BOCCONI UNIVERSITY - MILAN
Ph.D. in Economics

ON THE POLICY PREFERENCES OF THE US FEDERAL RESERVE

Paolo SURICO

Advisor: Professor Carlo FAVERO

Academic year 2003-2004

On the Policy Preferences of the US Federal Reserve

Paolo Surico
Bocconi University

Academic year: 2003-2004

Contents

Acknowledgment

> "Salomon saith: There is no new thing
> upon the earth. So that as Plato had an
> imagination, that all knowledge was but
> remembrance; so Salomon giveth his
> sentence, that all novelty is but oblivion."
>
> Francis Bacon: Essays, LVIII,
> from Jorge Luis Borges: L'Aleph

Writing a Ph.D. thesis is a unique experience from a scientific perspective as well as from a personal perspective. While I devote the entire book to motivating the former, I wish to spend a few words here to the personal aspects as in my view these actually determine the degree of success of such an important chapter of the profession. Needless to say, the success we are talking about is very much subjective and it has the non-trivial privilege of being free from any external approval such as publications and number of citations. As it is too frequently misunderstood or forgotten, it is important to stress that the latter are very much an instrument and anyone serious should resist the temptation to confuse them with the more noble objective of knowledge diffusion.

During the last four years, my research interests have brought me through several cities including Glasgow, Barcelona, Frankfurt and Milan. Each of them is associated to a distinguishable feeling that a few people made special. Vasso, Efrem, Giordano, Saverio, Gaia, Nico and Joana have greatly contributed to my serenity such as these lines, on purpose, can simply be obscuring. By no means, this reflects the influential role they played and they continue to play for my life. I feel Lucky for this!

On the academic side, Carlo Favero had the burden of taking care of my research. Without his extremely warm support and promptly encouraging advice, the ideas presented here would have difficulty turned into a Ph.D. thesis. Jordi Galí and Anton Muscatelli were much more than members of the thesis committee. I wish to sincerely thank them for having forced my production to be at the highest possible level given the endowment constraint. The invaluable guidance of Guido Tabellini has very much shaped the spirit to which I wish my future research be committed.

Besides them, I had the honor to know a number of interesting people. Pierpaolo Battigalli, Tommaso Monacelli, Pietro Garibaldi, Eliana La Ferrara, Massimo Rostagno, Bruno Sitzia and Ulf Söderström provided me with important feedback and timely encouragement. Efrem Castelnuovo and Saverio Simonelli proved to be great in research as much as in spirit. Their friendship has tenderly followed all important moments of my Ph.D. Co-authoring and confronting with Efrem has been not only a very fruitful experience but also a stimulating avenue for future research together.

Being among the first Ph.D. at Bocconi University, I feel very much 'responsible' and pleased. I am grateful to all fellows in the program for the friendly and stimulating environment they provide me with. A special thanks to Nicola Curci, who was unlucky enough to be a very good friend and not simply a colleague. Giuseppe Attanasi provided superb musical assistance, Angela Baldassarre excellent secretarial assistance. Gavino and Silvio gave me intellectually stimulating discussions, and not just about economics.

As a matter of life, any experience could have been different if at some cross we picked up an alternative path. These years are not an exception. Being content of your present does not exclude a critical inspection of the past and demanding problems inevitably require painful solutions. Evaluating the latter is beyond me. Accepting their implications has been an important step towards my maturity. Along this way, I have learned by now to recognize an irremovable pattern in a constantly changing background. This is a bunch of people including Annalisa, Antonio, Claudia, Massimo, Valeria and Vito. I am very proud they are my closest friends.

Dulcis in fundo my family, who has taught me the most difficult subject of education: life. This thesis is dedicated to my brother Dario. I owe to him most of my strength. I devote to him most of my thoughts.

Milan, April 2004

Paolo Surico

Preface

Welcome to my Ph.D. thesis. Handle it carefully. I have put in it a lot of efforts!

It consists of three chapters. Hope you enjoy them.

The first chapter - *Inflation Targeting and Nonlinear Policy Rules: the Case of Asymmetric Preferences* - derives and evaluates the testable predictions of a model of sticky prices and monopolistic competition in which the central bank is allowed to weight differently positive and negative deviations of inflation and output from the target values. Asymmetric preferences translate into a nonlinear reaction function and they are shown to induce a systematic inflation bias through the private sector expectations of a larger policy response in recessions than in booms. Reduced-form and structural estimates of the central bank first order condition indicate that the preferences of the Fed have been highly asymmetric only before 1979, with the response to output contractions being larger than the response to output expansions of the same magnitude. This asymmetry is shown to induce an average inflation bias of 1.11% that appears to have substantially contributed to the great inflation of the 1960s and 1970s.

The second chapter - *Measuring the Time-Inconsistency of US Monetary Policy* - identifies the novel inflation bias associated with the presence of asymmetric preferences using a Lucas aggregate supply curve as the model of the economy. The inflation bias occurs in spite of the fact that the monetary authorities target output at its natural level and therefore it represents an alternative form of time-inconsistency relative to the Barro-Gordon model. Using the identification strategy developed in Chapter 1, this paper quantifies the relative contribution of a change in the inflation target and a shift in the asymmetric preferences to the rise and fall of postwar US inflation. The average inflation bias, whose estimate is around one percent before 1979, tends to disappear over the last two decades whereas the inflation target declines from 3.42% to 1.96%. This result can be rationalized in terms of the preference on output stabilization which is found to be large and asymmetric in the pre- but not in the post-Volcker regime.

The third chapter - *Model Uncertainty, Optimal Monetary Policy and the Preferences of the Fed* (coauthored with Efrem Castelnuovo) - investigates to what end the policy makers' uncertainty about the structure of the economy can account for the observation that the Fed smooths interest rate changes. Using a novel 'thick modeling' approach to retain and pool the information embodied in a large class of given models, this paper incorporates model uncertainty into the identification of central bank's preferences. The robust policy rule implied by the thick modeling shows the kind of smoothness observed in the data without resorting to a theoretically implausible weight for interest rate changes in the central bank objective function.

Inflation Targeting and Nonlinear Policy Rules: the Case of Asymmetric Preferences

February 2004

Abstract

This paper investigates the empirical relevance of a new framework for monetary policy analysis in which the decision makers are allowed to weight differently positive and negative deviations of inflation and output from the target values. Reduced-form and structural estimates of the central bank first order condition indicate that the preferences of the Fed have been highly asymmetric only before 1979, with the response to output contractions being larger than the response to output expansions of the same magnitude. This asymmetry is shown to induce an average inflation bias of 1.11% that appears to have substantially contributed to the great inflation of the 1960s and 1970s.

JEL Classification: E52, E58.

Keywords: asymmetric objective, nonlinear monetary policy rules, average inflation bias

---

[0] I am grateful to Efrem Castelnuovo, Carlo Favero, Jordi Galí, Tommaso Monacelli, Anton Muscatelli, Roberto Perotti, Giorgio Primiceri, Ulf Söderström, Guido Tabellini and participants at various seminars and conferences for very helpful feedbacks.

1 Inflation targeting and nonlinear policy rules: the case of asymmetric preferences

A popular method of monetary model building is to regard policy interventions as the solution of an optimal control problem in which the central bank minimizes some quadratic criterion subject to a linear structure of the economy. The quadratic characteristic of the objective and the linear feature of the constraints give rise to a linear first order condition, usually referred to as a targeting rule (see Svensson, 1999), that describes the optimal response of the central bank to the developments in the economy. While the quadratic specification implies that monetary authorities evenly weight positive and negative deviations of inflation and output from the target values, such a modeling choice has been questioned by several practitioners at the policy committees of various central banks on the ground that it has little justification beyond analytical tractability.[1]

Blinder (1997, p. 6) argues that '*academic macroeconomists tend to use quadratic loss functions for reason of mathematical convenience, without thinking much about their substantive implications. The assumption is not innocuous, [...] practical central bankers and academics would benefit from more serious thinking about the functional form of the loss function*'. Describing his experience as Fed vice-Chairman Blinder (1998, pp. 19-20) pushes the argument even further and claims '*in most situations the central bank will take far more political heat when it tightens pre-emptively to avoid higher inflation than when it eases pre-emptively to avoid higher unemployment*', suggesting that political pressures can induce asymmetric central bank interventions. Similar concerns appear to emerge also at other central banks like the ECB and in the occasion of an interest rate cut of 50 point basis Duisenberg (2001) states '*the main-*

---

[1] The few notable exceptions include Rotemberg and Woodford (1999) and Woodford (2003, ch. 6), who show that the quadratic form can be obtained as a second order approximation of the representative agent's utility function.

tenance of price stability remains our first priority. [...] today's action could be taken "without prejudice to price stability", and it thereby supported the other goals of EMU, such as economic growth'.

On the academic side, several recent studies explore novel mechanisms through which the costs of the business cycle can be asymmetric. Persson and Tabellini (1999) combine retrospective voting with imperfect information about the incumbent's talent to show that career concerned politicians can make reappointment more likely by endowing the central bank with an asymmetric objective that requires a larger monetary policy response in periods of poor economic performance.[2] Galí, Gertler and Lopez-Salido (2003a) construct a theoretical measure of welfare gap that is based on price and wage markups, and find that the costs of output fluctuations for the US have been historically large and asymmetric. Erosa and Ventura (2002) introduce transaction costs and heterogeneity in portfolio holdings in an otherwise neo-classical model and show that these frictions can make the costs of inflation variation asymmetric. Lastly, the psychology of choice reveals that people tend to place a greater weight on the prospect of losses than on the prospect of gains in decision making under uncertainty (see Kahneman and Tversky, 1979), suggesting that also policy makers, who aggregate over individual welfare, may be loss-averse.

Despite its intuitive appeal, only a few studies have attempted to identify asymmetric central bank behaviors and the relevance of this alternative framework remains to be assessed. Cukierman and Muscatelli (2003), Martin and Milas (2004), and Kim, Osborne and Sensier (2004) show some international evidence that supports the notion of nonlinear interest rate rules. Ruge-Murcia (2003 and 2004), and Cukierman and Gerlach (2003) adopt an inflation rate reaction function that is nonlinear in either inflation

---

[2] De Long (1997) forcefully argues that US monetary policy during the 1970s was highly sensitive to the political pressures for a higher money growth and lower interest rates, and provides extensive narrative evidence about the influence of Nixon's administration on the Chairmanship of Arthur Burns at the Fed.

or the output gap, and using data for some OECD economy they favor the hypothesis of an asymmetric objective. Dolado, Maria-Dolores and Ruge-Murcia (2004) estimate an optimal interest rate rule that is drawn upon the existence of asymmetric preferences on inflation only, and find that US monetary policy can be characterized by a nonlinear function after 1983.

This paper contributes to the literature on monetary policy rules in several respects. First, it proposes a general, potentially asymmetric specification for *both* the inflation and the output objectives that nests the quadratic form as a special case. Accordingly, the optimal policy rule is nonlinear if and only if the preferences of the central bank are asymmetric. Second, the analytical solution of the optimal control problem allows us to identify the degree of nonlinearity and asymmetry with respect to *both* objectives, a result that to our knowledge of the existing literature comes as new. Third, the model generates the testable prediction that the monetary authorities respond not only to the level of inflation and output gaps as suggested by Taylor (1993) but also to their squared values. Fourth, reduced-form and structural estimates of US monetary policy rules indicate that nonlinearity is a robust feature of the postwar data only before 1979 and with respect to the output gap. While this finding is consistent with the notion of a Fed's policy regime shift, it provides an explanation for the great inflation of the 1960s and 1970s as the model predicts that asymmetric preferences over the output gap generate an average inflation bias. The latter is found to move from 1.11% before 1979 to a value not statistically different from zero over the last two decades.

The road map of the paper is as follows. Section 2 presents the model and derives the interest rate rule as the first order condition of the central bank optimization problem. Section 3 reports the estimates of both the policy rule coefficients and the preference parameters, and conducts a robustness analysis. The following section shows that asymmetric preferences on the output gap induce an average inflation bias, and proposes a simple strategy to decompose the actual inflation mean into a target and a bias argument. Section 5 concludes.

