between hawk and dove regimes providing further evidence against the split sample approach and supporting the framework adopted in this paper.[4] While the regime-switching framework has been a popular technique to model parameter instability in the literature, it is not the natural choice here as we want to allow for potential gradual changes.[5]

I find that the Fed generally put lower weight on inflation relative to the output gap in the 1970s, though even within Arthur Burns' tenure there is some variation in preferences. The weight starts rising gradually around the appointment of Paul Volcker but the rise is temporary as the weight falls in the early 1980s before rising again in the mid-1980s. Under Greenspan there is a steady decline in this weight until the end of his tenure when the weight begins to increase again. There is a growing literature that focuses on estimating dynamic macroeconomic models with a framework of optimal monetary policy. These results suggest that researchers should be careful to account for the dynamics of preferences. To further illustrate the importance of the dynamics of the estimated preferences I perform the following exercise. I first take the average of the preferences under Volcker's and Greenspan's regimes and using impulse responses show that they imply pretty similar reactions of the Federal Reserve to inflation shocks. However these average differences in preferences mask the considerable variation within each chairman's tenure. If instead I use the preference parameter from specific points during each chairman's time at the Fed, I find that their reactions to inflation shocks differ significantly.

[4] Although they take a more reduced form approach where the preferences capture, among other things, changes in the weight parameter used here and changes in the inflation target. Additionally they use a stylized model that implies restricted dynamics for the economy.

[5] Furthermore, a regime-switching estimation would be highly intractable for the model in this paper as discussed in section 3.

I then use changes in the estimated preference parameter to construct a novel measure of monetary policy shocks. In their seminal survey on monetary policy shocks Christiano et al. (1999) point to exactly these changes in preferences as an interpretation of monetary policy shocks estimated in Vector Autoregressions (VARs). I use two alternative identification strategies to embed my measure of monetary policy shocks in a structural VAR and evaluate its effects on the economy. The results suggest that the preference shocks have a smaller contemporaneous but more persistent effect on output, relative to using a more traditional shock measure in a VAR.

Since the results show that the Fed became more hawkish with the election of Paul Volcker, it is important to understand the role of preferences in the high inflation episode of the 1970s. To shed light on this issue I consider a counterfactual exercise where I evaluate the effect of electing Paul Volcker earlier in the 1970s. The results suggest that a stronger preference for inflation relative to output gap would have led to lower inflation but this alone would not have been enough to avoid the Great Inflation. This suggests that the shocks hitting the economy played a big role in the inflation episode of the 1970s. I also consider the role of preferences in explaining the fall in volatilities of macroeconomic variables since the early 1980s, the so called Great Moderation. I construct counterfactual histories after fixing the preference parameter to its average value in the pre-1984 sample. I find a limited role for preferences to affect the volatilities of aggregate variables. Additionally, since the model is estimated allowing the standard deviation of the shocks to be different in the pre- and post-1984 samples, I can examine the role of the size of the shocks in explaining the Great Moderation. I construct a similar counterfactual experiment and find that unlike preference changes, the fall in the volatility of the shocks hitting the economy have played a major role in contributing to the Great Moderation.

An alternative way to model time variation in the preferences of the Federal Reserve would be to allow the inflation target to vary over time. In this model a change in the inflation target only affects the constant term in the optimal interest rate rule without changing the responses to inflation or output gap, whereas a change in the weight parameter affects both the constant term and the responses. Given the empirical evidence of changes in the coefficients of the monetary policy rule, it seems more natural to consider a model where Fed preferences are allowed to affect the responses to inflation and output gap.[6]

The rest of the paper is organized as follows. In the next section, I setup the model and explain the optimization problem. In Section 3, I provide an outline of the estimation strategy and discuss the Bayesian estimation algorithm. The details are included in the appendix. I discuss the main results in section 4 and conduct counterfactual analyses in section 5. In section 6, I use the estimated preference parameter series to back out a new measure of monetary policy shocks and assess its impact on the economy. Section 7 provides a variety of robustness checks including using real-time data to construct an alternative measure of the output gap. In section 8 I offer some concluding remarks.

2 The Model

I use a simple model where the central bank minimizes a quadratic loss function subject to linear constraints that characterize the behavior of the economy.

2.1 Constraints

I use the backward-looking model outlined in Rudebusch and Svensson (1998). The main advantages of this setup are that it is parsimonious and fits the data well. Previous work has also used this model to study the preferences of the Fed (Dennis (2006), Favero and Rovelli (2003) and Ozlale (2003)). The backward looking equations can be thought of as representing adaptive expectations.[7] The model incorporates two basic equations describing the behavior of inflation and the output

[6]Furthermore, in an online appendix I discuss how the inflation target in this model relates to other uses in the literature and that the overall framework of this model is consistent with some models that use a time-varying inflation target.

[7]This backward looking model implies that agents form expectations based on past realizations of data. This is not a trivial assumption and incorporating forward looking behavior may have important implications. However there is disagreement about whether to model agents' expectations in a completely rational framework or as following some limited information form of learning. In the current paper any of these extensions would significantly increase the computational burden in the estimation and are thus left for future work.

gap. The first equation is the aggregate supply curve (or the Phillips curve) which relates inflation to lagged inflation and the output gap.

$$\pi_t = b_0 + b_1\pi_{t-1} + b_2\pi_{t-2} + b_3\pi_{t-3} + (1 - b_1 - b_2 - b_3)\pi_{t-4} + b_4\tilde{y}_{t-1} + s_t \quad (1)$$

where π_t is annualized quarterly inflation and $\tilde{y}_t$ is the output gap. The coefficients on the lags of inflation are restricted to sum to one. This implies that there is no long run tradeoff between inflation and the output gap. This assumption is consistent with the natural rate hypothesis and means that the Fed cannot manipulate inflation to have a permanent effect on the output gap.[8] The error term (s_t) can be interpreted as a supply shock. The second equation is the IS curve which relates the output gap to lagged output gaps and the real interest rate as follows

$$\tilde{y}_t = a_0 + a_1\tilde{y}_{t-1} + a_2\tilde{y}_{t-2} + a_3\left[i^a_{t-1} - \pi^a_{t-1}\right] + g_t \quad (2)$$

where i_t is the annualized quarterly fed funds rate and $i^a_t = \frac{1}{4}\sum_{j=0}^{3} i_{t-j}$ and $\pi^a_t = \frac{1}{4}\sum_{j=0}^{3}\pi_{t-j}$ are the one year trailing averages of the fed funds rate and inflation respectively. Here the error term (g_t) represents a demand shock. In the estimation the variables are constructed in the following manner. $\pi_t \equiv 400 * [\ln(P_t) - \ln(P_{t-1})]$ is quarterly inflation of GDP chain-weighted price index at an annualized rate and $\tilde{y}_t \equiv 100 * [\ln(y_t) - \ln(y^*_t)]$ is the output gap. y_t is quarterly real GDP and y^*_t is the Congressional Budget Office's measure of potential output. In section 7.1, I present and discuss results using alternative measures of the output gap, including a real-time version.

2.2 Loss Function

The loss function of the central bank is assumed to be quadratic.

$$L = \tilde{E}_t \sum_{j=0}^{\infty} \beta^j \left[\alpha_t \left(\pi^a_{t+j} - \pi^*\right)^2 + \tilde{y}^2_{t+j} + \nu(i_{t+j} - i_{t+j-1})^2\right] \quad (3)$$

[8] Note that the estimated coefficients will imply that there is a short-run tradeoff.

The first two terms of the objective function are standard and represent the Federal Reserve's dual mandate of price stability and full employment. The central bank wants to keep annual inflation close to an inflation target π^* and the output gap close to zero.[9] This quadratic loss function is widely used in the monetary policy rules literature (Taylor (1999), Rudebusch and Svensson (1998)) and makes it easy to compare my estimates to the literature. Furthermore, a quadratic loss function with linear constraints keeps the model tractable which is especially important given the involved nature of the estimation procedure.

For a theoretical motivation of the third term $\nu(i_{t+j} - i_{t+j-1})^2$, based on the desirable effects on private-sector expectations, see Woodford (2003). This term can also be motivated by the central bank's desire to reduce volatility of asset prices by avoiding big changes in the interest rate. Rudebusch (2002), Rudebusch (2006), Castelnuovo and Surico (2004) have argued that the central bank does not value smoothing directly but rather the term arises as an artefact of other things like autocorrelated shocks hitting the economy or potential concern about model uncertainty. However, in recent empirical work Borger et al. (2011) and Coibion and Gorodnichenko (2012) find that the central bank does indeed care about smoothing interest rates. For the purpose of this paper, the debate about the true source of observed sluggishness in interest rates is not important. I include the interest rate smoothing term in the loss function to improve the fit of the empirical model but the primary focus is on the weight the central bank puts on inflation relative to output gap.

Here, the main departure from the standard loss function is that the weight on inflation relative to output gap (α_t) is time-varying.[10] There are several reasons to model this parameter as time-varying. First, this weight can change with the changing composition of the FOMC committee where new committee members may be more hawkish or dovish. An extremely clear and uncontroversial example of this weight parameter increasing is the appointment of Chairman Volcker who was a self-subscribed inflation hawk, Meltzer (2006) says the following when talking about the

[9] Changing annual inflation to contemporaneous inflation does not change the results.

[10] The implicit assumption is that the weight on interest rate smoothing is stable over the sample period. I also estimated a version of the model with time-varying ν but the main results are very similar and are available upon request.

changes that Paul Volcker brought to the Federal Reserve: "... he changed the weights on inflation and unemployment" on page 186. Furthermore, other committee members of the FOMC change over time. Second, there is differing degree of political pressure on the Federal Reserve. There are accounts of Presidents Johnson and Nixon putting pressure on Fed chairmen Martin and Burns to refrain from monetary tightening, see Meltzer (2011). There is also evidence of political pressure on the Fed during the current crisis, see Di Maggio (2010). This changing political pressure can be captured by time variation in the weight the Fed puts on inflation relative to output gap, with lower values of α_t representing the case when the Fed gives in to more pressure from politicians to be dovish. The loss function specified here with the time-varying weight on inflation relative to output gap is a flexible and convenient way to capture time variation in central bank preferences.

2.3 Optimal Policy

In each period t, the central bank committee convenes to choose interest rates. The committee agrees on collective preferences within each period, as captured by α_t. I assume that the committee expects this preference parameter to remain constant at the current value in the future. Thus they do not account for the possibility of future updates and the associated uncertainty when formulating the optimal interest rate rule. This assumption is similar to the one made in the learning literature, see Sargent (1993), Sargent (2001) or Evans and Honkapohja (2001) for a detailed exposition and Primiceri (2006) and Sargent et al. (2006) for recent uses. Intuitively this means that the FOMC members do not try to compensate for the possibility that future committee members may be more hawkish or dovish than the current committee members.[11] With this assumption in place the optimization problem can be handled by standard linear quadratic procedures (Sargent (1987)).

The central bank chooses the interest rate i_t to minimize the loss function (3) subject to the equations governing the economy (1) and (2). Appendix A gives the details of the setup of the model in state space form and derives the optimal policy rule. Given the loss function and

[11]This assumption may not be innocuous. In a theory paper Debortoli and Nunes (2014) show that current central bankers change their behavior to account for the possibility that future committees may be more hawkish or dovish than they are. But the empirical quantitative relevance of this effect has not been explored. Here I take a pragmatic approach and ignore this effect to keep the model tractable.

constraints, the optimal policy rule takes the following form.

$$i_t = f_t + F_{1,t}\pi_t + F_{2,t}\pi_{t-1} + F_{3,t}\pi_{t-2} + F_{4,t}\pi_{t-3} + F_{5,t}\tilde{y}_t + F_{6,t}\tilde{y}_{t-1} + F_{7,t}i_{t-1} + F_{8,t}i_{t-2} + F_{9,t}i_{t-3} \quad (4)$$

The coefficients of this rule, $F_{i,t}$, depend on the constant parameters of the constraints $(a_0, a_1, a_2, a_3, b_0, b_1, b_2, b_3, b_4)$, the constant policy preference parameters (β, π^*, ν) and the time-varying preference parameter (α_t). The coefficients of this rule are non-linear functions of these parameters and vary over time as governed by changes in α_t. It is important to emphasize that π^* only affects the intercept term f_t while the weight parameter α_t affects both the intercept f_t and the coefficients $F_{i,t}$. In the estimation we append this equation with a shock e_t.

3 Estimation

The three equations (1),(2) and (4) can be written as a system in the following manner (detailed derivation is in Appendix B).

$$A_{0,t}y_t = A_{1,t} + A_{2,t}y_{t-1} + A_{3,t}y_{t-2} + A_{4,t}y_{t-3} + A_5 y_{t-4} + L\Psi_t\varepsilon_t \quad (5)$$

The $A_{i,t}$ matrices are functions of the time varying preference parameter α_t which is modeled as a random walk

$$\alpha_t = \alpha_{t-1} + v_t \quad (6)$$

where the shock $v_t \sim N(0, Q)$ is assumed to be independent of the shocks to the structural model. The random walk process, quite standard in the literature, is a flexible and parsimonious way of modeling time-varying parameters. It can capture permanent shifts in the preference parameter and involves estimating fewer parameters than a general autoregressive process. This specification even provides a decent approximation in the case that the true data generating process displays a discrete shift. As mentioned earlier, a potential alternative to the time-varying parameter approach is to use a regime-switching framework. Given the anecdotal evidence on preference changes, I discussed above the desirability of having a continuous time-varying parameter approach as opposed to discrete shifts. Additionally, the parameter matrices $A_{i,t}$ are non-linear functions of α_t and the estimation would be considerably more challenging.

As mentioned above the error vector ε_t consists of three shocks: a supply shock (s_t), a demand shock (g_t), and a shock to the interest rate equation (e_t). The innovations in the model are assumed to be jointly normal with the following variance matrix

$$Var\left(\begin{bmatrix} \varepsilon_t \\ v_t \end{bmatrix}\right) = \begin{bmatrix} I & 0 \\ 0 & Q \end{bmatrix}$$

It has been well documented that there is heteroskedasticity in the exogenous shocks in macroeconomic models that use inflation, output and the fed funds rate, for example see Sims and Zha (2006). To account for this, I allow the standard deviations of the shocks to be different pre and post 1984:Q1. This is a well agreed upon date where the economy has experienced a drop in volatility, see Stock and Watson (2003). A triangular decomposition of the variance matrix of the structural errors gives $\Omega_t = L\Psi_t\Psi_t L'$ where

$$\Psi_t = \begin{bmatrix} \sigma_{s,t} & 0 & 0 \\ 0 & \sigma_{g,t} & 0 \\ 0 & 0 & \sigma_{e,t} \end{bmatrix}, L = \begin{bmatrix} 1 & 0 & 0 \\ l_{2,1} & 1 & 0 \\ l_{3,1} & l_{3,2} & 1 \end{bmatrix}$$

where $\sigma_{i,t} = \sigma_{i,1}$ for the pre-1984:Q1 sample while $\sigma_{i,t} = \sigma_{i,2}$ corresponds to the post 1984:Q1 sample, for $i \in \{s, g, e\}$

The model characterized by the constraints (1) and (2), the optimal interest rate rule (4) and the random walk process for the time-varying parameters (6) can now be written in state space form. It is effectively a three equation structural VAR with cross-equation restrictions, drifting parameters in the interest rate equation and standard deviations of the shocks allowed to be different for the split sample. The only source of drift in the coefficients of the interest rate rule is the time variation in the preference parameter α_t. Sims and Zha (2006) estimate unrestricted reduced form VARs and when they allow the coefficients to drift they find the specification that fits the data the best is one where only the coefficients of the monetary policy rule change, consistent with the model here. Appendix B shows how these above equations are transformed to the following nonlinear state space system

$$y_t = h(\alpha_t, X_t, \Gamma, \varepsilon_t) \tag{7}$$

$$\alpha_{t+1} = \alpha_t + v_{t+1} \tag{8}$$

where $\Gamma = [\delta, \nu, L, \Psi_t]$. I estimate the following set of parameters: α_t: time-varying weight on inflation, δ: coefficients of the constraints in (1) and (2), ν: weight on interest rate smoothing, Ψ_t: standard deviations, L: covariance terms, and the hyperparameter Q which represent the variances of the shocks to the preference parameter. I outline the estimation algorithm in the next section.

3.1 Bayesian MCMC Estimation and Priors

Bayesian estimation treats the parameters to be estimated as random variables. Then a prior distribution about these parameters is combined with the likelihood to form the posterior distribution which can be used for inference. The MCMC methodology involves numerically sampling from the posterior distribution which is done using a random walk Metropolis Hastings algorithm. The full details of the estimation procedure are provided in Appendix B and and an online appendix. Here I give a brief overview of the departure from standard state-space estimation techniques. For linear state space models the standard Kalman filter can be used to evaluate the likelihood function. Here the time-varying preference parameter enters non-linearly in the observation equation (7). I use the Extended Kalman Filter (EKF) to tackle the non-linear state-space system. The EKF linearizes the observation equation at each point in time using a first order Taylor expansion and then the standard filtering techniques of the Kalman Filter can be applied. The performance of the EKF depends crucially on the the linearization errors being "small." I show in an online appendix that the non-linearity in this model is not extreme and thus the EKF performs reasonably well. In the same appendix, I also discuss alternative non-linear filtering methods used in the literature.

I use 10 years of data from 1955:Q2 to 1965:Q1 as a training sample to set up the priors. For the training sample I estimate an unrestricted time invariant VAR with OLS which is similar in setup to the full model to compute the prior parameters. The priors for the coefficients of the constraints, δ are assumed to be normal with the prior mean equal to the OLS estimates from the training sample and the prior variance set high enough so that the prior is effectively non-informative. The prior for the interest rate smoothing term ν, is assumed to be uniform on the positive real line. The prior distributions of $\sigma_{i,t}$ are assumed to be inverse-gamma which is a common way of modeling standard deviation parameters. The shape and scale of the inverse-gamma distribution is set to 2 and 1 respectively, implying a fairly loose prior. Finally, the prior for Q, which governs the amount of prior time variation in the preference parameter is set to uniform over the positive real line. My assumptions over the priors can be summarized as the following:

$$
\begin{aligned}
\delta &\sim N(\delta_{OLS}, 10.V_{\delta,OLS}) \\
\sigma_{i,t} &\sim IG(2,1) \\
L &\sim N(0,10) \\
Q &\sim U(0,\infty) \\
\nu &\sim U(0,\infty)
\end{aligned}
$$

4 Results

Two parameters of the loss function (β and π^*) are fixed and the remaining parameters of the model are estimated. I fix the discount factor (β) at 0.99 and the unconditional inflation target (π^*) at 2 percent. The value for β is standard in the literature; decreasing it to 0.95 does not change the results much. The value of 2 for the unconditional inflation target has been used by several others in the literature (Primiceri (2006) and Sargent et al. (2006)). Additionally this number is often reported in the news media as reflecting the Federal Reserve's unofficial target and since 2012 the FOMC has officially announced an inflation target of 2 percent. The data sample goes from 1965:Q2 to 2007:Q2. I stop before the financial crisis as dealing with the zero lower bound and unconventional monetary policy is outside the scope of this paper.

The bottom panel in figure 1 shows the data used in the estimation. The top panel in figure 1 plots the smoothed value of the time varying preference parameter at the posterior mean of the estimates along with one standard deviation confidence bands. These confidence bands take into account both the filter uncertainty and parameter uncertainty.[12] As expected there is a rise in the weight on inflation around the appointment of Paul Volcker. But interestingly the rise starts just before Volcker's appointment. The weight falls in the early 1980s before rising again towards the end of Volcker's tenure. With Alan Greenspan stepping in, there is a steady decline in the weight till the end of his tenure where there is a rise that is continued in the early Ben Bernanke years. There is a drastic rise at the end of the sample, but the preference parameter is not accurately estimated towards the end of the sample as is evident by the wide confidence bands. It may seem surprising that when inflation was near its peak in the early 1980s is also when the weight on inflation is relatively high. But these two are consistent and suggests that the shocks hitting the economy drove inflation up despite the high Federal Reserve weight on inflation. The counterfactual exercise below explores this idea in more detail. In section 7, I discuss several robustness exercises including using alternative measures of the output gap and an alternative specification for the state equation (8). Those estimates show a similar overall trend in the time-varying preference parameter. Thus the results suggest that it would be misguided to treat the preference parameter as constant or allow a one-time discrete change, as is common in the literature. Even if a discrete change were imposed it is not clear when the discrete change should be modeled as there are long periods of gradual changes. Given the increased interest in estimating dynamic models with optimal monetary policy, careful attention needs to be paid to the specification of preferences.

To explore the implications for the optimal policy rule, consider the long-run response to

[12] I employ the Monte Carlo algorithm outlined in Hamilton (1986) to calculate the confidence bands.

inflation in this model given by

$$\beta_t \equiv \frac{F_{1,t} + F_{2,t} + F_{3,t} + F_{4,t}}{1 - F_{7,t} - F_{8,t} - F_{9,t}}$$

There are two components: the direct response to a change in inflation captured by $F_{1,t}+F_{2,t}+F_{3,t}+F_{4,t}$ and the component coming from the inertial response of the interest rate which is governed by $F_{7,t} + F_{8,t} + F_{9,t}$. In figure 2 below, I plot both these components along with the long-run response to inflation. The long-run response averages a little over 3 in the full sample, and is higher in the 80s and early 90s relative to the mid 70s and mid 90s. The direct response to inflation has the same dynamic pattern as the time-varying weight on inflation α_t. However, the amount of inertia implied by the optimal interest rate rule is inversely related to α_t. Thus even though α_t falls in the early Volcker years, we see that the inertial component of the interest rate rule actually rises, keeping the long-run response to inflation from falling much. Similarly, the long-run response to inflation stays above 1 under Greenspan in the mid 90s when there is a considerable fall in α_t.

To further highlight the importance of the dynamics of the estimated preference parameter I compare the impulse responses to inflation under Alan Greenspan and Paul Volcker. The commonly used assumption of a one-time discrete shift in preferences with the appointment of Paul Volcker implies that the behavior of the Federal Reserve under Volcker and Greenspan regimes is the same. Panel (a) in figure 3 shows the impulse response to a 1 unit inflation shock for Volcker and Greenspan regimes taking the average of the estimated weights on inflation for each Chairman's regime.[13] The responses are almost identical as the average of the weight under Volcker is not much higher than under Greenspan. But this figure does not reveal the complete picture since there is time variation within each regime. Depending on the the specific time period chosen, there can be a big difference between the response of the Fed. To show how large this difference can be, I compare the response to an inflation shock using the lowest value of the preference parameter estimated under Greenspan (mid 1990s) and the highest value under Volcker (last year of his regime). Panel (b) in figure 3 shows that the responses are quite different. Under the"high weight" Volcker regime, the fed funds

[13]A unit shock is used instead of a one standard deviation shock to make the proper comparison as the standard deviations of the shock to inflation are different across the two time periods.

rate would have been raised considerably more to bring down inflation. Whereas the response of the the "low weight" Greenspan regime would have been to avoid the fall in the output gap at the expense of higher inflation. In fact using the maximum weight estimated under Greenspan and minimum under Volcker gives similar differences in the implied responses. The main point is to highlight the fact that treating the preferences as the same under the two chairmen can be somewhat misleading when analyzing the historical behavior of the Federal Reserve.

The parameter estimates along with 95^{th} and 5^{th} percentiles of the posterior distribution are listed in Table 1. First, I focus on the parameters of the private sector. To facilitate a comparison with the literature, table 2 presents results from two other papers that have used this model: Rudebusch and Svensson (1998) and Dennis (2006) represented by RS and Dennis. To keep the comparison as clean as possible, I re-estimate these models using the same data and also including the split in the sample for standard deviation as in the benchmark model. Additionally, the last two columns (RS Original) and (Dennis Original) present the parameter estimates from the published versions using the original data, 1961:Q1 to 1996:Q2 for Rudebusch and Svensson (1998) and 1982:Q1 to 2000:Q2 for Dennis (2006). For the Phillips curve the parameter estimates from the different models and samples are quite similar. For the IS curve, these parameters are also similar except for a_3, which governs the sensitivity of the output gap to the real interest rate. This sensitivity parameter is lower for the models with optimal policy relative to the RS specifications. Moreover, allowing for time variation in the weight on inflation lower it even more relative to the Dennis specification.[14]

Second, for the loss function parameters, the estimate of the weight on interest rate smoothing term, ν might seem high. However, to compare this value to the existing literature we need to divide by the weight on inflation.[15] This implies that the parameter comparable to the existing literature averages around 30. The magnitude of this parameter in the related literature is very sensitive to the specified model. For example, estimates of this weight range from 0.0051 in Favero and Rovelli (2003) to 37.168 in Dennis (2006) to 2131 in Primiceri (2006). Moreover, Castelnuovo

[14]For the "Dennis" column, I fix $\beta = .99$ and $\pi^* = 2$ following the benchmark specification.

[15]This is because of the different normalization employed in this paper.

(2006) finds that specifying a forward looking model reduces the estimated weight on interest rate smoothing. Finally, adding an additional term in the loss function involving the squared deviation of the interest rate (interest rate variability) reduces this estimated value even further. As mentioned earlier there is a debate in the literature regarding the true source of the observed sluggishness in the policy instrument. Here the interest rate smoothing term is not important as we are primarily concerned with the dynamic behavior of the weight on inflation versus output gap.