1.1 Theoretical model

We assume that the central bank conducts monetary policy through a *targeting rule* according to the terminology of Svensson (1999). Thus, all available information are used to bring at each point in time the target variables in line with their targets by penalizing any future deviation of the former from the latter. The policy rule is modeled as the discretionary outcome of an intertemporal optimization problem in which the decision makers minimize a given criterion subject to the constraints provided by the structure of the economy. The optimizing device allows us to back out the objectives of the monetary authorities, which are unobserved, from the observed path of policy rates implying that evidence on the latter can be interpreted as informative about the former. Since our identification strategy relies on the estimation of a model-based specification for the reaction function, we challenge the assumption of symmetric policy preferences in the context of a popular framework for monetary policy analysis. This is a version of the New-Keynesian model of the business cycle derived in Yun (1996), and Woodford (2003, chs. 3 and 4), among many others.[3]

1.1.1 The structure of the economy

This subsection describes an aggregate, log-linearized version of the New-Keynesian forward-looking model with sticky prices that has been recently summarized by Clarida, Galí and Gertler (1999). The evolution of the economy is compactly represented by the

---

[3] Surico (2003) shows that both the theoretical and the empirical results obtained here using a New-Kynesian model are robust to the specification of a Lucas aggregate supply curve as structure of the economy.

following two-equation system:

$$\pi_t = \theta E_t \pi_{t+1} + k y_t + \varepsilon_t^s \tag{1}$$

$$y_t = E_t y_{t+1} - \varphi (i_t - E_t \pi_{t+1}) + \varepsilon_t^d \tag{2}$$

Equation (1) captures the staggered feature of a Calvo-type world in which each firm adjusts its price with a constant probability in any given period, and independently from the time elapsed from the last adjustment. The discrete nature of price setting creates an incentive to adjust prices by more the higher is the future inflation expected at time t. The inflation level is π_t whereas the output gap is denoted by y_t and captures the movements in marginal costs associated with variations in excess demand. For analytical convenience, the aggregate supply curve is assumed purely forward-looking. Galí and Gertler (1999), Ireland (2001), Galí, Gertler and Lopez-Salido (2003b), and Smets and Wouters (2003a) provide empirical support for this choice as a good first approximation to the dynamics of US inflation.

Equation (2) is a standard Euler equation for consumption combined with the relevant market clearing condition. It basically brings the notion of consumption smoothing into an aggregate demand formulation by making the output gap a positive function of its future value and a negative function of the real interest rate, $i_t - E_t \pi_{t+1}$. Lastly, ε_t^s and ε_t^d are respectively cost and demand disturbances that obey an autoregressive, mean reverting process.

1.1.2 An asymmetric specification of the loss function

An important aspect of monetary policy making is that policy actions are taken before the realization of economic shocks and therefore before the variables in the system are determined. Accordingly, the problem of the central bank is to choose the interest rate at the beginning of period t conditional upon the information available at the end of the previous period. This timing device is captured by the following intertemporal criterion:

$$\underset{\{i_t\}}{Min} \; E_{t-1} \sum_{\tau=0}^{\infty} \delta^{\tau} L_{t+\tau} \tag{3}$$

where δ is the discount factor and L stands for the period loss function.

Our framework differs from the conventional quadratic set up in that we employ a more general specification of the monetary authorities' objectives. Indeed, the quadratic form may approximate reasonably well a number of different functions and in the absence of a rigorous theoretical foundation any specific nonquadratic proposal is destined to be unsatisfactory against the wide range of plausible alternatives. Hence, rather than attempting to uncover the correct functional form of policy makers' preferences, we evaluate the symmetric quadratic setup upon the empirical merits of the monetary policy rule that this specification implies. With this descriptive scope in mind, we write L_t as follows:

$$L_t = \frac{1}{2}\left[(\pi_t - \pi^*)^2 + \frac{\alpha}{3}(\pi_t - \pi^*)^3\right] + \frac{\lambda}{2}\left[y_t^2 + \frac{\gamma}{3}y_t^3\right] + \frac{\mu}{2}(i_t - i^*)^2 \tag{4}$$

The coefficients λ and μ represent the central bank's aversion towards output fluctuations around potential and towards interest rate *level* fluctuations around the target i^*. The policy preference towards inflation stabilization is normalized to one and therefore λ and μ are expressed in relative terms. The inflation target is π^* whereas the parameters α and γ capture any asymmetry in the objective function of the monetary authorities.

The cubic specification (4) departures from the quadratic in that policy makers are allowed, but not required, to treat differently positive and negative deviations of inflation and output from the target. A negative value of γ implies that, everything equals, an output contraction relative to the potential level is weighted more severely than an output expansion. To see this notice that whenever $y_t < 0$ the cubic term, γy_t^3, is positive and amplifies the penalty due to the quadratic component. Conversely, for values of output above potential the quadratic and the cubic terms move in opposite directions implying that a positive deviation of a given amount is associated with a smaller loss than a negative deviation of the same size. Figure 1 compares the standard quadratic with the asymmetric cubic function using the historical values of the output gap and the estimates of γ reported below.

A similar reasoning holds for the coefficient α that captures any asymmetry in the policy preferences for stabilizing inflation around the target. However, if the monetary authorities are more concerned about overshooting π^* rather than undershooting it, the value of α would be positive meaning that high inflation relative to the target is more costly than low inflation. It should be noted that while these sign predictions seem plausible given the sample we use, the cubic specification does not prevent α to be negative corresponding to a case in which the risk of deflation outweighs the risk of inflation.[4]

The cubic loss function nests the quadratic form as a special case such that $\alpha = \gamma = 0$ corresponds to the symmetric parametrization $L_t = \frac{1}{2}\left[(\pi_t - \pi^*)^2 + \lambda y_t^2 + \mu (i_t - i^*)^2\right]$. The latter can be obtained as a second order approximation of the utility-based welfare function in a New-Keynesian model of the business cycle that involves a zero lower bound for nominal interest rate (see Woodford, 2003, ch. 6). Accordingly, the policy preferences would be functions of some primitive parameters of the model implying that potential evidence of asymmetries in the central bank objective could be tracked into evidence of asymmetries in the representative agent's utility. Indeed, as argued by Clarida, Galí and Gertler (1999), the representative agent approach can be misleading

---

[4] The cubic specification can also be interpreted as some third-order approximation around $(\pi_t - \pi^*) = 0$ and $y_t = 0$ to the linex function proposed by Nobay and Peel (2003), and employed by Chadha and Schellekens (1999), Geraats (1999) and Ruge-Murcia (2003 and 2004). The advantage of using the cubic form as the primitive function is that it does not require any approximation of the optimal monetary policy rule. Nevertheless, for a realistic range of values for $(\pi_t - \pi^*)$ like $[-0.04, 0.09]$ and for y_t like $[-0.08, 0.06]$, and given the estimates of α and γ reported below, the cubic and the linex function behave very similarly.

as a guide to welfare analysis and in the absence of complete markets it is likely that some groups suffer more in recessions than others. This suggests that an asymmetric utility-based specification of the loss function may be a desirable representation of the social costs associated with the business cycle.

1.1.3 A nonlinear policy rule

We solve for the optimal monetary policy under discretion. Because no endogenous state variable enters the model, the intertemporal problem reduces to a sequence of static optimization problems. This amounts to choosing in each period the instrument i_t such as to minimize:

$$E_{t-1}\left\{\frac{1}{2}\left[(\pi_t-\pi^*)^2+\frac{\alpha}{3}(\pi_t-\pi^*)^3\right]+\frac{\lambda}{2}\left[y_t^2+\frac{\gamma}{3}y_t^3\right]+\frac{\mu}{2}(i_t-i^*)^2\right\}+F_t$$

subject to $\pi_t = ky_t + f_t$ and $y_t = -\varphi i_t + g_t$, where $F_t \equiv E_{t-1}\sum_{\tau=1}^{\infty}\delta^{\tau}L_{t+\tau}$, $f_t \equiv \theta E_t\pi_{t+1} + \varepsilon_t^s$ and $g_t \equiv E_t y_{t+1} + \varphi E_t\pi_{t+1} + \varepsilon_t^d$ are taken as given reflecting the fact that the monetary authorities cannot directly manipulate expectations. The first order condition reads

$$-k\varphi E_{t-1}(\pi_t-\pi^*)-\lambda\varphi E_{t-1}y_t-\frac{\alpha k\varphi}{2}E_{t-1}(\pi_t-\pi^*)^2-\frac{\lambda\varphi\gamma}{2}E_{t-1}y_t^2+\mu(i_t-i^*)=0 \tag{5}$$

and it implicitly describes the optimal, potentially nonlinear response of the central bank to the developments in the economy. Equation (5) nests the linear form as a special case and whenever $\alpha = \gamma = 0$ the reaction function collapses to an implicit interest rate rule of the type analyzed in Rudebusch (2002), and Clarida, Galí and Gertler (2000):

$$-k\varphi E_{t-1}(\pi_t-\pi^*)-\lambda\varphi E_{t-1}(y_t)+\mu(i_t-i^*)=0$$

This feature is attractive as it delivers a joint restriction on policy makers' preferences that can be formally tested for. The parameters α and γ are indeed crucial for the analysis of optimal monetary policy not only because they introduce an asymmetric motive in the central bank objective function but also because, more importantly, they make nonlinear an otherwise conventional policy rule. This suggests that the hypothesis of symmetric central bank preferences can be tested simply by evaluating the functional form of the interest rate reaction function as the latter would correspond to test whether the structural parameters α and γ are significantly different from zero.

1.2 Empirical results

This section reports the estimates and the relevant tests of the optimal policy rule. The analysis is conducted on US quarterly data spanning the period 1960:1-2003:2. The data set has been obtained in July 2003 from the web site of the Federal Reserve Bank of St. Louis and embodies alternative measures of inflation and output gap. In the baseline case, inflation is measured as the changes in the log of the consumer price index (CPI) while the output gap is constructed using the series of potential output provided by the Congressional Budget Office (CBO). Figure 2 plots the baseline series. As a way to provide a robustness check, we also report the results for two alternative measures of inflation and output gap, namely the GDP deflator and the Hodrick-Prescott filtered real GDP.[5]

We divide the full sample around the third quarter of 1979 which corresponds to the appointment of Paul Volcker as Fed Chairman. This lines up with a number of empirical studies that demonstrate a significant difference in the way monetary policy

---

[5] The use of a low frequency filter to obtain estimates of the target level of real activities does not contrast with the model-based definition of flexible-price level of output. As argued by Woodford (2003, ch. 7), the central bank can make society better off by accommodating technology and preference shocks while offsetting disturbances to inflation and wage mark-ups. In this vein, Smets and Wouters (2003b) show that if the monetary authorities wish to hedge against shocks of unknown nature, they would regard persistent disturbances as the only shocks affecting the target level of output. When applied to an estimated New-Keynesian model for the Euro area, they find that the counterfactual flexible-price level of output, which is the one responding to all non-monetary shocks in the economy, is indeed extremely volatile, whereas the target level of output, which is the one only affected by supply and demand disturbances, actually follows a relatively smooth path.

was conducted pre- and post-1979 (see Clarida, Galí and Gertler, 2000, and Favero and Rovelli, 2003 among many others). Moreover, we remove from the second subsample the period 1979:3-1982:3 when, as documented by Bernanke and Mihov (1998), the operating procedure of the Fed temporarily switched from federal funds rate to non-borrowed reserves targeting. Finally, we address the issue of subsample stability by re-evaluating the model over the Chairmanship of Alan Greenspan, namely 1987:3-2003:2.

We estimate a version of the central bank Euler equation using the Generalized Method of Moments (GMM) with an optimal weighting matrix that accounts for possible heteroskedasticity and serial correlation in the error terms (see Hansen, 1982). In practice, we employ a four lag Newey-West estimate of the covariance matrix. Starting from date $t-1$, four lags of the explanatory variables, the federal funds rates and the measure of inflation left out from the regression are included as instruments corresponding to a set of 19 overidentifying restrictions that can be tested for.

1.2.1 Preliminary Analysis

The quadratic terms in (5) stem from asymmetric central bank preferences but we cannot exclude in principle that some alternative source like a nonlinear Phillips curve might also return evidence of nonlinearity in the policy rule (see Schaling, 1999). A simple way to discriminate between nonquadratic objectives and nonlinear constraints is to perform the REgression Specification Error Test (RESET), which is designed to detect incorrect functional forms, on the New-Keynesian Phillips curve. Accordingly, we estimate equation (1) over the full sample using Instrumental Variables and a twelve-lag Newey-West variance covariance matrix. The set of instruments dated at time $t-1$ includes four lags of the GDP deflator inflation, the CBO output gap, the long-short interest rate spread, and the CPI inflation. When the squared, and then the squared and the cubes of the predictions $\hat{\pi}_t$ are added to the original equation, the corresponding

F-tests show that the null hypothesis of *non-misspecification* is not rejected. This suggests that the US aggregate supply curve is well approximated by a linear relation, consistently with the findings in Dolado, Maria-Dolores and Ruge-Murcia (2004), and Dolado, Maria-Dolores and Naveira (2004).