5 Counterfactual Analysis

In this section I conduct two counterfactual analyses to gain insight into two important phenomena of the post war period; the Great Inflation episode of the 1970s and the ensuing Great Moderation.

5.1 The Great Inflation

Here I pose the following hypothetical question: Would the early appointment of Paul Volcker have avoided the inflation episode of the 1970s? To answer this question I simulate the path of inflation by freezing preferences at the average value estimated for the Volcker regime.[16] First I compute the structural shocks for the estimated model. Next the counterfactual inflation is initialized to the actual level of inflation in the first year of the sample. I then simulate a path of inflation using "Volcker-style" preferences but fix shocks hitting the economy at the estimated shocks from the original results. This gives the simulated path of inflation with an aggressive Paul Volcker hypothetically heading the Fed, but with the same true observed shocks hitting the economy; thus the only difference between the paths is the change in preferences.

Figure 4 shows this simulated path of inflation along with the actual inflation path. The results suggest that inflation would have been lower under Volcker in the 1970s (and also in the 1980s), but not low enough to avoid the high inflation episode. This is complementary to the

[16] Using the highest estimated value for preferences under Volcker's tenure produces a slightly bigger fall in the simulated path of inflation.

findings of Liu et al. (2011) who show that the high inflation episode cannot be explained by the Federal Reserve having a high inflation target in the 1970s. The results suggest that the structural shocks hitting the economy seemed to have played a big role in the 1970s. However, since the economy is described by a backward looking model, agents adapt their expectations of Volcker's "appointment" with a lag. If agents were instead forward-looking they would lower their inflation expectations sooner. For example, Bianchi (2013) performs a counterfactual where he imposes beliefs on the agents such that they expect that an extremely hawkish chairman is going to be appointed in the future. He finds this change in expectations dramatically lowers the simulated path of inflation. Thus I interpret the fall in inflation found here to be the lower bound of the effects that different Fed preferences would have had on inflation in the 1970s.

5.2 The Great Moderation

The causes of the fall in the volatilities of macroeconomic variables since the early 1980s have been widely debated in the literature. Three main explanations have emerged: good luck, good policy and structural change. The good luck hypothesis states that the economy has been subject to more fortuitous shocks since the 1980s (Stock and Watson (2003)). The good policy explanation mainly focuses on the improved policy of the Federal Reserve (Clarida et al. (2000), Boivin and Giannoni (2006)). The structural change explanation attributes the lower volatility to other shifts in the structure of the economy such as better inventory management (Ramey (2006), McConnell and Perez-Quiros (2000)). The model here allows me to test the importance of the first two channels.

The first row of Table 3 documents the decline in volatility in the data. The standard deviations of inflation, output gap and fed funds rate have each fallen by at least 30%. The next row lists the model implied values and shows that the model does a good job of capturing the fall seen in the data. In the first counterfactual I assess the effect of removing the observed fall in the standard deviation of the residuals. The three macro variables are simulated using the estimated shocks and parameters but fixing the standard deviation of the shocks to their value in the pre 1984:Q1 period. The thought experiment is the following: How would the aggregate variables have behaved if there had not been the observed decline in the standard deviations of the shocks hitting the economy? If the volatilities of the counterfactual values do not display a similar fall between the pre-1984:Q1 and post-1984:Q1 samples, then we can conclude that the size of the shocks played a big role in the Great Moderation. What we observe in the third row is indeed that the fall in the volatilities of counterfactual values is much smaller for inflation and the counterfactual volatilities for output gap and the fed funds rate are in fact higher. This suggests that without the fall in the standard deviation of the shocks, instead of the Great Moderation we would have experienced even higher volatilities than in the pre-1984:Q1 sample.

The second counterfactual is similarly performed but here the value of the preference parameter is fixed to its pre-1984:Q1 mean. The results in the fourth row suggest that change in preferences did not seem to play an important part in the Great Moderation. We see a drop in the standard deviation of all three variables that is similar to the observed data.

It is important to keep in mind the details of this specific model when interpreting the counterfactual results regarding the Great Moderation. First, this model uses the output gap while most of the literature has documented the Great Moderation using real output. Second, even though the change in preferences does not appear crucial, it does not necessarily imply that monetary policy did not play a role. For example, it is possible that the monetary authority became better at evaluating and forecasting economic conditions, which improved the overall conduct of policy. This could have helped lower the volatility of aggregate variables but this channel is not captured here.

6 A New Measure of Monetary Policy Shocks

There is a large literature that tries to evaluate the effects of monetary policy. But systematic monetary policy decisions are endogenous with respect to developments in the economy and cannot help identify the effects of monetary policy on the economy. Thus there has been much interest in identifying exogenous measures of monetary policy shocks to help us understand how monetary policy decisions affect the economy. Vector Autoregressions (VAR) are quite often used in identifying monetary policy shocks. While there is disagreement about the identifying assumptions and the specific variables to use in the VAR, there is reasonable consensus regarding the use of the residual of the interest rate equation as the measure of monetary policy shock. However, it is not clear how one should interpret these shocks. I propose using the exogenous change in the preference parameter as a new measure of monetary policy shocks. This has intuitive appeal because this represents changes in the fundamental behavior of monetary policy decision making that are not endogenous to economic developments. Additionally Christiano et al. (1999) suggest this measure as one of their interpretations for the monetary policy shock in VARs. They say that one interpretation of monetary policy shocks is "... exogenous shocks to the preferences of the monetary authority, perhaps due to stochastic shifts in the relative weight given to unemployment and inflation". Other candidate explanations that they suggest are technical factors like measurement error and strategic factors related to private agents' expectations. Using the estimates in this paper, I can evaluate the response of the economy to preference shocks and compare that to the response to the composite monetary policy shock measure used in the literature.

I start with a simple quarterly reduced-form VAR with the following variables: 1) Y_t: log of real GDP, 2) P_t : Log of GDP Deflator and 3) R_t: monetary policy indicator,[17] estimated using 8 lags. This is a simplified version of the benchmark VAR used in Christiano et al. (1999) and many other papers in the literature. Identifying restrictions are required to estimate the response to a structural monetary policy. A common approach is to put zero restrictions on the matrix that governs the contemporaneous relationship between measures of economic activity and the monetary policy indicator. As discussed in Christiano et al. (1999), a popular identification scheme is the recursive or Cholesky assumption. In this framework the ordering of the variables determines which variables are allowed to react contemporaneously to the others. The most common recursive identification assumption is to order the variables $[Y_t, P_t, R_t]$. Thus monetary policy cannot contemporaneously affect output or prices but does respond to their current values.

To compare the effects of preference shocks to this recursive approach, I use two different

[17] A constant and a time trend are added in the estimation of the reduced form VAR.

methods. The first one is the external instruments identification used in recent work by Stock and Watson (2012) and Gertler and Karadi (2015). This methodology requires an instrument which has to be correlated with the structural shock of interest (shock to the monetary policy indicator) and uncorrelated with the other structural shocks in the VAR. Given this assumption, the impulse responses to a monetary policy shock can be estimated. I will use the change in the estimated preference parameter ($\hat{v}_t$) as the instrument. One concern with the external instruments methodology is the weak instruments problem. To check this, I regress the reduced-form residual in the fed funds rate equation on $\hat{v}_t$. The F-statistic calculated using robust standard errors is 6.62 and the R-squared from the regression is 0.21. Thus the instrument seems to be sufficiently correlated with the reduced form residuals and explains a non-negligible proportion of the variation. Second, I follow the analysis in Romer and Romer (2004) (R&R) where they construct an exogenous measure of monetary policy shocks and then embed it in a VAR. Specifically, they cumulate their measure of monetary policy shocks to represent the stance of monetary policy. This cumulated measure is ordered last and the ordering is justified based on the exogenous nature of the shocks. To apply this methodology, I enter the cumulated shock series $\lambda_t = (\sum_{j=1}^{t} \hat{v}_j)$ as the monetary policy indicator. Thus the Cholesky and External Instruments VAR both try to identify the effects of a "conventional" monetary policy shock (i.e. a shock to the fed funds rate equation) where the latter uses the preference estimates as an instrument for identification. On the other hand, the R&R methodology captures the sole effect of preference shocks.

The top panel of figure 5 plots the estimated structural monetary policy shocks from the VAR. The solid blue line shows the case where the Cholesky (recursive) identification scheme is used, while the dashed red line shows the case with the external instruments identification. The two shock series are quite similar and the correlation between them is 0.93. The bottom panel plots the estimated residuals of the preference parameter equation, $\hat{v}_t = \alpha_t - \alpha_{t-1}$ obtained from the filtered estimates. A positive value for the shock can be interpreted as a more hawkish stance of monetary policy, similar in concept to a contractionary monetary policy shock. There are large negative shocks in the mid 1970s followed by large positive and negative shocks in the early 80s after which the shocks are smaller in magnitude until early 2000s.

Figure 6 shows the impulse responses to a one standard deviation monetary policy shock. The response of output to a shock using the Cholesky identification (solid blue line) is familiar; there is a humped shape response of output that peaks around two years. The response with the external instruments identification is similar, but output falls less in the short-run and a little more after the 3 year mark compared to the Cholesky case. Looking at the response using the R&R methodology, there is a delayed response of output and the peak effect is reached after three years. The middle panel shows the response of the price level. The Cholesky response displays the so-called "price puzzle" where prices initially rise after a contractionary monetary policy shock before they start falling. The External Instruments response shows a more severe version of this response, where prices don't start falling till the 3 year mark. The R&R methodology leads to a smaller impact on prices, with a fall initially. Finally, the third panel shows the response of the monetary policy indicator. For the Cholesky and External Instruments VAR this is the fed funds rate, while for shows the persistent response of the fed funds rate under the Cholesky and External Instruments case and even more persistent response of the cumulated preference shock measure. Thus overall the preference shocks appear to have a more persistent effect on output relative to prices.

7 Robustness Checks

In this section, I estimate the model using alternative measures of the output gap and using different dates for the break in the variance of the shocks and compare it with the baseline results.

7.1 Alternative and Real Time Measures of Output Gap

The assumption in the benchmark estimates is that the Federal Reserve is conducting optimal policy using "revised" data, i.e. data that is available to the econometrician at the end of the sample. But in reality the Federal Reserve has access to only preliminary data when they make monetary policy decisions and these data typically get revised, sometimes significantly. Orphanides (2001) and the related literature has pointed out that this was a concern with Federal Reserve policy in the 1970s. To check the robustness of my results to this concern, I estimate the model using a measure of output gap constructed from real-time GDP data. Unfortunately there is no publicly available series for real-time output gap. The Federal Reserve Bank of Philadelphia does maintain a database of real time macroeconomic variables, but their series for the output gap starts in the late 1980s. To get around this problem, I use real-time GDP data and construct measures of real-time output gap. Specifically I use a HP filter and a quadratic trend to construct measures of real-time output gap. Since these measure are not directly comparable to the CBO's measure of the output gap (which uses revised data), in the top two panesl of figure 7 I plot the output gap constructed with the quadratic trend and HP filter using both revised and real-time GDP data. The top left panel shows the quadratic trend estimates and the top right panel shows the HP filtered estimates. Overall there are significant differences between the CBO measure and the output gaps constructed using revised and real-time GDP data, but the broad patterns look similar.[18]

I re-estimate the model using these alternative measures of the output gap. For the output gap estimates with revised data, the estimation just involves replacing the CBO's measure of the output gap with the corresponding quadratic trend or HP filtered measures. But it is more involved when using the real-time estimates. The model with real-time data implies that the Federal Reserve is setting policy based on their assessment of the economy using real-time data. But the dynamics of the real economy are governed by revised data. Here I take a simple approach and assume that the Federal Reserve ignores any uncertainty about the underlying structure of the economy and data revisions. Thus the Fed is setting the fed funds rate plugging in real-time data in the optimal policy rule. But there is no effect of this use of "wrong" data on the economy other than through the direct effect of the fed funds rate.

The literature has dealt with this issue in different ways. One way to setup this economy is to explicitly model the learning behavior of the Federal Reserve and its interaction with the economy to find a self-confirming equilibrium, as in Sargent et al. (2006) and Primiceri (2006).

[18]Note that these measures only deal with the uncertainty stemming from data revisions and not the uncertainty stemming from measurement of the trend of output, see for example Orphanides and van Norden (2002). This extension is left for future research.

An alternative approach is to have a situation where the state of the economy is unknown but the structural model is known with certainty, as in Givens and Salemi (2015). They assume that the econometrician observes the real-time data in addition to the ex post (final release) data while the agents never observe the revised data. My assumption is closer in spirit to this latter setting and is made to keep the model tractable.[19] This assumption gives rise to a simple model which is estimated plugging in the real-time output gap in the optimal interest rate equation (4) but still using the revised data in the equations that describe the dynamics of the structural economy ((1) and 2).