An additional form of nonlinearity comes from the policy makers' (mis)perception of the state of the economy. Suppose that on the basis of the estimates available in real-time the Fed believed for part of the sample that the output gap was larger than the revised data indicates. Then, the policy interventions during that period may appear surprisingly activist given the values of the gap from the 2003 vintage. However, using real-time data Orphanides (2004) finds that the Fed response to the output gap was actually more activist in the 1970s when the misperceptions on potential output turned out to be more severe. Moreover, Kuha and Temple (2003) show that measurement error in quadratic regressions tends to hide the presence of nonlinearities. In the view of these arguments, this paper takes an essential step towards asymmetric preferences by extending the available evidence on monetary policy rules using revised data.

A further reason for nonlinearity is associated with the point estimates of the natural rate of real activity. Meyer, Swanson and Wieland (2001) show that in periods of heightened uncertainty about the NAIRU, the central bank may face an incentive to move policy rates only for sufficiently large deviations of unemployment from the target. While potentially relevant, this hypothesis testing would require a real-time series for potential output such as to reflect the policy makers' beliefs about the state of the economy at the time decisions were taken. For reasons discussed above, however, we use the official estimates of potential output, which are actually revised by the CBO on a regular basis. As these revisions sensibly reduce the uncertainty about the historical measures of the output gap, this form of nonlinearity is likely to play only a marginal role in our analysis.

1.2.2 Reduced-form estimates

We solve equation (5) for i_t and prior to GMM estimation we replace expectations with realized values. As customary in the empirical studies, we introduce a lagged dependent variable to capture interest rate smoothing for which a number of explanations are provided in the literature (see Woodford, 1999, Sack and Wieland, 2000, and Castelnuovo, 2003). Accordingly, we estimate the following policy rule:

$$i_t = (1-\rho)\left[i^* + c_1(\pi_t - \pi^*) + c_2 y_t + c_3(\pi_t - \pi^*)^2 + c_4 (y_t)^2\right] + \rho i_{t-1} + v_t \qquad (6)$$

where the coefficients are given by the expressions

$$c_1 \equiv \frac{k\varphi}{\mu}, \quad c_2 \equiv \frac{\lambda\varphi}{\mu}, \quad c_3 \equiv \frac{\alpha k \varphi}{2\mu}, \quad c_4 \equiv \frac{\lambda\varphi\gamma}{2\mu}$$

and the error term is defined as

$$v_t \equiv -(1-\rho)\left\{ \begin{array}{c} c_1(\pi_t - E_{t-1}\pi_t) + c_2(y_t - E_{t-1}y_t) + \\ +c_3\left[\pi_t^2 - E_{t-1}(\pi_t)^2\right] + c_4\left[y_t^2 - E_{t-1}(y_t)^2\right] \end{array} \right\}$$

The term in curly brackets is a linear combination of forecast errors and therefore v_t is orthogonal to any variable in the information set available at time $t-1$.

Equation (6) makes clear that the reaction function parameters can only be interpreted as convolutions of the coefficients representing policy makers' preferences and those describing the structure of the economy. Although it is not possible to recover all structural parameters from a reduced-form single equation, the estimates of the policy rule can identify the asymmetric preferences as $\alpha = 2c_3/c_1$ and $\gamma = 2c_4/c_2$. In particular, the feedback coefficients c_3 and c_4 embody the relevant information such that the joint restriction $c_3 = c_4 = 0$ with $c_1 \neq 0$ and $c_2 \neq 0$ implies $\alpha = \gamma = 0$. Hence, testing the hypothesis $H_0' : c_3 = c_4 = 0$ in (6) is equivalent to testing the hypothesis $H_0 : \alpha = \gamma = 0$ in (5). Under the null of a linear reaction function, which fully corresponds to the null of symmetric preferences, the statistics has an asymptotic χ^2 distribution with as many degrees of freedom as the number of restrictions, and it can be successfully evaluated through a standard Wald test. As we are considering the auxiliary null $H_0' : c_3 = c_4 = 0$ rather than the original hypothesis $H_0 : \alpha = \gamma = 0$, the statistics is usually referred to as Wald-*type*.

In the absence of further assumptions our method only identifies the structural parameter on output gap asymmetry, γ, but neither the one on inflation, α, nor the target π^*, separately. As the focus of our analysis is on asymmetric preferences, we choose to fix a value for π^*. Specifically, we conduct a grid search in the 1% neighborhood of the subsample inflation mean, which is 4.5% for the pre- and 2.8% for the post-1979 period respectively, and we select the value that provides the best fit. Moreover, as restricting i^* appears beneficial for the convergence of the optimization algorithm, we assume that the subsample average of the interest rate provides a reasonable approximation for the target.

Table 1 reports the GMM estimates of the interest rate rule coefficients and the asymmetric preference parameters for the baseline case, which corresponds to the CBO output gap and CPI inflation. The squared output gap term, c_4, is highly significant over the pre-Volcker regime in the second column but loses most of its explanatory power during the later period in the third column (disregard the last column for the time being). The squared inflation term, c_3, appears relatively more relevant in the post-Volcker sample, though it is never statistically different from zero at the 5% significance level.

The estimates of the asymmetric preferences parameters are recovered from the feedback coefficients and the standard errors are computed using the delta method. Interestingly, α and γ take the expected signs and, in accord to the reduced-form estimates, the asymmetric preference on output is the significant parameter before 1979.[6] Specifically, a 0.3 estimate of γ implies on impact a 75 point basis cut of the interest rate

---

[6] The results are robust to letting the pre-Volcker sample begin in 1966:1 when the Federal funds rate first traded consistently above the discount rate.

in response to a negative 2% output gap but only a 42 point basis rise in response to a positive 2% gap. By contrast, after 1982 both coefficients become of limited importance and the Wald-type statistics in the second but last row indicates that the null hypothesis of symmetric preferences is not rejected at the 5% significance level, although it is rejected at the 10% level.

Finally, in order to gauge the forecasting advantages of the nonlinear (as opposed to the linear) monetary policy rule, we perform a version of the Diebold and Mariano (1995) test, which is designed to detect any difference in the predictive accuracy of two competing forecasts. To this end, we first compute the dynamically simulated fitted values of the two models and then we calculate the corresponding root-mean-squared error (RMSE) over both sub-samples. The RMSE of the linear model is 0.96 in the pre-Volcker period and 0.65 in the post-Volcker period, while the values of the nonlinear model are 0.78 and 0.63, respectively. The Diebold-Mariano test rejects the null hypothesis of *no difference* in the accuracy of the two specifications only during the pre-1979 regime, and it thus corroborates the results of the Wald-type tests for the presence of asymmetric preferences.

1.2.3 Robustness checks

We assess now in turn the robustness of our findings to subsample stability and to alternative measures of inflation and output gap. The last column of Table 1 displays the estimates for the sample 1987:3 - 2003:2, which corresponds to the tenure of Alan Greenspan as Fed Chairman. The squared inflation and output gaps do not have any explanatory power and translate into values of α and γ that are not statistically different from zero at any conventional level. This holds true also for their joint significance as shown by the p-value of the Wald test. Moreover, the parameter on inflation takes now a negative sign consistently with the view that deflation may have recently become the most imminent risk for the Fed.

Table 2 reports the estimates obtained using, everything equals, the rate of change in the GDP deflator as measure of inflation. The squared terms line up with those in Table 1 and translate into meaningful preference parameters. Specifically, the coefficient on output gap, γ, always takes a negative sign and is significant only during the pre-Volcker era, while the coefficient on inflation, α, is never statistically different from zero. Lastly, the Wald statistics confirm that asymmetric preferences matter before 1979, but not after 1982.

We re-estimate the policy rule (6) using CPI inflation and the Hodrick-Prescott filtered output. The results are shown in Table 3 and they bear out those from the previous tables. A significant, negative value of the feedback coefficient c_4 over the first sub-sample maps into a significant, negative value of the asymmetric preference on output, whereas no asymmetry is detected for inflation. Once more, the null hypothesis of symmetric preferences is rejected only during the pre-Volcker regime.

1.2.4 Structural estimates

One econometric issue we must confront with is that, in small samples, nonlinear GMM may be sensitive to the normalization of the orthogonality conditions (see Fuhrer, Moore and Schuh, 1995). Moreover, specific parameterizations of the central bank Euler equation may allow us to draw direct inference on the structural parameters α and γ. To address these issues, we rearrange the targeting rule in two alternative forms that we view as most natural for the problem at hand. To keep consistency with the reduced-form specification, we introduce a lagged interest rate. The first specification normalizes the coefficient on the inflation level to unity:

$$E_{t-1}\{[-\frac{\mu}{k\varphi}(i_t - i^*) + (1-\rho)\left((\pi_t - \pi^*) + \frac{\lambda}{k}y_t + \frac{\alpha}{2}(\pi_t - \pi^*)^2 + \frac{\gamma\lambda}{2k}y_t^2\right) + \rho(i_{t-1} - i^*)]\mathbf{z}_{t-1}\} = 0 \tag{7}$$

while the second normalizes the coefficient on the output gap level:

$$E_{t-1}\{[-\frac{\mu}{\lambda\varphi}(i_t - i^*) + (1-\rho)(\frac{k}{\lambda}(\pi_t - \pi^*) + y_t + \frac{\alpha k}{2\lambda}(\pi_t - \pi^*)^2 + \frac{\gamma}{2}y_t^2) + \rho(i_{t-1} - i^*)]\mathbf{z}_{t-1}\} = 0 \tag{8}$$

The latter specifications make it possible to estimate α and γ directly, and since these are the structural parameters of the model, we refer to the values inferred upon (7) and (8) as structural estimates.

An advantage of these normalizations relative to the reduced-form (6) is that they do not implicitly impose a non-zero value for the weight on the interest rate level stabilization μ. Moreover, to the extent that the inflation level and the output gap level significantly enter the central bank policy rule, as they virtually do in all empirical literature, the reduced-form coefficient on the interest rate gap $(i_t - i^*)$ is informative about μ such that a positive, significant value of the convolutions $(\frac{\mu}{k\varphi})$ and $(\frac{\mu}{\lambda\varphi})$ implies a positive, significant value for μ. While it is not possible to identify this policy preference parameter, we can evaluate whether it is statistically different from zero and since the test is performed on the convolution rather than on μ directly, we refer to it as a *t-type* test.

We estimate α and γ using nonlinear GMM and the set of instruments, $\mathbf{z}_{t-1}$, which includes the measures of inflation and output gap in the baseline case. The reduced-form coefficients are recovered from the estimates of the conditions (7) and (8) while the standard errors are computed using the delta method. The results for the first and the second normalization are reported in Table 4 and Table 5 respectively.

The structural estimates confirm, by and large, the reduced-form evidence. The implied c_is (i = 1, 2, 3 and 4) are in most cases not statistically different from the estimates of the previous tables and they provide empirical support for the presence of asymmetric preferences. The squared variables do never have explanatory power with the exception of the output gap in the pre-Volcker sample, whose estimate, c_4, is negative and significant. The structural parameter α is never statistically different from zero whereas the significant values of γ over the first sample are in line with the reduced-form estimates. In accord with the results of the previous tables, the joint null of symmetric central bank preferences, which is now directly tested on α and γ, is rejected before but not after 1979. Lastly, the t-type statistics for the null hypothesis $\mu = 0$ indicate that the central bank penalizes also the fluctuations of the interest rate level and therefore they validate the restriction implicitly imposed by the reduced-form representation (6).

1.2.5 Discussion

It is useful at this point to compare our estimates with the results from some recent studies that also focus on the policy regime shift of 1979. Clarida, Galí and Gertler (2000) estimate a forward-looking linear reaction function for the pre-Volcker period and report values of 0.68 for the coefficient on CPI inflation (s.e.= 0.06) and 0.28 for the coefficient on CBO output (s.e.= 0.08). Their estimates suggest that neglecting the squared output gap, which significantly enters our empirical specification with a negative sign, introduces a downward bias in the linear estimate.[7] Turning to the nonlinear specifications, Dolado, Maria-Dolores and Ruge-Murcia (2004) use a Clarida, Galí and Gertler-type of rule augmented with a generated regressor for the conditional variance of inflation and find no evidence for this form of nonlinearity. Kim, Osborne and Sensier (2004) use a semi-parametric method of estimation and show that only the asymmetry over the output gap has been sizable.

The post-Volcker estimates of the parameters on the inflation level and the output gap level are not statistically different from the values reported in Clarida, Galí and Gertler (2000), and therefore they confirm a limited role for nonlinearity during the last two decades. These results are consistent with those in Kim, Osborne and Sensier

---

[7] This result holds true also for the alternative measures of inflation and output gap.

(2004) while they are only marginally so with those in Dolado, Maria-Dolores and Ruge-Murcia (2004). The absence of an output gap objective in the latter however seems a natural candidate to explain the difference. Lastly, we line up with earlier contributions in that the coefficient on the inflation level becomes bigger than one moving from the pre- to the post-1979 period.

1.3 The average inflation bias

The estimates of the previous section support the notion of a novel inflation bias due to Cukierman (2002). In the presence of an asymmetric objective over the output gap and uncertainty about the state of the economy, the monetary authorities face an incentive to respond more aggressively to output contractions of a given amount than to output expansions of the same magnitude. The reason is that the expected marginal benefit of a policy intervention is convex in the output gap, meaning that to satisfy the Euler equation and stimulate aggregate demand the policy makers cut the interest rate by more the worse the economic outlook is. As the private sector correctly anticipates such an incentive, the precautionary stance of the monetary policy generates a systematic boost in inflation expectations even though, unlike in Barro and Gordon (1983), the central bank targets output at potential.[8]

1.3.1 A model-based measure of the inflation mean

This section proposes a simple strategy to measure the asymmetric preferences induced inflation bias, which is defined as the difference between the model-based inflation mean and the inflation target. The resulting expression is isomorphic to the one that Surico

---

[8] In the theory of consumption, a precautionary motive emerges from the interaction between non-quadratic preferences and labor income risks such as to generate above-average saving rates in periods of high uncertainty. As shown by Kimball (1990), a necessary and sufficient condition for a precautionary saving is that the expected marginal *utility* be *convex* in consumption. Analogously here, the above-average inflation comes from the interaction between an asymmetric central bank objective and uncertainty about the state of the economy. Moreover, as the expected marginal *loss* is *concave* in the output gap, this motive can be thought as a precautionary demand for expansions.

(2003) derives as the difference between the optimal policies under discretion and under commitment using an asymmetric central bank objective and a Lucas aggregate supply.

On the basis of the empirical results presented in the previous section, we impose the restriction $\alpha = 0$ into the first order condition of the central bank optimization problem (5).The corresponding augmented targeting rule writes

$$E_{t-1}\{[-(i_t - i^*) + (1-\rho)(c_1(\pi_t - \pi^*) + c_2 y_t + c_4 y_t^2) + \rho(i_{t-1} - i^*)]z_{t-1}\} = 0 \qquad (9)$$

where the parameters are written in reduced-form for expositional convenience.

The maintained assumption that the target i^* equals the sample mean of interest rate, combined with the empirically grounded restriction of a symmetric preference over inflation allow us to uniquely identify the inflation target. To see this, notice that the constant in the above expression becomes nothing but the convolution $(-c_1\pi^*)$.[9] The average inflation bias can then be computed by taking the unconditional expectation of equation (9). According to the model, the inflation mean corresponds to the following expression:

$$E(\pi_t) = \pi^* - \frac{c_4}{c_1}\sigma_y^2 = \pi^* - \frac{\gamma\lambda}{2k}\sigma_y^2 \qquad (10)$$

where we have used the fact that the output gap has an unconditional distribution with zero mean and variance σ_y^2.

The average inflation bias arises here because policy preferences are asymmetric with respect to the output gap rather than because the desired level of output is above potential like in Barro and Gordon (1983). The distortion increases with the degree of asymmetry, and to the extent that the penalty associated to an output contraction is larger than the penalty associated to an output expansion of the same size, the model predicts $\gamma < 0$. As λ and k are positive, the difference between the model-based inflation

---

[9] It is worth noticing that the assumption on the interest rate target should bias, if any, the inflation target towards the sample mean of inflation. This suggests that our estimates are likely to understate the contribution of the asymmetric preferences induced bias to the actual mean of US inflation.

mean and the inflation target represents an inflation bias rather than a deflation bias. When γ is equal to zero, the expected marginal benefit of a policy intervention becomes linear and the inflation bias disappears together with the precautionary motive.

The average inflation bias is proportional to the variance of the output gap and, as shown by the first equality in (10), it is inversely related to the inflation slope of the targeting rule (9). Hence, the model is general enough to confront the explanatory power of a change in the asymmetric preference parameter over the output gap, γ, with two alternative interpretations of the behavior of US inflation. The first is a shift in the response to the inflation level as captured by c_1. The second is a difference in the variance of the shocks as proxied by σ_y^2.

1.3.2 Measuring the bias

We estimate equation (9) using GMM with a four lag Newey-West estimate of the covariance matrix. The measures of inflation and output gaps and the instrumental variables refer to the baseline case. The only difference relative to Table 1 is that, in line with the restriction $c_3 = 0$, the four lags of the squared inflation are not included here as instruments. The results are shown in Table 6 and they turn out to be sufficiently close to those reported in the previous tables that we do not comment further. The restrictions discussed above allows us to identify the inflation target, which is found to move from 3.61% before 1979 to a statistically lower 2.77% during the last two decades. Interestingly enough, this result contrasts with most of the empirical literature on monetary policy rules that, neglecting asymmetric preferences on the output gap and therefore imposing a linear reaction function, usually find a difference in π^* across subsamples of two-to-three percentage points.

We use the estimates of table 6 to compute the inflation bias implied by the model, $\left(-\frac{c_4}{c_1}\sigma_y^2\right)$, and the delta method to obtain the standard errors. Table 7 displays the results. The average inflation bias, which is reported in the second row, is sizable and statistically different from zero only in the pre-Volcker period. The model-based inflation mean in the fourth row confirms that we effectively decompose the actual inflation mean into a target and a bias argument. Moreover, a shift in the policy preferences on output stabilization appears to account for a larger fraction of the difference in the sub-samples mean of inflation relative to a reduction in the inflation target.

The results in Table 6 and Table 7 suggest that while a different interest rate response to the inflation level, as described by the rise of c_1, and a more favorable macroeconomic environment, as summarized by the decline in the standard deviation of the output gap, have also played a role, a change in the policy preference on output from asymmetric to symmetric appears crucial to account for the observation that US inflation has been on average higher during the 1960s and 1970s than during the 1980s and 1990s.

1.4 Conclusions

The contribution of this paper is twofold. At the theoretical level it derives the analytical solution of the central bank optimization problem when the policy preferences are asymmetric in *both* inflation and output gaps, and the monetary transmission mechanism is New-Keynesian. The specification of the policy objectives is general enough to nest the quadratic form as a special case and therefore it translates into a potentially nonlinear targeting rule. This feature forms the basis of our hypothesis testing for the presence of asymmetric preferences as it allows to reversely engineer potential evidence of nonlinearities in the reaction function into evidence of asymmetries in the policy objective.

At the empirical level this paper shows that US monetary policy can be effectively characterized by a nonlinear policy rule only during the pre-Volcker regime, with the interest rate response to the state of the business cycle being the dominant type of nonlinearity. In particular, the Fed appears to have historically attached a larger weight to output contractions than to output expansions of the same magnitude such as to induce an average inflation bias of 1.11%. The latter can account for a sizable fraction of the inflation rise observed during the 1960s and 1970s. These findings are robust across alternative measures of inflation and output gap, as well as across alternative estimation strategies.

Altogether, this paper provides empirical support for asymmetric preferences and suggests some caution about using symmetric loss functions as a guide to policy analysis. Promising strands of literature have recently emphasized that political pressures, labor market frictions and heterogeneity in portfolio holdings can make the costs of business fluctuations and inflation variation asymmetric. Along these lines, a stimulating avenue for future research is to derive an utility-based welfare function within richer models of the business cycle in order to provide a formal microfoundation for an asymmetric central bank objective.

Measuring the Time-Inconsistency of US Monetary Policy

December 2003

Abstract

This paper offers an alternative explanation for the behavior of postwar US inflation by measuring a novel source of monetary policy time-inconsistency due to Cukierman (2002). In the presence of asymmetric preferences, the monetary authorities end up generating a systematic inflation bias through the private sector expectations of a larger policy response in recessions than in booms. Reduced-form estimates of US monetary policy rules indicate that while the inflation target declines from the pre- to the post-Volcker regime, the average inflation bias, which is about one percent before 1979, tends to disappear over the last two decades. This result can be rationalized in terms of the preference on output stabilization, which is found to be large and asymmetric in the former but not in the latter period.

JEL Classification: E52, E58

Keywords: asymmetric preferences, time-inconsistency, average inflation bias, US inflation

---

[9] I wish to thank Alberto Alesina, Filippo Altissimo, Efrem Castelnuovo, Sheng-Hung Chung, Günter Coenen, Carlo Favero, Jordi Galí, Tommaso Monacelli, Anton Muscatelli, Jorges Rodrigues, Massimo Rostagno and Guido Tabellini for very useful comments. This paper has been prepared while the author was visiting the European Central Bank whose kind hospitality is gratefully acknowledged.

2 Measuring the time-inconsistency of US monetary policy

The behavior of postwar US inflation is characterized by two major episodes. The first is an initial rise that extends from the 1960s through the early 1980s. The second is a subsequent fall that lasts from the early 1980s to the present day. The important change that underlies such a path can be exemplified by the average rates reported in the second column of Table 1. Inflation is measured as the annualized quarterly increase in the log GDP chain-type price index whereas the output gap is constructed as the log deviation of real GDP from the Congressional Budget Office potential output. The difference of the average inflation rates across the two sub-samples is above 2% and it is echoed by the decline in the volatility of the output gap displayed in the third column.

While a more favorable macroeconomic environment and a better policy management during the last two decades or a persistent error in the real-time estimates of potential output during the 1970s are also likely to have played a role, an important strand of the literature has investigated whether the time-consistency problem can explain the behavior of US inflation.

In a stimulating contribution, Ireland (1999) shows that Barro and Gordon's (1983) model of time-consistent monetary policy imposes long-run restrictions on the time series properties of inflation and unemployment that are not rejected by the data. In the absence of a commitment technology, the monetary authorities face an incentive to surprise inflation in an effort to achieve a lower level of unemployment through an expectations-augmented Phillips curve. However, such an optimal plan is not time-consistent in the sense of Kydland and Prescott (1977), and private agents, who rationally understand such a temptation, adjust their decisions accordingly. In equilibrium, unemployment is still at its first-best level but the rate of inflation is inefficiently higher than it would otherwise be. This is the celebrated inflation bias result, according to which the higher the natural rate of unemployment the more severe the time-consistency problem of monetary policy is.

As Persson and Tabellini (1999) make clear, the central bankers' ambition of attaining a level of unemployment below the natural rate is crucial to generate the kind of inflation bias a la Barro and Gordon (1983), and both researchers and policy makers have challenged such an assumption on the ground of realism. McCallum (1997) argues that were this the case, the monetary authorities would learn by practicing the time-inconsistency of their actions and eventually would revise their objective. Describing his experience as vice-Chairman, Blinder (1998) claims that the Fed actually targets the natural rate of real activity, thereby suggesting that overambitious policy makers cannot be at the root of any kind of inflation bias. While this may rationalize the failure of the theory to account for the short-run inflation dynamics (see Ireland, 1999), it does not necessarily imply that the time-consistency problem has been unimportant in the recent history of US monetary policy.

In an intriguing article, Ruge-Murcia (2003) constructs a model of asymmetric central bank preferences that nests the Barro-Gordon model as a special case. When applied to the full postwar period, the hypothesis that the Fed targets a level of real activity different from the natural rate is rejected but the hypothesis that it weights more severely output contractions than output expansions is not. This suggests the existence of a novel *average inflation bias* that according to Cukierman (2002) comes from the private sector expectations of a more vigorous policy response in recessions than in booms.

Specifically, the average inflation bias is a function of both the preferences of the central bank and the volatility of the output gap. To the extent that a significant policy regime shift has occurred at the beginning of the 1980s after the appointment of Paul Volcker as Fed Chairman, it is likely that the degree of asymmetry and therefore the degree of time-inconsistency has also changed during the last four decades. Hence, rather than focusing on the full postwar period like Ireland (1999) and Ruge-Murcia (2003), we study the sub-samples that are typically associated with a shift in the conduct of US monetary policy according to the reasoning that the time-inconsistency problem and the relative inflation bias are best interpreted as regime-specific. The difference in the sub-sample volatility of the output gap shown in the third column of Table 1 also seems consistent with this view.