The bottom panel of figure 7 shows the estimate of the time-varying parameter with the quadratic trend estimate of the output gap using both revised and real-time data.[20] The estimate of the variance of the shock to the weight parameter is lower with these alternate measures of output gap, which explains the less volatile estimates seen in the figure. But the overall pattern is similar to the baseline estimates using the CBO output gap. The weight on inflation rises in the late-1970s. There is a steady decline in the weight starting in the late 1980s, and then a rise towards the end of the sample. Finally, the results of the counterfactual analyses and the effects of monetary policy shocks on output from these alternate estimates are very similar to the baseline results.

7.2 Break dates for shock variances

In this section I re-estimate the model changing the break date for the variance of the shocks. To check the robustness of the baseline results, I estimate the model with 3 different break dates: i) 1979:Q3 to correspond to Paul Volcker's appointment, ii) 1987:Q3 to correspond to Alan Greenspan's appointment and iii) break at both 1979:Q3 and 1987:Q3 (so that there are three sets of parameters to be estimated for the shock variances). The dashed blue lines in panels (b), (c) and (d) of figure 8 show the smoothed estimate of the weight parameter for these three specifications. Specifications ii) and iii) display very similar dynamics to the baseline case (shown in solid

[19] Modeling the self-confirming equilibrium setup would make the non-linear estimation procedure intractable.
[20] The results from the HP filtered output gaps are similar and not included here.

black line) with a slightly bigger rise in the mid-80s. For specification i) with a break in 1979:Q3 we see that while the general pattern is similar, there are more high-frequency movements in the early 1980s. When I compare the fit of these specifications to the baseline case using the Bayesian Information Criterion (shown in table 4), it becomes clear that the data prefer a break in the shock variances around 1984. The baseline specification has the best fit while specification i)has the worse fit.

7.3 Alternative Specification of State Equation

In the baseline specification, the dynamics of the time-varying parameter is modeled using a random walk equation. As discussed above, this is a parsimonious approach to capture time variation that is the most popular method in this literature. Here we consider an alternative specification following Ireland (2007), where the time-varying parameter is allowed to respond to the exogenous shocks in the structural model. Thus the new process is specified as

$$\alpha_t = \alpha_{t-1} + \theta_s s_t + \theta_g g_t + v_t$$

The estimate of the constant parameters of the model from this specification are very similar to the baseline case and are not shown here. The posterior mean of two new parameters, θ_s and θ_g are 0.14 and 0.69 with 95% credibility intervals $[-.14, \quad .49]$ and $[.37, \quad 1.07]$ respectively. The positive sign of the posterior mean suggests that positive supply and demand shocks are associated with a rise in the weight on inflation. Note however that only θ_g is statistically significantly different from 0. The top left panel of figure 8 shows the smoothed estimate of α_t from this specification, together with the benchmark specification. The broad pattern of the estimates with the new specification is very similar to the benchmark case but with some more high frequency movements around the late 1970s. Moreover, the fit of the baseline model is almost identical to this new specification, as computed using the Bayesian Information Criterion (shown in table 4). Thus I conclude that this specification produces very similar results to the baseline case.

8 Conclusion

Using a simple model of optimizing central bank behavior, I estimate a continuously time-varying series of the weight on inflation relative to output gap. This parameter enters non-linearly in the model and is estimated with a Bayesian Markov Chain Monte Carlo algorithm that uses non-linear filtering techniques. Consistent with the anecdotal evidence, the results show that there is a large rise in the weight on inflation with the appointment of Paul Volcker. However, the weight parameter displays instability even in the pre- and post-Volcker periods. The estimation of the time-varying preference parameter also provides a novel measure of monetary policy shocks. I use this measure of monetary policy shocks in a VAR to identify the monetary transmission mechanism with two different identification strategies. The response of output to this new measure of monetary policy shocks is more persistent relative to the response to conventional measures constructed from the residuals of the short rate equation. This provides complementary evidence for the VAR literature which has extensively evaluated the impact of monetary policy shocks.

There has been a lot of work trying to understand how Federal Reserve behavior has changed. One simple way to think about these changes is to broadly categorize it into two fields: 1) changes coming from policy mistakes and 2) changes coming from policy preferences. The existing literature has mostly focused on the former (Primiceri (2006), Sargent et al. (2006) and Orphanides(2003)). The basic idea is that the Federal Reserve made mistakes in evaluating the state of the economy or the dynamics governing the economy. While existing literature has given ample evidence in support of this line of thinking, the main motivation for this paper is that not enough attention has been paid to the role of policy preferences and the implied consequences of its dynamics. The findings in this paper motivate a unifying framework that takes into account both policy preferences and policy mistakes as a promising area of future research.

References

Best, G. (forthcoming). Policy Preferences and Policymakers' Beliefs: The Great Inflation. *Macroeconomic Dynamics*.

Bianchi, F. (2013). Regime Switches, Agents' Beliefs, and Post-World War II U.S. Macroeconomic Dynamics. *The Review of Economic Studies 80*(2), 463–490.

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

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

Borger, S., J. D. Hamilton, and S. Pruitt (2011). Estimating the Market-Perceived Monetary Policy Rule. *American Economic Journal: Macroeconomics 3*(3), 1–28.

Burns, A. (1987). The Anguish of Central Banking. *Federal Reserve Bulletin 73*(3), 697–98.

Castelnuovo, E. (2006). The Fed's Preference for Policy Rate Smoothing: Overestimation Due to Misspecification? *Topics in Macroeconomics 6*(2).

Castelnuovo, E. and P. Surico (2004). Model Uncertainty, Optimal Monetary Policy and the Preferences of the Fed. *Scottish Journal of Political Economy 51*(1), 105–126.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary Policy Shocks: What Have We Learned and To What End? Volume 1, Part A of *Handbook of Macroeconomics*, pp. 65 – 148. Elsevier.

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

Cogley, T. and T. J. Sargent (2005). Drifts and Volatilities: Monetary Policies and Outcomes in the post WWII US. *Review of Economic Dynamics 8*(2), 262 – 302. Monetary Policy and Learning.

Coibion, O. and Y. Gorodnichenko (2012). Why Are Target Interest Rate Changes so Persistent? *American Economic Journal: Macroeconomics 4*(4), 126–162.

Debortoli, D. and R. Nunes (2014). Monetary regime switches and central bank preferences. *Journal of Money, Credit and Banking 46*(8), 1591–1626.

Dennis, R. (2006). The Policy Preferences of the US Federal Reserve. *Journal of Applied Econometrics 21*(1), 55 – 77.

Di Maggio, M. (2010). The Political Economy of the Yield Curve. *Working Paper*.

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

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.

Fernandez-Villaverde, J., P. Guerron-Quintana, and J. F. Rubio-Ramirez (2010). Fortune or Virtue: Time-Variant Volatilities Versus Parameter Drifting in U.S. Data. *Federal Reserve Bank of Philadelphia Working Paper*.

Gertler, M. and P. Karadi (2015). Monetary Policy Surprises, Credit Costs, and Economic Activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

Givens, G. E. (2012). Estimating central bank preferences under commitment and discretion. *Journal of Money, Credit and Banking 44*(6), 1033–1061.

Givens, G. E. and M. K. Salemi (2015). Inferring monetary policy objectives with a partially observed state. *Journal of Economic Dynamics and Control 52*, 190–208.

Greenspan, A. (2008). *The Age of Turbulence: Adventures in a New World*. Penguin.

Hamilton, J. D. (1986). A Standard Error for the Estimated State Vector of a State-space Model. *Journal of Econometrics 33*(3), 387 – 397.

Ilbas, P. (2012). Revealing the preferences of the US Federal Reserve. *Journal of Applied Econometrics 27*(3), 440–473.

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.

Kim, C.-J. and C. R. Nelson (2006). Estimation of a Forward-looking Monetary Policy Rule: A Time-varying Parameter Model Using Ex Post Data. *Journal of Monetary Economics 53*(8), 1949 – 1966.

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. A. and F. Schorfheide (1 March 2004). Testing for Indeterminacy: An Application to U.S. Monetary Policy. *The American Economic Review 94*, 190–217(28).

McConnell, M. M. and G. Perez-Quiros (2000). Output Fluctuations in the United States: What Has Changed Since the Early 1980's? *The American Economic Review 90*(5), pp. 1464–1476.

Meltzer, A. H. (2006). From Inflation to More Inflation, Disinflation, and Low Inflation. *The American Economic Review 96*(2), pp. 185–188.

Meltzer, A. H. (2011). Politics and the Fed. *Journal of Monetary Economics 58*(1), 39 – 48. Carnegie-Rochester Conference Series on Public Policy: The Future of Central Banking April 16-17, 2010.

Orphanides, A. (2001). Monetary Policy Rules Based on Real-Time Data. *The American Economic Review 91*(4), pp. 964–985.

Orphanides, A. and S. van Norden (2002). The Unreliability of Output-Gap Estimates in Real Time. *Review of Economics and Statistics 84*(4), 569–583.

Owyang, M. T. and G. Ramey (2004). Regime Switching and Monetary Policy Measurement. *Journal of Monetary Economics 51*(8), 1577 – 1597.

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. (2006). Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy. *Quarterly Journal of Economics 121*(3), 867–901.

Ramey, V. A. (2006). Declining Volatility in the U.S. Automobile Industry. *American Economic Review 96*(5), 1876–1889.

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

Rudebusch, G. D. (2002). Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia. *Journal of Monetary Economics 49*(6), 1161 – 1187.

Rudebusch, G. D. (2006). Monetary Policy Inertia: Fact or Fiction? *International Journal of Central Banking 2*(4), 85–135.

Rudebusch, G. D. and L. E. O. Svensson (1998, April). Policy Rules for Inflation Targeting. Working Paper 6512, National Bureau of Economic Research.

Salemi, M. K. (2006). Econometric Policy Evaluation and Inverse Control. *Journal of Money, Credit, and Banking 38.7*, 1737–1764.

Sargent, T. (1987). *Dynamic Macroeconomic Theory, 2nd ed.* Harvard University Press: Cambridge, MA.

Sargent, T. (1993). *Bounded Rationality in Macroeconomics: The Arne Ryde Memorial Lectures.* Oxford University Press.

Sargent, T. (2001). *The Conquest of American Inflation.* Princeton University Press.

Sargent, T., N. Williams, and T. Zha (2006). Shocks and Government Beliefs: The Rise and Fall of American Inflation. *American Economic Review 96*(4), 1193–1224.

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

Stock, J. H. and M. Watson (2012). Disentangling the channels of the 2007-09 recession. *Brookings Papers on Economic Activity: Spring 2012*, 81.

Stock, J. H. and M. W. Watson (2003). Has the Business Cycle Changed and Why? In *NBER Macroeconomics Annual 2002, Volume 17*, pp. 159–230. MIT Press.

Taylor, J. (1999). *Monetary Policy Rules.* University of Chicago Press: Chicago.

Woodford, M. (2003). Optimal interest-rate smoothing. *The Review of Economic Studies 70*(4), 861–886.

Appendix

Appendix A: Derivation of the optimal policy rule

To apply the linear quadratic regulator of Sargent (1987) I start by putting the constraints (1) and (2) in the following state space form where z_t is the state vector and x_t is the control variable.

$$z_{t+1} = C + Dz_t + Bx_t + u_{t+1}$$

where $z_t \equiv [\pi_t, \pi_{t-1}, \pi_{t-2}, \pi_{t-3}, \tilde{y}_t, \tilde{y}_{t-1}, i_{t-1}, i_{t-2}, i_{t-3}]'$, $x_t \equiv [i_t]$ and $u_{t+1} \equiv [s_t, g_t, 0]'$,

$$C = [b_0, 0, 0, 0, a_0, 0, 0, 0, 0]'$$
$$B = [0, 0, 0, 0, \frac{a_3}{4}, 0, 1, 0, 0]'$$
$$D = \begin{bmatrix} b_1 & b_2 & b_3 & 1-b_1-b_2-b_3 & b_4 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -\frac{a_3}{4} & -\frac{a_3}{4} & -\frac{a_3}{4} & -\frac{a_3}{4} & a_1 & a_2 & \frac{a_3}{4} & \frac{a_3}{4} & \frac{a_3}{4} \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

Following Dennis (2006) I rewrite the loss function (3) in the following way, see Sargent (1987) for more details.

$$\begin{aligned} L \;\; = \;\; & \tilde{E}_t \sum_{j=0}^{\infty} \beta^j [(z_{t+j} - \bar{z})' W_t (z_{t+j} - \bar{z}) + (x_{t+j} - \bar{x})' N (x_{t+j} - \bar{x}) \\ & +2(z_{t+j} - \bar{z})' H (x_{t+j} - \bar{x}) + 2(x_{t+j} - \bar{x})' G (z_{t+j} - \bar{z})] \end{aligned}$$

where

$$
\begin{aligned}
W_t &= P' R_t P \\
P &= \begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \\
R_t &= \begin{bmatrix} \alpha_t & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \nu \end{bmatrix} \\
H' &= G = \left[0_{1\text{x}6} \ , -\frac{\nu}{2}, 0_{1\text{x}2} \right] \\
N &= \nu
\end{aligned}
$$

The optimal rule for x_t is then given by

$$
\begin{aligned}
x_t &= \bar{x} - F_t \bar{z} \\
F_t &= -\left(N + \beta B' M_t B \right)^{-1} \left(H' + G + \beta B' M_t D \right) \\
M_t &= W_t + F_t' N F_t + 2 H F_t + 2 F_t' G + \beta \left(D + B F_t \right)' M_t \left(D + B F_t \right)
\end{aligned}
$$

where the first equation corresponds to equation 4 in section 2.3.