This paper contributes to the literature on optimal monetary policy by proposing a measure of the average inflation bias that arises in a model of asymmetric central bank preferences. To this end, it is developed a novel identification strategy that allows to recover the relevant parameters in the central bank objective function and, most importantly, to translate them into a measure of time-inconsistency. The comparison between the commitment and the discretionary solutions shows how the observed inflation mean can be successfully decomposed into a target and a bias argument, a result that to our knowledge of the existing literature comes as new. Reduced-form estimates of US monetary policy rules indicate that a significant regime shift has occurred during the last forty years as measured by the change in the Fed policy preferences. In particular, while the inflation target declines from 3.42% to 1.96%, the average inflation bias, which is estimated at 1.01% before 1979, is found to disappear over the last two decades. The result can be rationalized in terms of the policy preference on output stabilization, which is found to be large and asymmetric in the pre- but not in the post-Volcker period.

The paper is organized as follows. Section 2 sets up the model and solves for the optimal monetary policy. Section 3 derives its reduced-form version and reports the estimates of both the feedback rule coefficients and the average inflation bias. Section 4 concludes.

2.1 The model

Following the literature, the private sector behavior is characterized by an expectations-augmented Phillips curve:

$$y_t = \theta\left(\pi_t - \pi_t^e\right) + u_t, \ \ \theta > 0 \tag{11}$$

where y_t is the output gap measured as the difference between actual and potential output, π_t denotes inflation and π_t^e stands for the expectations on the inflation rate in period t from the standpoint of period $t-1$. The supply disturbance, u_t, obeys a potentially autoregressive process $u_t = \rho u_{t-1} + \varepsilon_t$ where $\rho \in [0,1)$ and ε_t is an i.i.d. shock with zero mean and variance σ_ε^2. The private sector has rational expectations

$$\pi_t^e = E_{t-1}\pi_t \tag{12}$$

with E_{t-1} being the expectation conditional upon the information available at time $t-1$.

Potential output is identified with the real GDP trend so that the mean of the output gap is normalized to zero. Moreover, y_t is also a random variable as it depends on u_t, and its variance, which is a positive function of both ρ and σ_ε^2, is denoted by σ_y^2.

As customary in the literature, the central bank is assumed to have full and direct control over inflation, which is chosen to minimize the following intertemporal criterion:

$$\underset{\{\pi_t\}}{Min} \ \ E_{t-1}\sum_{\tau=0}^{\infty}\delta^\tau L_{t+\tau} \tag{13}$$

where δ is the discount factor and L_t stands for the period loss function. The latter is specified in a cubic form:

$$L_t = \frac{1}{2}\left(\pi_t - \pi^*\right)^2 + \frac{\lambda}{2}\left(y_t^2 + \frac{\gamma}{3}y_t^3\right) \tag{14}$$

where $\lambda > 0$ and γ represent the relative weight and the asymmetric preference on output stabilization, respectively. The inflation target, π^*, is assumed stable enough to be approximated by a positive constant that possibly differs across sub-samples. Unlike in the Barro-Gordon model, the target level of output is not meant to overambitiously exceed potential. This is consistent with the empirical evidence reported by Ruge-Murcia (2003).

The objective function (14) departures from the quadratic form in that policy makers are allowed, but not required, to treat differently output contractions and output expansions. Indeed, under a cubic specification deviations of the same size but opposite sign yield different losses and a negative value of γ implies that negative output gaps are weighted more severely than positive ones. To see this notice that whenever $y_t < 0$ the cubic term, γy_t^3, is positive and amplifies the penalty due to the quadratic component whereas for $y_t > 0$ the quadratic and the cubic terms move in opposite directions.

The cubic form nests the quadratic specification as a special case and whenever γ is equal to zero the central bank objective function (14) reduces to the conventional symmetric parametrization $L_t = \frac{1}{2}\left[(\pi_t - \pi^*)^2 + \lambda y_t^2\right]$. This feature is attractive as it allows us to test whether the relevant preference parameter is statistically different from zero. Figure 1 compares the standard quadratic with the asymmetric cubic function using the historical values of the output gap and a value of γ that is consistent with the estimates reported below.[10]

The specification of an asymmetric loss with respect to the output gap only is motivated by empirical as well as theoretical considerations. At the empirical level, Surico (2003b) derives a general, nonlinear interest rate rule within a model of nonquadratic preferences over both inflation and output, and finds evidence of an asymmetric objective for the latter but not for the former variable. At the theoretical level, Geraats

---

[10] The cubic specification can also be interpreted as some third-order approximation around $(\pi_t - \pi^*) = 0$ and $y_t = 0$ to the linex function proposed by Nobay and Peel (2003), and employed by Geraats (1999), Ruge-Murcia (2003) and Surico (2003a). The advantage of using the cubic form as the primitive function is that it does not require any approximation of the optimal monetary policy rule. Nevertheless, for a realistic range of values of y_t like $[-0.08, 0.06]$, and given the estimates of γ reported below, the cubic and the linex function behave very similarly.

(1999) shows that the labor market flows over the business cycle provide a natural microfoundation for an asymmetric welfare criterion as the firms' hiring-firing decisions are mainly taken along the extensive margin during recessions but along the intensive margin during booms.

2.1.1 Commitment

This subsection solves for the optimal monetary policy under commitment. Because no endogenous state variable enters the model, the intertemporal policy problem reduces to a sequence of static optimization problems. Accordingly, the monetary authorities, who can manipulate inflation expectations, choose both planned inflation, π_t, and expected inflation, π_t^e, to minimize the asymmetric loss function (14) subject to the augmented Phillips curve (11) and to the additional constraint (12) imposed by the rational expectations hypothesis. The corresponding first order conditions are, respectively:

$$(\pi_t - \pi^*) + E_{t-1}\left\{\lambda\theta\left[y_t + \frac{\gamma}{2}y_t^2\right] - \mu\right\} = 0 \tag{15}$$

$$-E_{t-1}\left\{\lambda\theta\left[y_t + \frac{\gamma}{2}y_t^2\right]\right\} + \mu = 0 \tag{16}$$

with μ being the Lagrange multiplier associated to the rational expectation constraint. Combining the optimality conditions (15) and (16) to eliminate μ, and taking expectations of the resulting expression produce

$$E(\pi_t) = \pi^* \tag{17}$$

where we have used the law of iterated expectations to get rid of E_{t-1}. Equation (17) states that the planned inflation rate equals on average the socially desirable inflation rate and therefore it is independent of the output gap.

2.1.2 Discretion

If commitment is infeasible, the monetary authorities choose the inflation rate π_t at the beginning of the period after the private agents have formed their expectations but before the realization of the real shock u_t. Accordingly, the discretionary solution reads

$$(\pi_t - \pi^*) + E_{t-1}\left\{\lambda\theta\left[y_t + \frac{\gamma}{2}y_t^2\right]\right\} = 0 \tag{18}$$

It is instructive at this point to compare the solution obtained under asymmetric preferences with the solution obtained under the standard quadratic case. Whenever $\gamma = 0$, the optimal monetary policy becomes

$$(\pi_t - \pi^*) = -\lambda\theta E_{t-1}(y_t) \tag{19}$$

This implies that under quadratic preferences there exists a one to one mapping between the inflation bias and the output gap conditional mean. Moreover, in the face of white noise supply disturbances (i.e. $\rho = 0$) the inflation bias is zero reflecting the notion of potential output targeting.

To compute the *average inflation bias*, we take expectations of equation (18), and using the fact that the unconditional mean of the output gap is zero, we obtain the following expression:

$$E(\pi_t) = \pi^* - \frac{\lambda\theta\gamma}{2}\sigma_y^2 \tag{20}$$

The comparison between the expected rates under commitment (17) and under discretion (20) illustrates the source of a novel *average inflation bias*. Like in the Barro-Gordon model, the time-inconsistency of monetary policy arises here because the policy makers face an incentive to surprise inflation. However, the nature of the incentive in the two models is very different. In Barro-Gordon (1983) this is the central bank desire to push the economy beyond its potential level. Here, it is the asymmetric concern about the business cycle that associates a more aggressive policy response to output contractions than to output expansions (i.e. $\gamma < 0$). As the private sector correctly anticipates such an incentive, the inflation rate systematically exceeds the first-best solution attainable under commitment *even though* the monetary authorities target output to potential. Moreover, the bias is higher the larger and the more asymmetric the policy preference on output stabilization is.

Possible improvements to the discretionary solution include the appointment of a more conservative central banker, who is one endowed with a lower relative weight λ in the spirit of Rogoff (1985) and/or a lower inflation target than society, or the appointment of a more symmetric policy maker, who is one endowed with a smaller absolute value of γ. Lastly, the average inflation bias is proportional to the variance of the output gap as the marginal benefit of an inflation surprise is convex in the output gap. When γ is equal to zero as it is in equation (19), such a marginal benefit becomes linear and the *average* inflation bias disappears together with the precautionary motive. This feature parallels the precautionary motive result in the theory of consumption according to which non-quadratic preferences and labor income risks generate above-average saving rates in periods of high uncertainty.

2.2 The evidence

This section investigates the empirical merits of the asymmetric preference model to account for the behavior of postwar US inflation. The analysis spans the period 1960:1-2002:3 and it is conducted on quarterly, seasonally adjusted data that have been obtained in February 2003 from the web site of the Federal Reserve Bank of St. Louis. Inflation is measured as the annualized change in the log of the GDP chain-weighted price index, whereas the output gap is constructed as the difference between the log of the real GDP and the log of the real potential output provided by the Congressional Budget Office.

To make our results comparable with those reported by Ruge-Murcia (2003), we first consider the whole sample. Then, we use our identification strategy to estimate the asymmetric preference and to obtain a measure of the inflation bias for both the pre- and the post-Volcker regimes. We also address the issue of sub-sample stability by re-estimating the model over Greenspan's tenure, which begins in the third quarter of 1987. Indeed, equation (20) makes it clear that the inflation bias is a function of policy makers' preferences and therefore it can only be interpreted as regime-specific. To the extent that a significant break has occurred in the conduct of US monetary policy during the last forty years, our identification scheme provides a sharper evaluation of the model by measuring the time-inconsistency across the two eras.

2.2.1 Preliminary analysis

As a way to illustrate the potential relevance of the asymmetric preferences induced inflation bias, we consider a testable prediction of the quadratic preference model. According to equation (19), the conditional mean of the output gap is informative about the difference between the realized inflation and the inflation target. Moreover, in the face of i.i.d. supply shocks the conditional mean and therefore the inflation bias should be zero reflecting the notion of quadratic preferences and potential output targeting.

Figure 2 displays the kernel estimates of the output gap conditional mean (with the sign switched) over the full sample using the Nadaraya-Watson estimator, a second order Gaussian kernel and the likelihood cross validation procedure to obtain a value for the fixed bandwidth parameter. The results are unaffected by using the least squares cross validation criterion and an higher order kernel. Before proceeding however it is important to stress what we are not doing in this exercise. In particular, we are not using the output gap as the dependent variable while estimating the optimality condition (19). Rather, we are computing from the bivariate time-series model of inflation and output the conditional mean of the output gap which according to the model of quadratic preferences and potential output targeting is the measure of the inflation bias at each point in time.

A couple of interesting results emerge from Figure 2. First, the third quarter of 1982 appears to witness the beginning of a new era as represented by the intersection between the lower bound of the 95% confidence interval and the zero line. This is consistent with the conventional wisdom that a regime-switch in the conduct of US monetary policy has occurred at the beginning of the 1980s, especially with the end of the so-called 'Volcker experiment' of non-borrowed reserves targeting that Bernanke and Mihov (1998) date in 1982:3. Second, the measure of the inflation bias displays a fairly different pattern across the two periods moving from the significant estimates of the 1970s to values that are not statistically different from zero during the last two decades. Although also a change in the persistence of the supply shocks may account for part of the difference, we stress that the nonparametric evidence over the earlier sample rejects a model of quadratic preferences, potential output targeting and i.i.d. disturbances. Given the popularity of these assumptions in the literature, we interpret this finding as a call for an extension of the theory. We return to the identification of asymmetric preferences versus persistent supply shocks in the discussion of the empirical results.

2.2.2 The reduced-form

We solve equation (18) for π_t and prior to estimation we replace expected output gaps with actual values. The empirical version of the feedback rule is given by:

$$\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t \tag{21}$$

which is linear in the coefficients

$$\alpha = -\lambda\theta \quad \text{and} \quad \beta = -\frac{\lambda\theta\gamma}{2}$$

and whose error term is defined as

$$v_t \equiv -\left\{\alpha\left(y_t - E_{t-1}y_t\right) + \beta\left[y_t^2 - E_{t-1}\left(y_t^2\right)\right]\right\}$$

The term in curly brackets is a linear combination of forecast errors and therefore v_t is orthogonal to any variable in the information set available at time $t-1$.