Appendix B: Setup of model for Bayesian estimation

I first start by stacking equations 1,2 and 4 to get the following form

$$A_{0,t} y_t = A_{1,t} + A_{2,t} y_{t-1} + A_{3,t} y_{t-2} + A_{4,t} y_{t-3} + A_{5,t} y_{t-4} + L \Psi_t \varepsilon_t$$

$$\text{where} \quad y_t \equiv [\pi_t, \tilde{y}_t, i_t]' \ , \ \ \varepsilon_t \equiv [s_t, g_t, e_t]'$$

$$
\begin{aligned}
A_{0,t} &= \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -F_{1,t} & -F_{5,t} & 1 \end{bmatrix} \qquad A_{1,t} = \begin{bmatrix} b_0 \\ a_0 \\ f_t \end{bmatrix} \qquad A_{2,t} = \begin{bmatrix} b_1 & b_4 & 0 \\ -\frac{a_3}{4} & a_1 & \frac{a_3}{4} \\ F_{2,t} & F_{6,t} & F_{7,t} \end{bmatrix} \\
A_{3,t} &= \begin{bmatrix} b_2 & 0 & 0 \\ -\frac{a_3}{4} & a_2 & \frac{a_3}{4} \\ F_{3,t} & 0 & F_{8,t} \end{bmatrix} \quad A_{4,t} = \begin{bmatrix} b_3 & 0 & 0 \\ -\frac{a_3}{4} & 0 & \frac{a_3}{4} \\ F_{4,t} & 0 & F_{9,t} \end{bmatrix} \quad A_5 = \begin{bmatrix} 1 - b_1 - b_2 - b_3 & 0 & 0 \\ -\frac{a_3}{4} & 0 & \frac{a_3}{4} \\ 0 & 0 & 0 \end{bmatrix}
\end{aligned}
$$

Note an error term e_t is added to the interest rate equation. $F_{i,t}$ refers to the ith component of the optimal policy coefficient vector F_t and are functions of α_t, δ, ν as governed by the optimal policy restrictions. Now pre-multiply both sides by $A_{0,t}^{-1}$

$$\begin{aligned} y_t &= A_{0,t}^{-1}A_{1,t} + A_{0,t}^{-1}A_{2,t}y_{t-1} + A_{0,t}^{-1}A_{3,t}y_{t-2} + A_{0,t}^{-1}A_{4,t}y_{t-3} + A_{0,t}^{-1}A_{5,t}y_{t-4} + A_{0,t}^{-1}L\Psi_t\varepsilon_t \\ &= B_{1,t} + B_{2,t}y_{t-1} + B_{3,t}y_{t-2} + B_{4,t}y_{t-3} + B_{5,t}y_{t-4} + A_{0,t}^{-1}L\Psi_t\varepsilon_t \end{aligned}$$

$$\begin{aligned} B_{1,t} &= \begin{bmatrix} b_0 \\ a_0 \\ f_t + F_{5,t}a_0 + F_{1,t}b_0 \end{bmatrix}, \; B_{2,t} = \begin{bmatrix} b_1 & b_4 & 0 \\ -\frac{a_3}{4} & a_1 & \frac{a_3}{4} \\ F_{2,t} - \frac{F_{5,t}a_3}{4} + F_{1,t}b_1 & F_{6,t} + F_{5,t}a_1 + F_{1,t}b_4 & F_{7,t} + \frac{F_{5,t}a_3}{4} \end{bmatrix} \\ B_{3,t} &= \begin{bmatrix} b_2 & 0 & 0 \\ -\frac{a_3}{4} & a_2 & \frac{a_3}{4} \\ F_{3,t} - \frac{F_{5,t}a_3}{4} + F_{1,t}b_2 & F_{5,t}a_1 & F_8 + \frac{F_5 a_3}{4} \end{bmatrix} \\ B_{4,t} &= \begin{bmatrix} b_3 & 0 & 0 \\ -\frac{a_3}{4} & 0 & \frac{a_3}{4} \\ F_{4,t} - \frac{F_{5,t}a_3}{4} + F_{1,t}b_3 & 0 & F_{9,t} + \frac{F_{5,t}a_3}{4} \end{bmatrix}, \; B_{5,t} = \begin{bmatrix} 1 - b_1 - b_2 - b_3 & 0 & 0 \\ -\frac{a_3}{4} & 0 & \frac{a_3}{4} \\ -\frac{F_{5,t}a_3}{4} - F_{1,t}(b_1 + b_2 + b_3 - 1) & 0 & \frac{F_{5,t}a_3}{4} \end{bmatrix} \end{aligned}$$

Finally we write the model as a non-linear state space system

$$y_t = h(\alpha_t, X_t, \Gamma, \varepsilon_t) \tag{9}$$

$$\alpha_{t+1} = \alpha_t + v_{t+1} \qquad \text{where } v_t \sim N(0, Q) \tag{10}$$

where $h(\alpha_t, X_t, \Gamma\varepsilon_t) = B_{1,t} + B_{2,t}y_{t-1} + B_{3,t}y_{t-2} + B_{4,t}y_{t-3} + B_{5,t}y_{t-4} + A_{0,t}^{-1}L\Psi_t\varepsilon_t$ and $\Gamma = [\delta, \nu, L, \Psi_t]$

	IS Curve				Phillips Curve		
Parameter	Pos. Mean	95 percent	5 percent	Parameter	Pos. Mean	95 percent	5 percent
a_0	0.035	0.096	-0.020	b_0	-0.062	0.014	-0.141
a_1	1.269	1.355	1.179	b_1	0.458	0.586	0.327
a_2	-0.334	-0.256	-0.409	b_2	0.189	0.331	0.046
a_3	-0.009	-0.004	-0.018	b_3	0.172	0.307	0.036
				b_4	0.120	0.162	0.083

Standard Deviations							
	1965:Q2 - 1983:Q4				1984:Q1 - 2007:Q2		
Parameter	Pos. Mean	95 percent	5 percent		Pos. Mean	95 percent	5 percent
σ_v	1.369	1.575	1.195		0.766	0.873	0.674
σ_g	1.024	1.172	0.896		0.487	0.549	0.430
σ_e	1.107	1.293	0.934		0.308	0.367	0.255

	Covariance				Weight on interest rate smoothing		
Parameter	Pos. Mean	95 percent	5 percent	Parameter	Pos. Mean	95 percent	5 percent
l_{21}	-0.038	0.039	-0.115	ν	46.126	76.407	24.199
l_{31}	-0.003	0.055	-0.067				
l_{32}	0.007	0.073	-0.055				

	Variance of state equation		
Parameter	Pos. Mean	95 percent	5 percent
Q	1.234	1.981	0.633

Table 1: Parameter estimates from benchmark model: The a_i and b_i are estimates of the constraints, σ_i and l_i the standard deviation and covariance terms of the variance matrix, ν is the weight on interest rate smoothing in the loss function and Q is the variance of shocks in the state equation governing the time-varying weight.

IS Curve					
Parameter	Benchmark	RS	Dennis	RS Original	Dennis Original
a_0	0.035	0.116	0.086	0.000	0.035
a_1	1.269	1.243	1.187	1.160	1.596
a_2	-0.334	-0.318	-0.252	-0.250	-0.683
a_3	-0.009	-0.035	-0.030	-0.100	-0.021

Phillips Curve					
Parameter	Benchmark	RS	Dennis	RS Original	Dennis Original
b_0	-0.062	-0.056	-0.073	0.000	0.025
b_1	0.458	0.514	0.501	0.700	0.401
b_2	0.189	0.140	0.141	-0.100	0.080
b_3	0.172	0.168	0.171	0.280	0.407
b_4	0.120	0.109	0.114	0.140	0.144

Table 2: Parameter estimate comparison of benchmark model with Rudebusch and Svensson (1998) (RS) and Dennis (2006) (Dennis) using the same sample. The columns “RS Original” and “Dennis Original” show the parameter estimates from the respective published papers, see the main text for more details.

Standard Deviations			
	Inflation	Output Gap	Fed Funds Rate
Data			
pre 1984:Q1	2.284	3.298	3.617
post 1984:Q1	0.964	1.807	2.410
% fall: pre to post	0.578	0.452	0.334
Model			
pre 1984:Q1	1.981	3.228	3.658
post 1984:Q1	0.744	1.727	2.412
% fall: pre to post	0.624	0.465	0.340
Counterfactual (No SV)			
pre 1984:Q1	2.288	3.299	3.725
post 1984:Q1	1.759	3.738	4.495
% fall: pre to post	0.231	-0.133	-0.207
Counterfactual (No Preference Change)			
pre 1984:Q1	2.234	3.283	2.833
post 1984:Q1	0.901	1.781	1.777
percent fall: pre to post	0.597	0.458	0.373

Table 3: Counterfactual standard deviations. The exercise with no change in variance is labeled "No SV" and the exercise with fixed preference parameter is labeled "No Preference Change". See the main text for more details.

	Bayesian Information Criterion
Baseline	1215.57
79:Q3 (i)	1308.50
87:Q3 (ii)	1245.98
79:Q3 and 87:Q3 (iii)	1255.18
Alt State Eqn	1212.00

Table 4: Bayesian Information Criterion calculations for the the baseline specification and four different specifications. The second, third and fourth rows show estimates with different dates for the break in the shock variance parameters and the last row shows the estimate for the "Alt State Eqn" specification. See section 7 for more details.

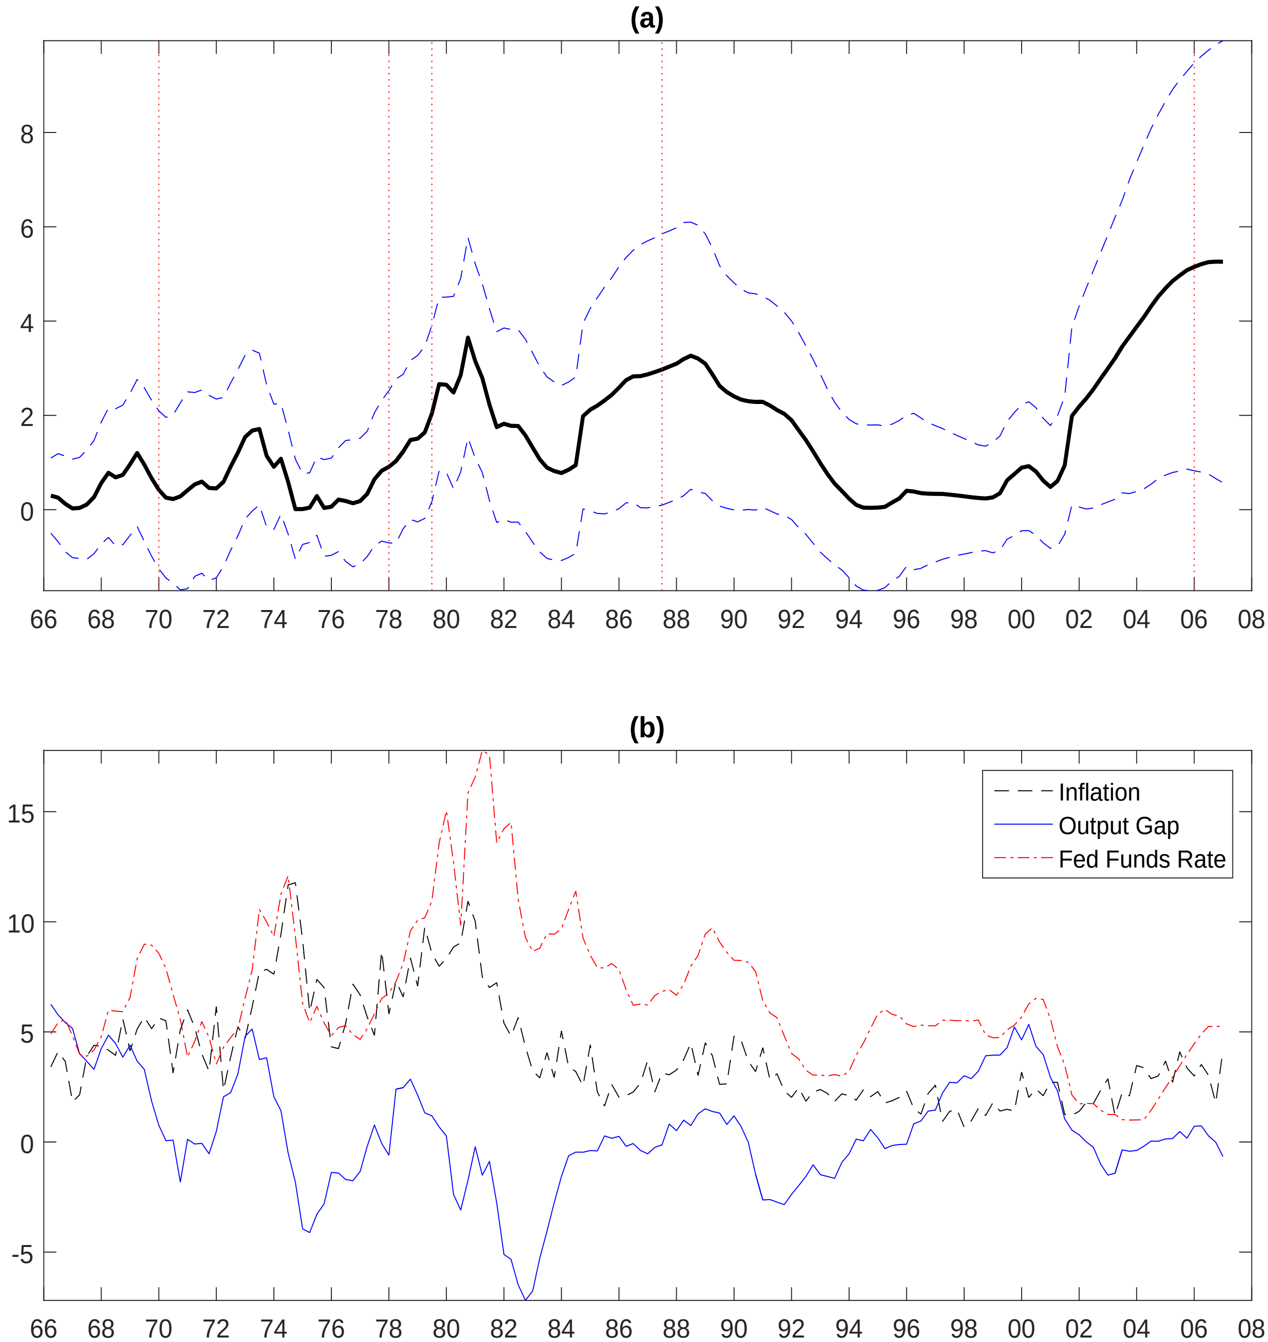

Figure 1: Panel (a) shows the time-varying weight on inflation with one standard deviation confidence intervals. The vertical red lines indicate a change in the chair of the FOMC. Panel (b) shows the data used in the estimation.

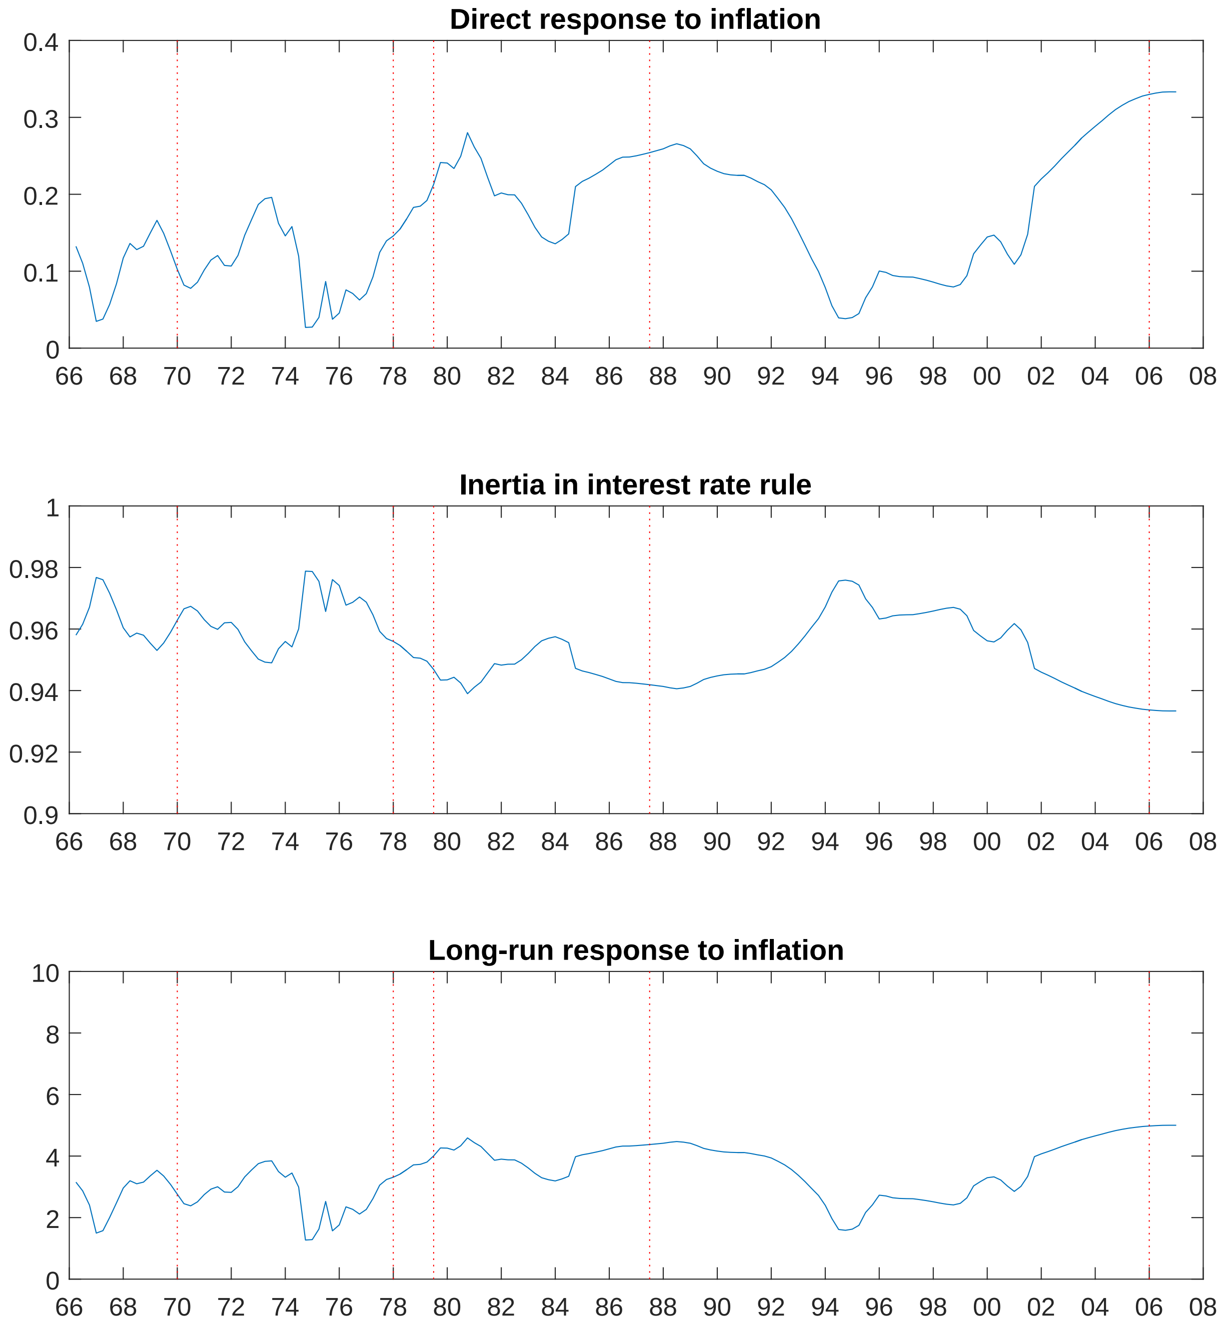

Figure 2: The three panels show properties of the optimal interest rate rule evaluated at the posterior mean of the parameter estimates. The top panel shows the direct response to inflation $(F_{1,t} + F_{2,t} + F_{3,t} + F_{4,t})$, the middle panel shows the inertial component $(F_{7,t} + F_{8,t} + F_{9,t})$ and the third panel shows the long-run response to inflation given by $\frac{F_{1,t}+F_{2,t}+F_{3,t}+F_{4,t}}{1-F_{7,t}-F_{8,t}-F_{9,t}}$

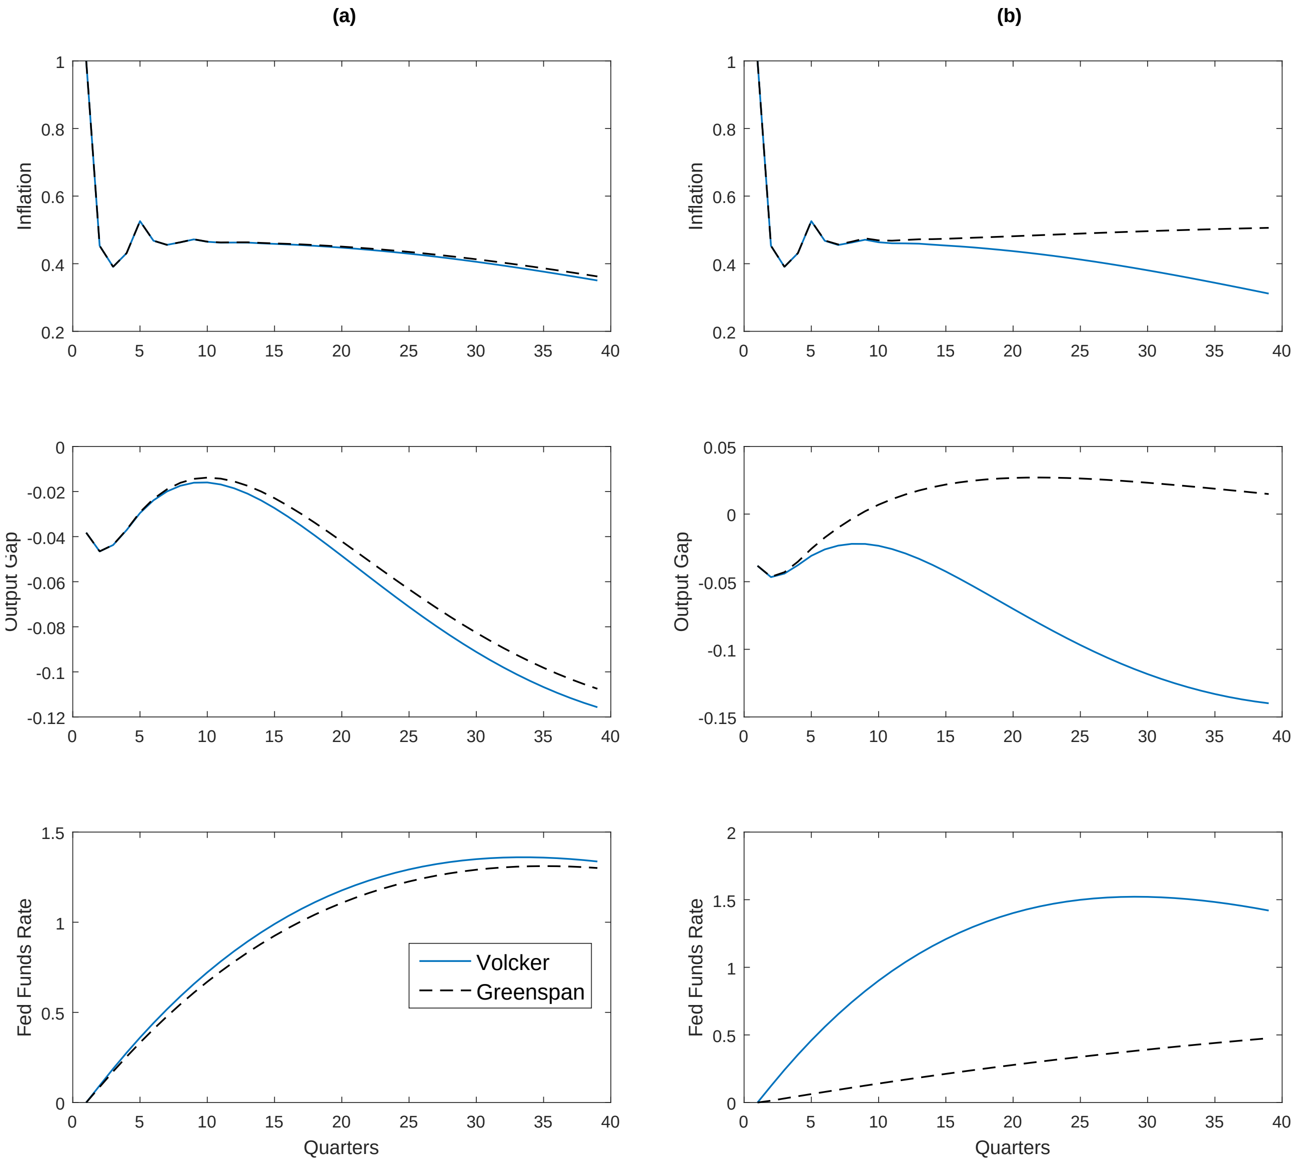

Figure 3: Impulse responses to a one unit shock in inflation. The solid blue line shows the responses using preferences estimated under Volcker's tenure, while the dashed green line shows those under Greenspan's tenure. Panel(a) uses the average of the estimated preferences for both chairmen while Panel (b) uses the highest value of preferences estimated under Volcker's tenure and lowest value of preferences estimated under Greenspan's tenure.