Equation (29) reveals that by assuming an optimizing central bank behavior the reaction function parameters can only be interpreted as a convolution of the coefficients representing policy makers' preferences and those describing the structure of the economy. Nevertheless, the reduced-form parameters allow us to identify both the asymmetric preference on the output gap and the average inflation bias. The asymmetric preference is $\gamma = 2\beta/\alpha$ while the bias amounts to $\beta\sigma_y^2$. The latter is obtained as the difference between the solution of the central bank optimization under commitment (17) and the solution under discretion (20).

2.2.3 Empirical results

To the extent that the penalty associated to an output contraction is larger than the penalty associated to an output expansion of the same size, the model predicts $\gamma < 0$, $\alpha < 0$ (since λ, $\theta > 0$), and $\beta > 0$. Moreover, while also persistent supply shocks imply a significant role for the level of the output gap, only asymmetric preferences are crucial for the prediction that the squared output gap is helpful to forecast inflation.

The orthogonality conditions implied by the rational expectation hypothesis makes the Generalized Method of Moments (GMM) a natural candidate to estimate equation (29). This has also the advantage that no arbitrary restrictions need to be imposed on the information set that private agents use to form expectations. To control for possible heteroskedasticity and serial correlation in the error terms we use the optimal weighting scheme in Hansen (1982) with a four lag Newey-West estimate of the covariance matrix. Three lags of inflation, output gap and squared output gap are used as instruments corresponding to a set of 7 overidentifying restrictions that can be tested for. The choice of a relatively small number of instruments is meant to minimize the potential small sample bias that may arise when too many overidentifying restrictions are imposed. We also check the robustness of our results to changes in the instrument set. In particular, we re-estimate the model using five lags of inflation and two lags of output gap and squared output gap. The F-test applied to the first stage regressions, which Staiger and Stock (1997) argue to be important in evaluating the relevance of the instruments, always rejects the null of weak correlation between the endogenous regressors and the variables in the instrument sets.

Table 2 reports the estimates of the feedback rule (29) for the full sample. Each row corresponds to a different set of instruments. The parameter on the output gap, α, is not statistically different from zero whereas the parameter on the squared output gap, β, is significant and positive. The estimates of the slope coefficients as well as the estimates of the inflation target are robust to the instrument selection and the hypothesis of valid overidentifying restrictions is never rejected. These results are similar to those reported by Ruge-Murcia (2003) as they confirm the presence of asymmetric preference using a different method of estimation and a different measure of real activity.

Table 3 reports the estimates for the pre- and post-Volcker regimes. We remove from the second sub-sample the period 1979:3-1982:3 when the temporary switch in the Fed operating procedure documented by Bernanke and Mihov (1998) appears to be responsible for the failure to gain control over inflation. The sample selection is also consistent with the nonparametric evidence reported in the preliminary analysis.

The first two rows of Table 3 refer to the pre-Volcker era and show large negative values for the level of the output gap besides to positive and significant parameters for its squared. The point estimates of the inflation target range from 3.42% to 3.69% while the asymmetric preference parameter is negative and statistically significant. These results sharply contrast with the post-1979 values that are displayed in the middle rows and the bottom rows of Table 3. Indeed, not only the inflation target statistically declines to values around 2%, but also the impact of the output gap level on inflation appears to be weaker, although still significant. To the extent that the structure of the economy has remained stable during the last forty years, a smaller value of α can only be rationalized by a decline in λ, which corresponds to a more conservative monetary policy stance. The most dramatic difference between the two regimes emerges however on the squared output gap, which actually loses explanatory power for both set of instruments as well as for both post-1979 samples. This translates into values of the policy parameter γ that are not statistically different from zero.

Turning to the measure of the asymmetric preference induced time-inconsistency, Table 4 reports the estimates of the average inflation bias. According to equation (20), the bias is a convolution of the structural parameters of the model and the variance of the output gap. Given the decline in the latter reported in the third column of Table 1, we expect also the inflation bias to decline moving from the pre- to the post-Volcker period. This seems consistent with the change in the volatility of the supply shocks documented by Hamilton (1996) between the 1970s and the 1980s.

The second column of Table 4 shows the measure of the average inflation bias implied by the reduced-form estimates of Table 3. The first block reports the pre-Volcker values whose point estimates range from 1.01% in the baseline case to 1.36% for the alternative instrument set. By contrast, the inflation bias is found to be not statistically different from zero over the post-1979 era, reflecting the fact that US monetary policy can be characterized by a nonlinear feedback rule during the former but not during the latter period. Empirical support for this form of regime shift can also be found in the cross-country evidence over 22 OECD economies reported by Cukierman and Gerlach (2003).

Lastly, the realized inflation mean over the pre-1979 sample falls in the range of estimates implied by the sum of the inflation target and the inflation bias while its post-Volcker counterparts appear to be higher than the model predicts. This suggests that the theory can effectively decompose the observed inflation mean into a measure of the target and a measure of the bias over the pre-1979 regime, though it needs to be extended to account more fully for the gap that appears in the data over the last two decades.

2.3 Concluding remarks

This paper develops a method to measure the time-inconsistency of monetary policy when the preferences of the central bank are asymmetric. As demonstrated by Cukierman (2002), if policy makers are more concerned about output contractions than output expansions, an inflation bias can emerge *on average* even though output is targeted at potential. In addition, both casual observations and formal empirical analyses challenge the predictions of the Barro-Gordon model by arguing that the Fed's desired level of output does not exceed the natural rate (see Blinder, 1998, and Ruge-Murcia, 2003).

Using a model of asymmetric preferences and potential output targeting, it is shown how the observed inflation mean can be successfully decomposed into a target and a bias argument. When applied to postwar US data, our identification method indicates that the target is 3.42% and the bias 1.01% during the pre-1979 policy regime. By contrast, over the last two decades the inflation target declines to 1.96% while the average inflation bias tends to disappear. This result can be rationalized by the fact that the policy preference on output stabilization is found to be large and asymmetric before but not after the appointment of Paul Volcker as Fed Chairman. Although other factors such as an inconvenient policy making and unfavorable supply shocks are also likely to have played a role, this paper provides empirical support and quantitative measures of a new, additional explanation for the behavior of US inflation during the 1970s.

While suggestive, the results reported in this paper are based on a simple model, and the specification of a richer structure of the economy is likely to produce also a state-contingent bias as well as a stabilization bias. However, as shown by Svensson (1997) and Cukierman (2002), the average inflation bias would then be larger than it is with a standard expectations-augmented Phillips curve. This suggests not only that our estimates are better interpreted as a lower bound but also that a richer specification of the private agents' behavior may account for the gap between the model-based average inflation and the actual average inflation during the last two decades. Given our limited knowledge of the channel(s) through which the time-consistency problem affects policy outcomes, measuring and disentangling the inflation bias remains a challenging topic for future research.

Model Uncertainty, Optimal Monetary Policy, and the Preferences of the Fed

(joint with Efrem Castelnuovo)

September 2002

Abstract

US monetary policy is characterized by a substantial degree of inertia. While in principle this may well be the outcome of an optimizing central bank behaviour, the ability of any derived policy rule to match the data relies on so large weights for interest rate smoothing into policy makers' preferences as to be theoretically flawed. In this paper we investigate whether such a puzzle can be interpreted as resulting from the concern of monetary authorities for potential misspecifications of the macroeconomic dynamics. Accordingly, we use a novel *thick modeling* approach to incorporate model uncertainty into the identification of central bank's preferences. The robust *thick* policy rule shows the kind of smoothness observed in the data without resorting to implausible values for the preference parameters.

JEL classification: C61, E52, E58.

Keywords: Model uncertainty, interest rate smoothing, Fed policy preferences, robust optimal monetary policy.

---

[10] Awarded the Scottish Economic Society's *Sir Alex Cairncross Prize* 2002 for the best paper by a young economist. We wish to thank Carlo Favero and Guido Tabellini for precious suggestions. Richard Dennis, Jordi Galí, Massimiliano Marcellino, Michael McAleer, Anton Muscatelli, Brian Sack, Tommaso Monacelli, Stefania Albanesi, Gaia Narciso, Marco Aiolfi and Saverio Simonelli provided very useful comments.

3 Model uncertainty, optimal monetary policy and the preferences of the Fed

The US Federal Reserve tends to change short-term interest rates by small steps that move in a particular direction over sustained periods and reverse only infrequently (see Rudebusch, 1995, and Goodhart, 1997). This prominent feature of policy rates, which is interchangeably referred to as interest rate smoothing, policy gradualism or policy inertia, characterizes the Fed response to inflation and output gaps as having been more moderate than an optimizing central bank behavior would predict.

In a recent survey of evidence, Sack and Wieland (2000) interestingly discuss several explanations to reconcile historical and optimal policy rules. A number of empirical studies find that uncertainty creates incentives to smooth policy rates, in the form of either parameter uncertainty or measurement error for inflation and output gap. Parameter uncertainty, which is the uncertainty on the monetary transmission mechanism, alters the knowledge of decision makers about the impact of policy action on the economy. Accordingly, a central bank that adjusted aggressively policy rates to the developments in the economy would be more likely to have unpredictable and therefore undesirable movements of output and inflation. Then, as shown in the VAR analyses by Sack (2000), Salmon and Martin (1999), and Söderström (1999), policy gradualism may be the optimal strategy to bring the relevant macroeconomic variables in line with the targets.

Another source of uncertainty comes from the measurement errors on inflation and output gap. Indeed, the evaluation of monetary policy in most empirical studies relies on the unrealistic assumption that policy makers know the state of the economy without error. However, monetary policy mainly involves decisions that are based on real-time available information, which are subject to frequent revisions after the initial release. Interestingly, Orphanides (1998) shows that whenever policy makers take data uncertainty into account the estimated policy response to inflation and output gaps is more moderate, thereby preventing the possibility of wide interest rate fluctuations due to measurement errors. This attenuation turns out to be particularly relevant under simple policy rules, although it also emerges for optimal policy rules.

These explanations have each proved to be statistically significant, although none alone has resulted to be quantitatively satisfactory (see Sack and Wieland, 2000). Moreover, interest rate smoothing is derived as the optimal policy rule of a central bank whose only concerns are to stabilize output and inflation and the possibility that policy makers have an explicit preference to penalize policy rate fluctuations is ruled out by assumption.

On the positive side, the inclusion of interest rate changes in the policy makers' loss function can be justified on several grounds (see Woodford, 2002, Ch. 7; Goodfriend, 1991 and Lowe and Ellis, 1997). The empirical model proposed by Rudebusch and Svensson (1999), which includes an explicit interest rate smoothing goal, has become by now a popular framework to analyze monetary policy under uncertainty (see Stock, 1999; Smets, 1999; Onatski and Stock, 2002; Rudebusch, 2001 and Favero and Milani, 2001). For example, Rudebusch (2001) argues that the interaction of several forms of uncertainty rather than a single one is likely to generate the kind of smoothness observed in the data and points towards measurement errors and model misspecifications as the most relevant candidates. In particular, the perturbation of some key structural relations such as the inflation dynamics and the output sensitivity to interest rate are shown, everything equals, to make smoother an otherwise volatile policy rate behavior, thereby being an excellent starting point for the present analysis.

On the negative side, the ability of any optimal policy rule to match the data badly relies on so large weights for the policy makers' aversion to interest rate changes as the theory cannot easily motivate. This suggests the potential for a strictly related issue, namely the identification of the Fed policy preferences. Indeed, several pioneering studies have proposed alternative strategies to estimate the structural parameters in a small empirical model à la Rudebusch and Svensson (see Favero and Rovelli, 2002; Dennis, 2001; Ozlale, 2001). While extremely promising, these estimates have left the *interest rate smoothing puzzle* unsolved in that any plausible set of preferences implies an optimal path for policy rates much more volatile than the observed one.

In this paper we bring together the literature on model uncertainty and the one on central bank's preferences by using the progresses made in the former to solve the puzzle emerged in the latter. To this end, we incorporate model uncertainty in the simple calibration method we propose to identify the Fed policy preferences. In so doing, we investigate whether the concern for model misspecifications can explain the inertial behavior of policy rates without resorting to implausible weights, if any, for an interest rate smoothing goal.