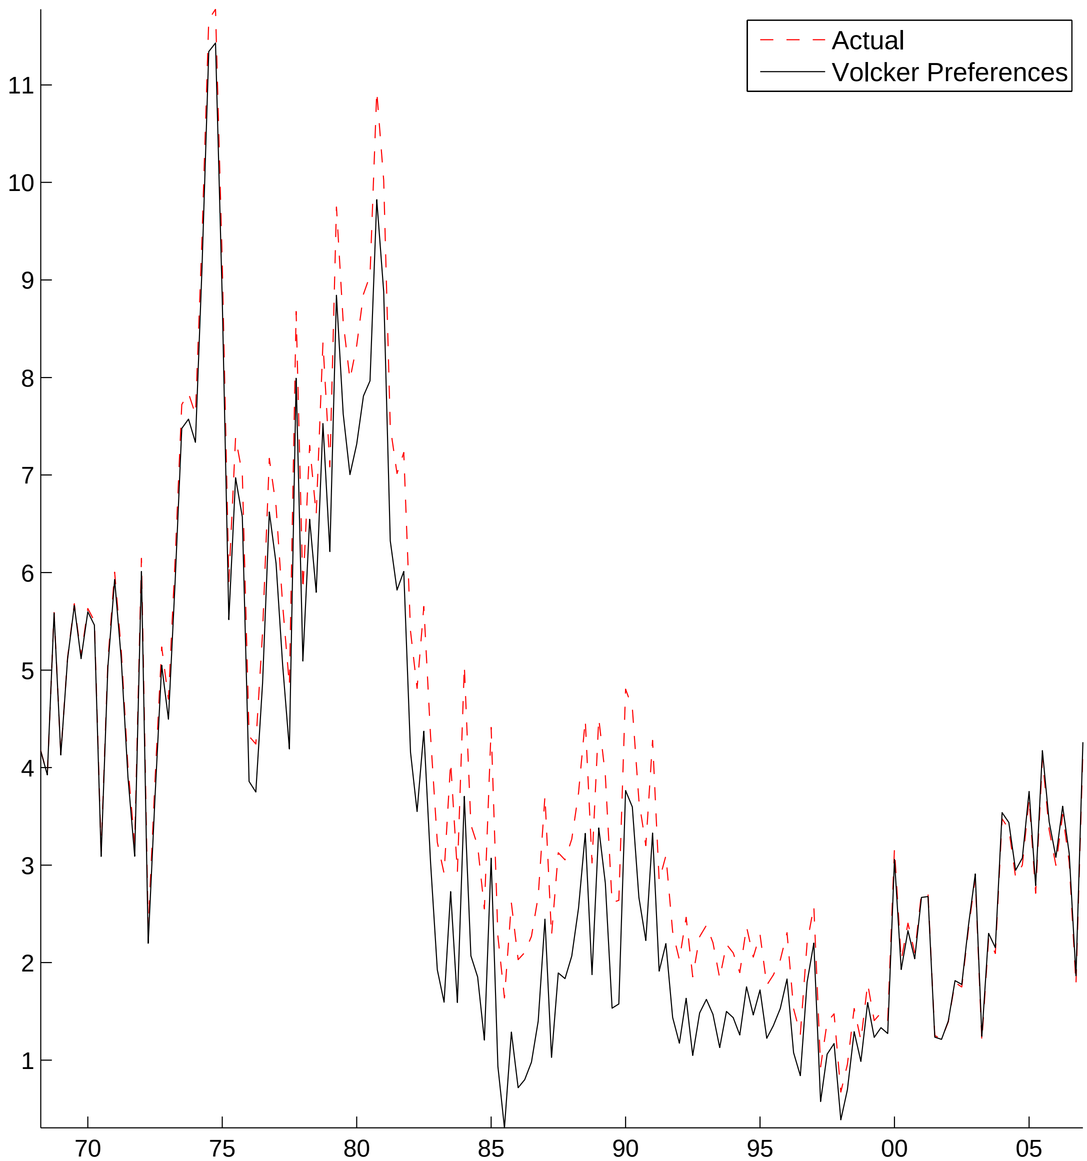

Figure 4: Counterfactual: Actual inflation (dashed) and simulated path of inflation with Volcker's preferences (solid)

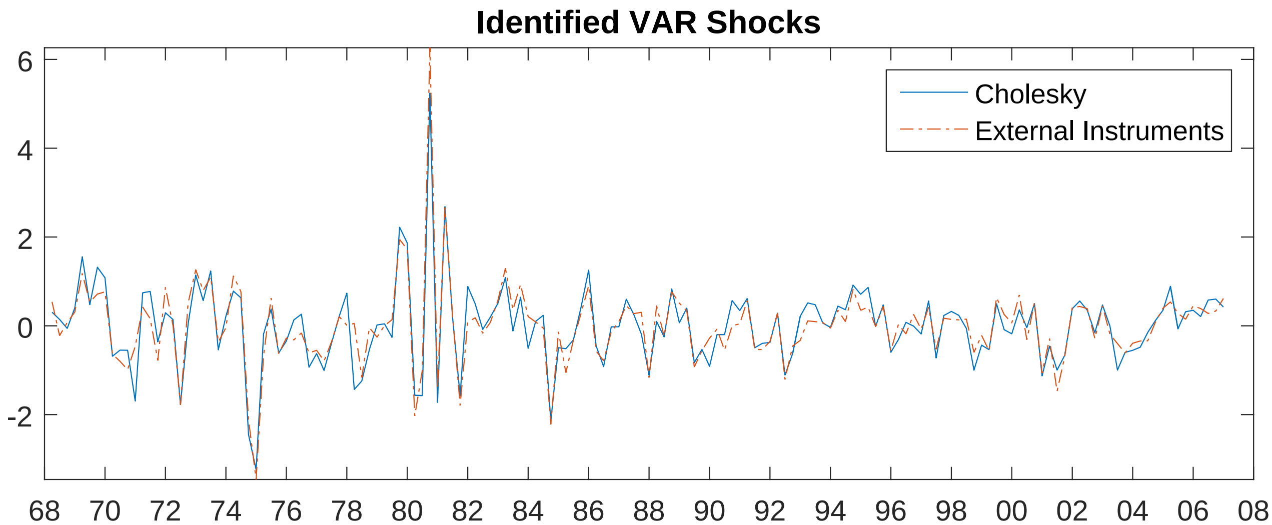

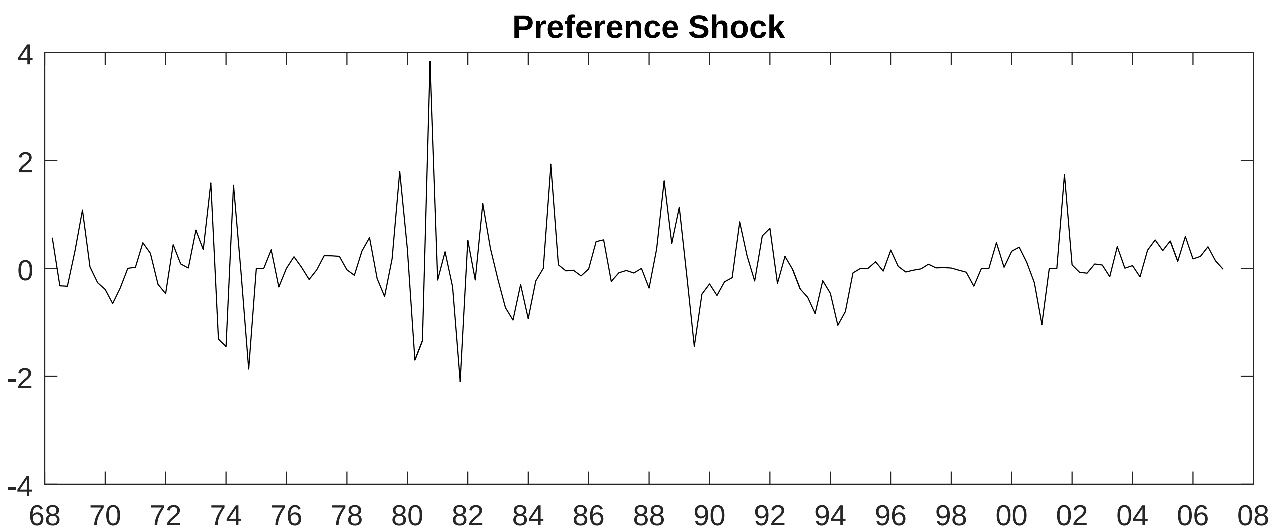

Figure 5: Monetary policy shocks: Top panel shows structural shocks from a VAR identified using the Cholesky and External Instruments methodology. The bottom panel show shocks constructed using change in the estimated preferences.