The intuition for having more moderate policy responses when the model is misspecified comes from the policy makers' agnosticism about what model provides the most accurate description of the economy. Accordingly, a policy rule, which is optimal under a single specification, may turn out to perform quite poorly if that model does not capture properly the 'true' macroeconomic dynamics. Then, the observation of smooth policy rates can simply reflect the choice of a policy rule that would perform reasonably well over various alternative policy scenarios.

A general strategy to take model uncertainty into account is to calculate a global optimal policy as some combination of the policy rules derived separately for each of the relevant specifications (see Stock, 1999). It is worthy to note that the *robust* rule we are interested in differs in scope from the one derived with robust control techniques. Indeed, here robustness has to be understood as a form of hedging against potential misspecifications of the macroeconomic dynamics rather than as a way of guarding against worst case scenarios. To this end, we follow the *thick* modeling proposed by Granger and Jeon (2001) to pool into a single policy rule a large number of specifications in a given class of nested models. In particular, we first let policy makers implement, at each point in time, some average of the optimal rates for each of the relevant specifications. Then, we identify among a large number of targeting policies the set of preference parameters that makes such a *robust* rule matching the data.

Our results shed new lights as well as confirm conventional wisdoms on the conduct of US monetary policy in the last decade. First, potential misspecifications of the macroeconomic dynamics is an important concern of the Fed such as to explain alone most of the observed inertial behavior of policy rates. Second, any identification method that did neglect model uncertainty would deliver a set of policy preferences that cannot be readily interpreted. Third, the stabilization of output over the cycle has not been a final concern of US monetary authorities whereas the stabilization of inflation has been a superior goal.

The paper is organized as follows. Section 2 sets up the model and presents the relative estimates. Section 3 identifies the preference parameters for the Greenspan's tenure and defines the *interest rate smoothing puzzle* from the comparison between our results and those obtained in several recent studies. The *thick* modeling approach to model uncertainty is introduced in section 4 and then it is used in the following section to re-identify the Fed policy preferences. The last section concludes while the appendix provides a guideline to solve numerically the optimal control problem.

3.1 A small empirical model of the US economy

The central bank faces a dynamic optimal control problem whose solution describes its policy actions. These are the optimal response of monetary authorities to the evolution of the economy as captured by the relations among the state variables. We describe such a dynamics by means of a simple closed economy-two equation framework made up of an aggregate supply and an aggregate demand, which actually represent the constraints of the policy makers' optimization problem.

3.1.1 The structure of the economy

The empirical evidence from VAR studies shows that monetary policy affects the economy at different lags (see Christiano, Eichenbaum and Evans, 1998, and Bernanke and Mihov, 1998). Furthermore, if the central bank faces an intertemporal optimization problem, then forecasting the behavior of the state variables becomes crucial to set policy rates as the optimal response to the developments in the economy. It follows that for the purpose of monetary policy making, which relies on forecasting methods, a backward-looking model may be a suitable characterization of the macroeconomic dynamics (see Fuhrer, 1997).

Accordingly, we let the structure of the economy evolve as follows:

$$\pi_{t+1} = \alpha_1 \pi_t + \alpha_2 \pi_{t-1} + \alpha_3 \pi_{t-2} + \alpha_4 \pi_{t-3} + \alpha_5 y_t + \varepsilon_{t+1} \tag{22}$$

$$y_{t+1} = \beta_1 y_t + \beta_2 y_{t-1} + \beta_3 \left(\bar{\imath}_t - \bar{\pi}_t\right) + u_{t+1} \tag{23}$$

where π_t is the quarterly inflation in the GDP chain-weighted price index, p_t, calculated at annual rate,. that is $4\left(p_t - p_{t-1}\right)$, and $\bar{\pi}_t$ is four-quarter inflation constructed as $\frac{1}{4}\sum_{j=0}^{3}\pi_{t-j}$. The quarterly average federal funds rate, i_t, is expressed in percent per year whereas the four quarter average federal funds rate, $\bar{\imath}_t$, is computed as $\frac{1}{4}\sum_{j=0}^{3} i_{t-j}$; Supply and demand i.i.d. shocks are denoted by ε_t and u_t respectively. All variables are demeaned. All variables but the funds rate are in logs and rescaled upward on a 100 point basis such that the output gap, say, is $y_t = 100 * (\log(Q_t) - \log(Q_t^*))$ where Q_t and Q_t^* are respectively actual and potential GDP, both in levels. Therefore, no constants appear in the equations.

On the one hand, the aggregate supply equation in (22), AS henceforth, captures the inflation dynamics by relating inflation to its lagged values and to current and lagged output gaps. On the other hand, the aggregate demand equation in (23), AD henceforth, explicitly models the monetary transmission mechanism by relating output gap to its lagged values and most importantly to past real interest rate (see Rudebusch and Svensson, 1999).

This empirical model of inflation and output, although parsimonious, embodies the minimal set of variables one may want to include for the analysis of monetary policy (see, for instance, Christiano, Eichenbaum and Evans, 1998), and, as argued in Rudebusch and Svensson (1999), it appears to be broadly in line with the view that policy makers hold about the dynamics of the economy (see the report of the Bank for International Settlements for 11 central bank models, 1995). Moreover, monetary policy affects (through the instrument i_t) aggregate demand with one lag and aggregate supply with two lags, in the spirit of the specifications in Ball (1999) and Svensson (1997). Finally, such a dynamics can be interpreted either as a structural relation or as a reduced-form restricted VAR with impulse responses that are consistent with those of the FRB-US model.

The AD-AS system is backward-looking and therefore it is subject to the Lucas critique (1976). It follows that the selection of an inappropriate sample may undermine the stability of the behavioral parameters of the economy, which is an important condition for drawing inference. For instance, Muscatelli and Trecroci (2001) show evidence that while the response of output to interest rate shocks has not significantly changed, the short-run correlation between output and inflation has shifted during the last two decades. To the extent that this can be ascribed to the productivity growth that has characterized the US economy since the late 80s, focusing on the sample 1987:3 - 2001:1, which corresponds to the tenure of Alan Greenspan as Fed chairman, it turns out to be beneficial to limit parameter variation. Indeed, one may argue that this period has been marked not only by an increasing macroeconomic stability and a lower inflation but also by the expectations of some form of inflation targeting (see Bernanke and Mihov, 1998), thereby reducing the significance of the Lucas critique.

We estimate individually equations (22) and (23) by OLS. The potential output is obtained from the Congressional Budget Office whereas all other data are taken from the web-site of the Federal Reserve Bank of St. Louis. In particular, we collect monthly time-series for the funds rate, quarterly data for the GDP chain-weighted 1996 commodity price index and quarterly data for the potential output. All series are seasonally adjusted. We then convert monthly data in quarterly data by taking end-of-quarter observations. Lastly, we de-mean all variables.

The estimates are as follows, standard errors in parenthesis:

$$\pi_{t+1} = \underset{(0.133)}{0.282}\pi_t - \underset{(0.134)}{0.025}\pi_{t-1} + \underset{(0.134)}{0.292}\pi_{t-2} + \underset{(0.136)}{0.385}\pi_{t-3} + \underset{(0.054)}{0.141}y_t + \hat{\varepsilon}_{t+1} \tag{24}$$

$$y_{t+1} = \underset{(0.136)}{1.229}y_t - \underset{(0.149)}{0.244}y_{t-1} - \underset{(0.078)}{0.073}\left(\bar{\imath}_t - \bar{\pi}_t\right) + \hat{u}_{t+1} \tag{25}$$

The system displays a reasonably good empirical fit with an Adjusted R^2 equal to 0.58 for the AS and 0.93 for the AD.[11] All estimates have the expected sign but the second lag of inflation in the AS, although it has not explanatory power. Furthermore, the coefficient for the real interest rate is not statistically significant. While undesirable, this result confirms the evidence from several studies for the US and the UK over recent samples (see for instance Muscatelli and Trecroci, 2001, and Neiss and Nelson, 2001). Finally, although these estimates suggest a minor initial role for monetary policy, the impact of the lagged values of the output gap in the AD is large implying that the response of aggregate demand to policy rates is much greater in the long-run.

---

[11] Moreover, the cross-correlation of the errors is 0.137, implying that the parameter estimates are not affected by the estimation method. Lastly, the Andrews' test (1993) cannot reject the null of stability for both equations.

3.1.2 The loss function and the optimal monetary policy

We assume that monetary authorities operate according to a *targeting rule* as defined in Svensson (1999). This corresponds to set the instrument rate so as to bring at each point in time the target variables in line with the targets by penalizing any future deviation of the former from the latter. Following Rudebusch and Svensson (1999), we let the central bank pursue the stabilization of the four-quarter inflation around the inflation target, the stabilization of the output around its potential value and potentially the smoothing of interest rate. The inflation target is assumed to be constant over time and it is normalized to zero because all variables are demeaned.[12] Then, policy rates are set to minimize the following objective function:

$$Var\left[\bar{\pi}_t\right] + \lambda Var\left[y_t\right] + \mu Var\left[\Delta i_t\right] \tag{26}$$

The quarterly average short-term interest rate, i_t, is regarded as the instrument under policy makers' control whereas Δi_t stands for its first difference. The parameters λ and μ represent the central bank's policy preferences towards output stabilization and interest rate smoothing respectively and unlike in Rudebusch and Svensson (1999), who set them exogenously, they will be determined within the model. The coefficient on inflation stabilization is normalized to one such that λ and μ are expressed in relative terms. Finally, we constrain both parameters to be non negative meaning that the central bank values both any deviation of output from its potential and any jump in interest rates as a *bad*.

On the positive side, the specification in (26) is empirically attractive since, unlike alternative monetary models as the FRB-US, it is able to predict an interest rate path that exhibits the kind of inertia observed in the data. On the negative side, the desire

---

[12] As argued in Dennis (2000), demeaning all variables does not affect the derivation of policy makers' preferences. Furthermore, our analysis is meant to identify the central bank parameters over the target variables rather than to estimate the targets per sè. A number of papers cover the issue, including Judd and Rudebusch (1998), Sack (2000), Favero and Rovelli (2001), and Dennis (2001).

for smoothing policy rates has little theoretical justification beyond the optimal delegation argument according to which the appointment of a central banker who pursues an alternative objective relative to the true social one may be welfare improving (see Woodford, 2002, Ch. 7).[13] However, it can be argued that high variability and frequent reversals in interest rate movements may lead to financial instability (see Goodfriend, 1991) as well as they may be interpreted by the private sector as an admission of earlier policy mistakes (see Lowe and Ellis, 1997), thereby being undesirable.

The optimal control problem described in (22), (23) and (26) has a convenient state space representation that is characterized by a quadratic objective and a linear transition law. This specification leads to the *stochastic optimal linear regulator problem* according to which the decision rule for interest rates is a linear function of the state variable vector:

$$X_t' = \left[\begin{array}{ccccccccc} \pi_t & \pi_{t-1} & \pi_{t-2} & \pi_{t-3} & y_t & y_{t-1} & i_{t-1} & i_{t-2} & i_{t-3} \end{array} \right] \tag{27}$$

In particular, the central bank minimizes the loss (26) subject to the dynamic constraints (22) and (23). In so doing, it determines an optimal reaction function that can be expressed in the compact form[14]:

$$i_t = f X_t \tag{28}$$

The coefficients in the vector f represent some convolution of the central bank's preferences, λ and μ, and the behavioral parameters of the economy, αs and βs, such that for any given distribution of weights in (26) there exists a different optimal f in (28).

Then, we make the model consistent with our implementation by the timing assumption that the Fed sets policy rates after the realization of the state variables,

---

[13] Alternatively, monetary authority may wish to stabilize the level, rather than the change, of policy rates. Then, the presence of transaction frictions and/or a zero nominal interest-rate lower bound result in an utility-based loss function with an interest rate term which enhances social welfare (see Woodford, 2002, Ch. 6)

[14] The appendix provides a full derivation of the feedback rule that solves the stochastic optimal linear regulator problem.

which occurs at the beginning of the period. Hence, we estimate by OLS the stochastic version of the optimal rule derived in (28). The estimates yield the following results:

$$\begin{aligned} i_t &= \underset{(0.07)}{0.212}\pi_t + \underset{(0.08)}{0.043}\pi_{t-1} + \underset{(0.08)}{0.151}\pi_{t-2} - \underset{(0.09)}{0.177}\pi_{t-3} + \underset{(0.10)}{0.346}y_t + \\ & -\underset{(0.11)}{0.265}y_{t-1} + \underset{(0.14)}{1.259}i_{t-1} - \underset{(0.20)}{0.398}i_{t-2} - \underset{(0.12)}{0.008}i_{t-3} + \hat{v}_t \end{aligned} \tag{29}$$

with an Adjusted R^2 of 0.96.[15] The significant parameters show that the monetary authorities operate in a gradual manner by changing the funds rates in response to both inflation and output gaps. In particular, the first lag of the policy rate implies that the Fed tends to move its instrument in a particular direction over sustained periods, while the second lag confirms the potential for few reversals (see Rudebusch, 1995, and Goodhart, 1997). Finally, the coefficients on the interest rate lags sum up to 0.85 consistently with much of the literature on partial adjustment policy rules. This suggests that the observed policy inertia is greater than systematic responses to output and inflation fluctuations would imply.