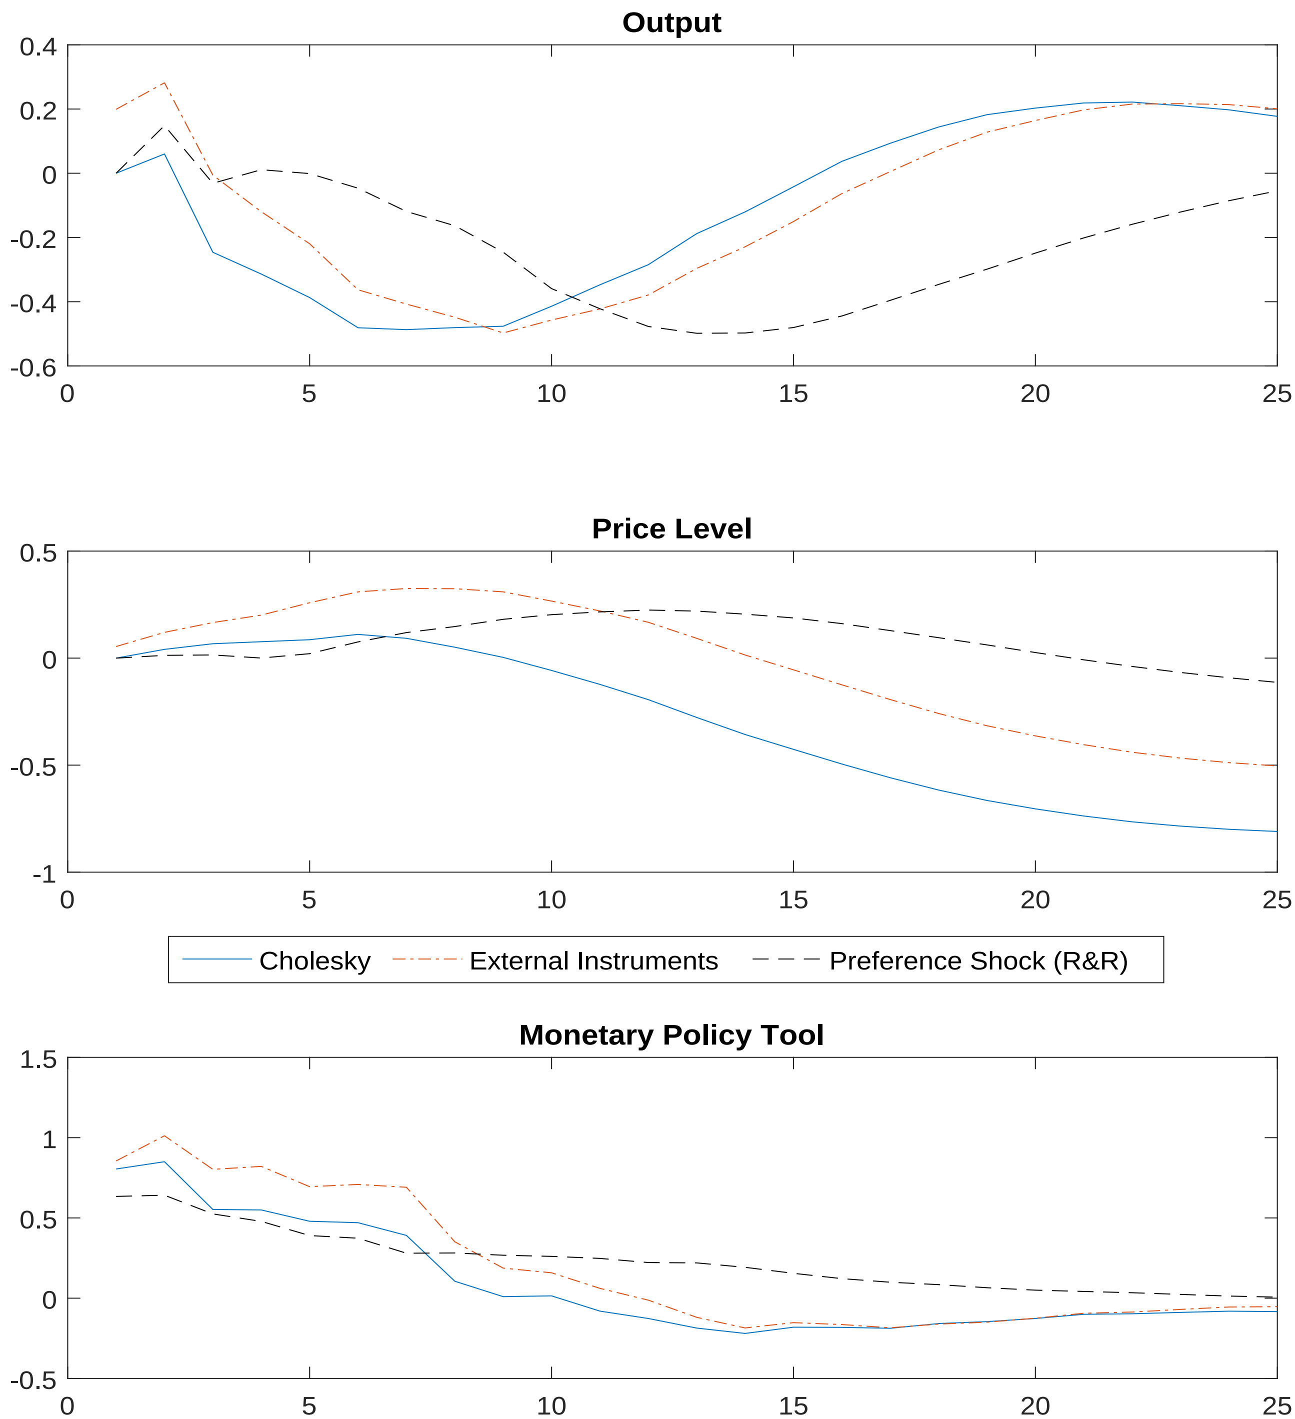

Figure 6: Impulse response to a monetary policy shock. The Cholesky and External Instruments VAR use the fed funds rate as the monetary policy indicator, while the Preference Shock (R&R) uses the cumulated preference shocks as the monetary policy indicator, see the text for more details.

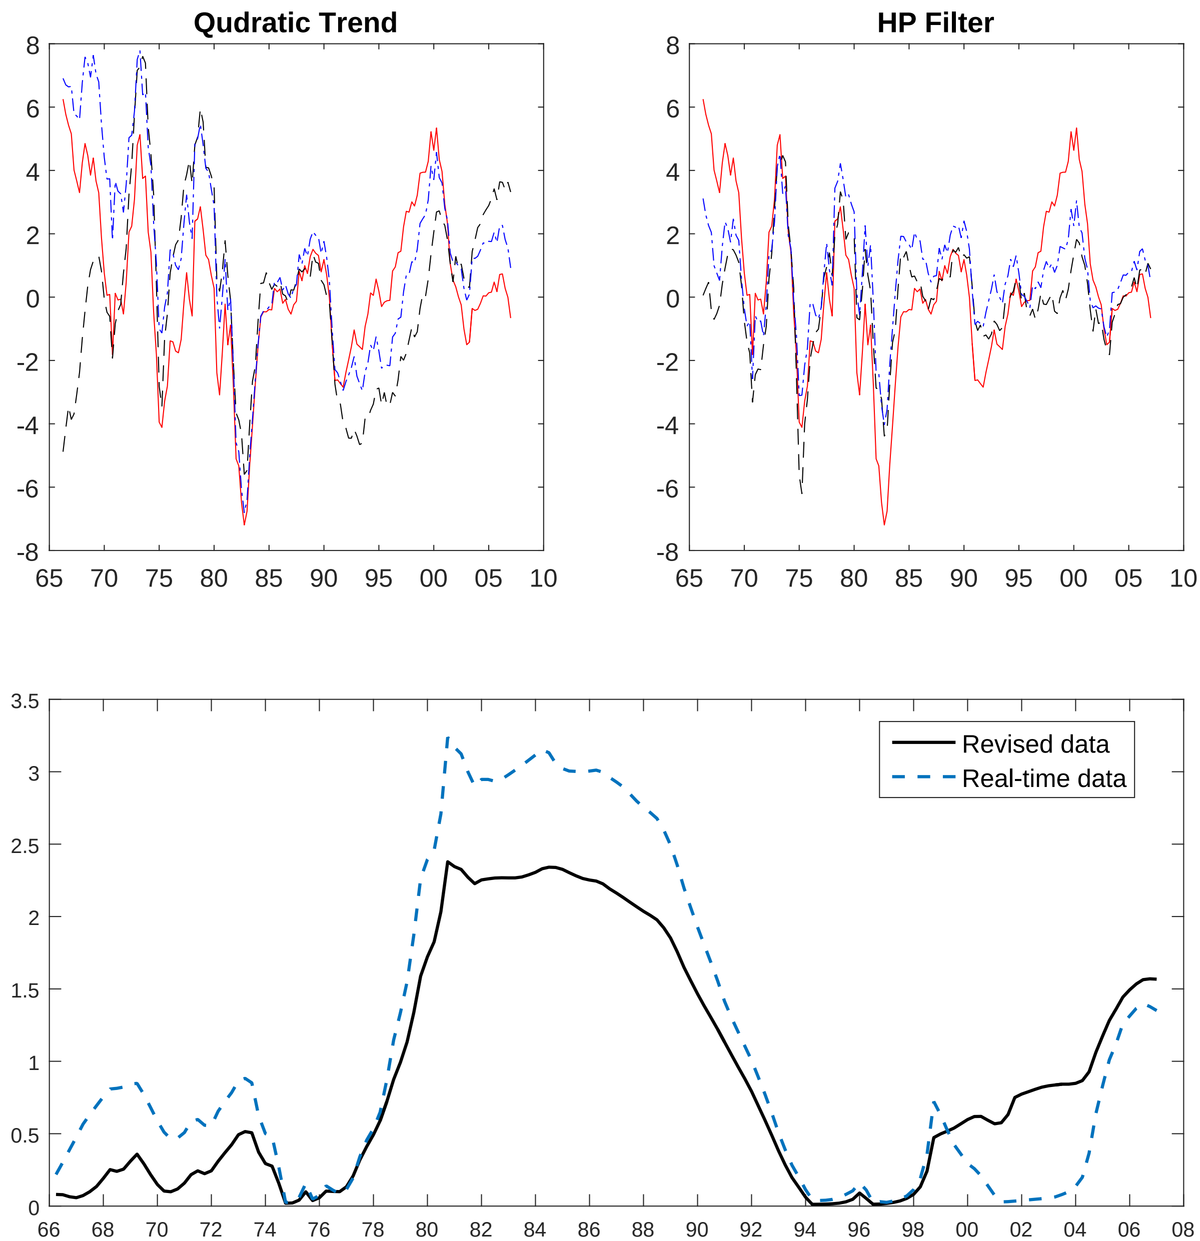

Figure 7: The top left panel shows the output gap calculated using a quadratic trend, while the top right panel uses the HP-Filter. In both these panels the dashed black lines use real-time GDP data, the dotted blue line uses revised GDP data and the red line shows the CBO output gap measure which is used in the baseline specification. The bottom panel shows the smoothed estimates of the time-varying weight on inflation using a quadratic trend to construct the output gap.

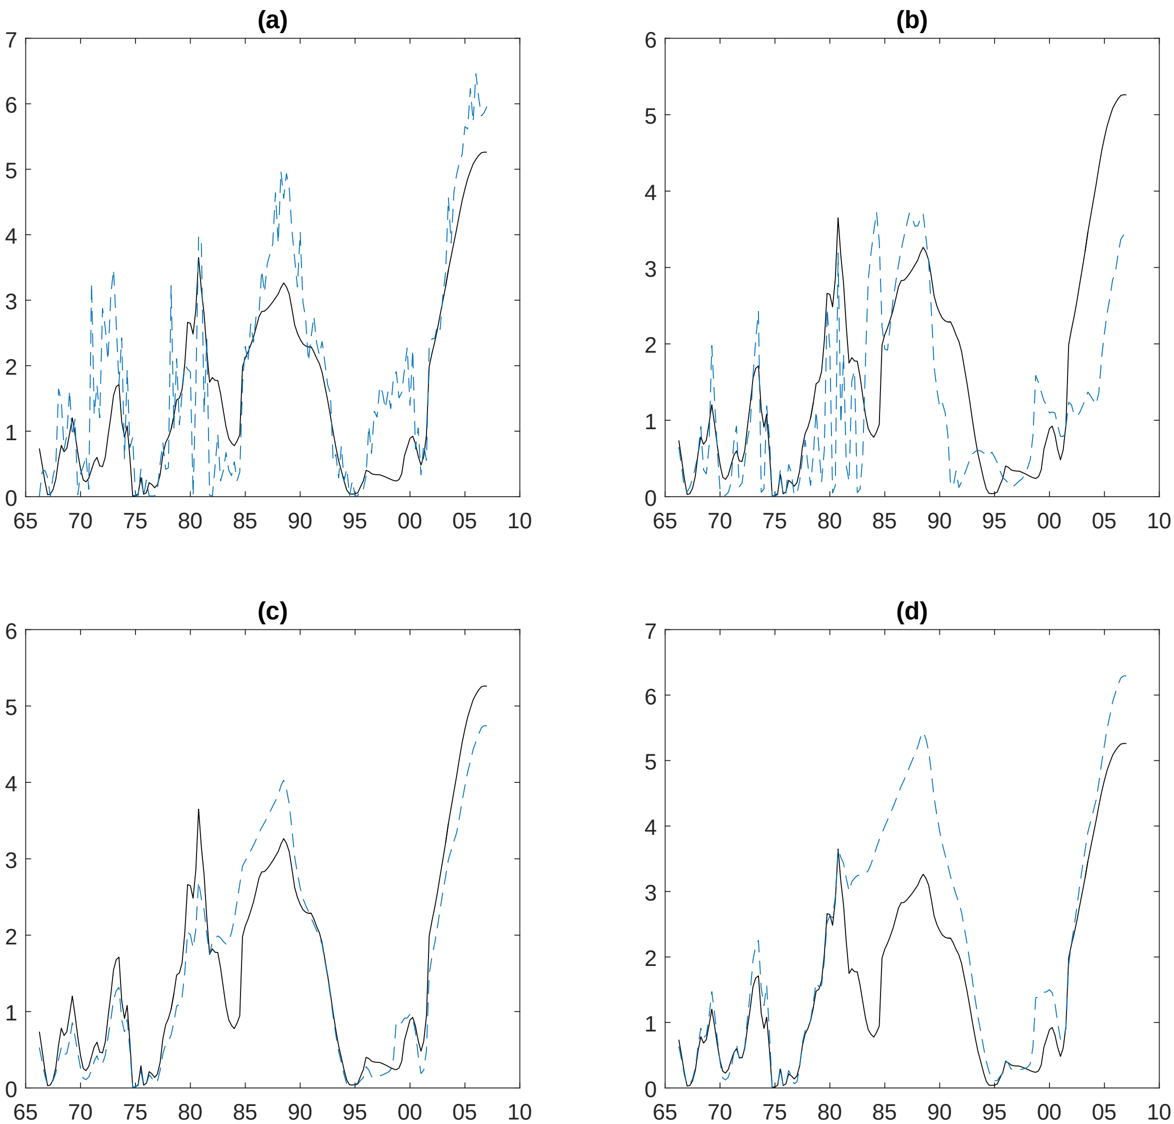

Figure 8: The solid black line in each panel plots the smoothed estimate of the time varying weight on inflation from the baseline specification, while the dashed blue line plots the weight for four different specifications. Panels (a) show the estimates from the "Alt State Eqn" specification. Panels (b)-(d) show the estimates with different dates for the break in the shock variance parameters: c) 1979:Q3, c)1987:Q3 and d)1979:Q3 and 1987:Q3. See section 7 for more details.