3.2 The Fed policy preferences with no model uncertainty

The design of monetary policy depends upon the targeting strategy adopted by the central bank. This strategy describes a set of policy preferences, which are actually the structural parameters that characterize the aversion of monetary authorities towards inflation, output and potentially interest rate volatility. Then, a simple way to recover these preferences is to assume that policy makers are acting optimally and, as a kind of revelation principle, to extract the relevant information from the observed policy decisions. The control problem described above shows that the reaction function estimates can be interpreted as convolutions of the behavioral parameters of the economy and

---

[15] McCallum and Nelson (1999) argue that in operational policy making the central bank does not observe (and respond to) the current state of the economy. Using four lags of funds rate, GDP inflation and CBO output gap as instruments does not change significantly neither the point estimates nor the standard errors of the feedback coefficients.

those describing the central bank's preferences and therefore they are natural candidates for the purpose at hand.[16] Accordingly, given the point estimates in (24) and (25), we calibrate the preference parameters $[\lambda, \mu]$ such as to minimize the distance between the optimal policy and the fitted path of interest rates in (29), where the distance is measured by the sum of squared deviations over time.[17] The optimal policy describes the path that the funds rates would have followed if the Fed had historically implemented the optimal rule and therefore, given the actual values of state variables at the beginning of the sample, it is derived by substituting, period by period, the simulated dynamics of the X into the reaction function (28). Our identification method applied to the sample 1987:3 - 2001:1, which corresponds to the Greenspan chairmanship, returns values of $\lambda = 1.00$ and $\mu = 8.00$ for the preferences on output stabilization and interest rate smoothing respectively. One may be tempted to conclude that while output and inflation stabilization have received an equal concern, interest rate smoothing has been the major objective of the Fed. However, we show below that these results can be highly misleading in that they miss an important feature of actual monetary policy making.

At this point, it is useful to relate our results to several recent studies since there exists interesting differences and similarities. Favero and Rovelli (2002) identify central bank's preferences by estimating via GMM the Euler equations for the solution of alternative specifications of the optimization problem. Cecchetti and Ehrmann (1999) capture the dynamics of the economy in a VAR framework and then recover policy makers' preferences from the estimates of the output-inflation variability frontier and those obtained via VAR. Dennis (2001) and Ozlale (2001) use respectively a full information approach and the Kalman filtering to jointly estimate with maximum likelihood

---

[16] Moreover, our optimal control problem satisfies the three necessary and sufficient conditions derived in Dennis (2000) to identify central bank policy preferences .

[17] By defining our measure of distance upon fitted rather than actual rates we restrict our attention to the systematic component of policy rate behaviour, that is, to the component we can explain within an optimal control framework. Moreover, our results do not change significantly when actual rates enter the calibration because of the good empirical fit of the feedback estimates.

the structural model of the economy and the loss function. These studies but the ones by Cecchetti and Ehrmann (1999) are built upon a common empirical model of inflation and output, namely the one by Rudebusch and Svensson (1999), and therefore their findings turn out to be directly comparable to ours. Table 1 brings together our revealed preferences and the estimates from the different contributions. The reported values refer to the Greenspan's tenure, although Favero and Rovelli (2002) do not distinguish between the Volcker's and the Greenspan's chairmanship.[18] In particular, Panel A shows the first two moments of the fitted policy rates whereas Panel B displays in columns the Fed policy preferences, the first two moments of the optimal paths and the average distance between optimal and fitted rates. Figure 1 plots the optimal and the fitted path of policy rates for the four studies.

The first two lines of Panel B in Table 1 refer to the present work and the one by Dennis (2001).[19] On the one hand, these sets of policy preferences predict a path for policy rates capable to replicate the kind of smoothness observed in the data (see the top panels of Figure 1). Indeed, the first two moments are broadly consistent in both cases with those of the fitted path in Panel A and the average distance, which is computed on squared values, is fairly low. On the other hand, they rely upon extremely large parameters for interest rate smoothing which cannot be easily motivated within the optimal monetary policy literature.[20]

By contrast, the last two lines of Table 1, which refer to the works by Favero and

---

[18] Understanding whether the two periods may be described by a single set of policy preferences is beyond the scope of this paper. However, to the extent that no monetary regime shifts have occured in the post-Volcker period (see Clarida, Galì and Gertler, 2000), the preference parameters in Favero and Rovelli (2002) can be taken as a rough approximation of those in the restricted sample for Alan Greenspan only. As we are interested only in a qualitative comparison between our optimal policy rule and those from other studies, we consider such an approximation only as a minor in the interpretation of the results.

[19] We thank Richard Dennis for having kindly offered the FIML estimates for the Greenspan's period.

[20] For instance, the utility based loss function in Woodford (2002, Ch. 6 and 7), albeit derived in a different class of models, implies a theoretical value of μ no greater than 0.28, which is based on stuctural estimates for the US economy.

Rovelli (2002) and Ozlale (2001), return more plausible weights for the inertial coefficient in the loss function. However, the bottom panels of Figure 1 show that this can be done only at the cost of an optimal policy rule that is so volatile as to contradict the evidence on the funds rates.

The results at this stage seem to call for a sort of *interest rate-smoothing puzzle.* A trade-off between an inertial behavior of policy rates and a plausible value for the relative preference parameter seems to emerge, thereby suggesting that the source of interest rate smoothing has to be found elsewhere.

The structure of the economy proposed by Rudebusch and Svensson (1999), while empirically attractive, is indeed very simple and the omission of any relevant variable may turn out to be an issue for the results obtained so far. Moreover, as discussed in the introduction, the lack of knowledge about the 'true' model of the economy may lead policy makers to consider various alternative policy scenarios, each one corresponding to a different specification of the underlying macroeconomic dynamics. We explore such an alternative in the next section to assess the potential of model uncertainty to account for the observed interest rate smoothing.

3.3 Model uncertainty

A common observation across central banks is that interest rates are moved in a more moderate fashion than certain equivalent optimal monetary policies predict. The difficulty of standard models to rationalize policy inertia has led to incorporate various forms of model uncertainty into the policy makers' optimization problem. In practice, monetary authorities know far less about the dynamics of the economy than simple policy experiments presume and model parameters are likely to be better viewed as random. In particular, suppose that monetary authorities know the distribution of parameters but not the realization; then, uncertainty can be introduced at different levels. A Brainard-style multiplicative uncertainty (1967) considers parameter distributions that are centered around the estimates of a specific model. This means that policy makers know the parameter first moments on an ex ante basis, although they do not know the values that realize in any given quarter. Rudebusch (2001), Estrella and Mishkin (1999), and Peersman and Smets (1999) find that parsimonious structural models and simple policy rules predict only negligible attenuations of policy action in the context of such an uncertainty. By contrast, Sack (2000), Salmon and Martin (1999) and Söderström (1999) show using unrestricted VARs and unrestricted policy rules that the response of monetary authorities may result quantitatively more moderate, although they conclude that multiplicative parameter uncertainty alone is not enough to replicate the kind of smoothness observed in the data.

Another way to think of model uncertainty is to regard also the parameter mean as unknown. In fact, if policy makers fear that a small structural model is misspecified, they would have no reason to believe that the 'true' parameters coincide, even on average, with the least square estimates. A valuable robustness check is then to vary the values of some key model parameter to understand whether this is the relevant form of uncertainty that central banks face. Rudebusch (2001) shows that the slope coefficients on inflation and output gap are indeed crucial as the perturbation of each of them, everything equals, results in a significant, but not exhaustive, attenuation of the policy stance.

These results altogether are very promising in that they point towards model uncertainty, in a *broad* sense, as the relevant source of the observed policy gradualism. Moreover, they suggest that the policy preference reported above may be 'misleading' as no identification method takes such an uncertainty into account and only the point estimates of the model parameters enter the analyses. By contrast, this section incorporates model specification uncertainty into the calibration of the Fed policy preferences.

In so doing, we attempt to solve for the *interest rate smoothing puzzle* by assessing the potential of a *broad* type of uncertainty for explaining the inertial behavior of policy rates.

Our approach departures from previous studies along three lines. First, we regard the point estimates of our benchmark model only as one set of possible realizations. In other words, we allow the average value of the distributions to be different from the estimated parameters. Moreover, rather than assuming that these distributions are known ex-ante, we let them be shaped ex-post by the point estimates obtained for each of the possible models. Lastly, in addition to the kind of slope coefficient uncertainty in Rudebusch (2001), we also allow for simultaneous perturbations of all parameters as potentially omitted variables are likely to affect each of the point estimates in the model.

In practise, we follow Granger and Jeon (2001) and we label this approach to model uncertainty *thick modeling.* We keep all close specifications according to some statistical criterion, find their outputs that relate to the design of optimal monetary policy and pool these values. The label 'thick', as opposed to 'thin', reflects the fact that if one estimates and plots each model-specification she will get a 'thick' representation of the optimal monetary policy, that is, a curve whose width is made up of as many 'thin' curves as the number of specifications that survive the trimming of the outliers.

Before discussing our 'thick' strategy, we consider worthwhile to describe how model uncertainty has been traditionally approached.

3.3.1 Traditional approaches

The robustness of monetary policy to model uncertainty has been the focus of a number of recent empirical studies. The goal has been to assess the performance of optimal rules moving from the model in which they are derived to a set of alternative specifications as well as to establish the efficiency of simple policy rules (see Taylor, 1999). For example, McCallum (1998) shows that monetary-based instrument rules overperform optimal ones over a range of possible macroeconomic dynamics. Moreover, simple partial adjustment policy mechanisms and simple forecast-based instrument rules responding to an inflation horizon no longer than one year are found to efficiently stabilize inflation and output in a variety of forward-looking models (see Levine, Wieland and Williams,1999 and 2001). Essentially, these rules set the change in the funds rate rather than the level as the optimal value of the lagged policy rate coefficient is close to one. The intuition is that the central bank, which has established a reputation of conducting monetary policy in a gradual manner, can achieve its goals while maintaining a low level of interest rate volatility through the expectations of policy inertia (see also Goodfriend, 1991 and Woodford, 2002, Ch. 7).

An alternative approach to solve for model uncertainty is provided by the techniques of robust control (see Hansen and Sargent, 2001, chapters 6 and 8). This method specifies a risk function and a minimax criterion that serve to form a non-parametric set of perturbations around the policy makers' model. The latter is assumed to be an approximation that belongs to a potentially time varying and state dependent bounded neighborhood of the 'true' model of the economy. Then, given the least favorable scenario, that is roughly speaking the maximum value that the loss function can take in that neighborhood, the robust optimal rule is chosen so as to minimize the maximum value function. Interestingly, Stock (1999), Onatski and Stock (2002), and Tetlow and von zur Muehlen (2001) show that model uncertainty may call for a more activist policy stance, although the worst possible models for the kind of historical Fed policy rule may not describe plausible structures of the economy (see Onatski, 2000). The intuition for this result comes from the fact that the central bank plays a game against a malevolent nature in which only worst case scenarios matter for policy making. This implies that an aggressive rule may be the optimal response of monetary authorities to large departures of inflation and output from the target values.

3.3.2 A novel approach: 'thick modeling'

The standard practice of econometric modelling is to choose among a set of relevant specifications the best according to some model selection criterion like *adjusted* R^2, *Akaike* or *Schwarz*, discarding any information in the alternative specifications. In practical policy making, however, it is not clear that this may be a good strategy and policy makers, who are uncertain about the future state of the economy, may find retaining and combining all information in a number of close specifications a superior strategy. The reason for that mirrors the results in the literature of optimal forecasting (and portfolio allocation) which demonstrate that the combination of forecasts (assets) is often a better procedure than using the best single forecast (asset). Then, *mutatis mutandis*, the monetary authority may prefer to consider the range of a wide number of optimal monetary policies, each one corresponding to the solution of the control problem associated to a different structure of the economy, rather than to come up with a single policy rule which is optimal only within the model specification in which it has been derived. In so doing, they may end up with as many policy prescriptions as the number of relevant macroeconomic scenarios. To the extent that the latter differ in the lag specification of the monetary transmission mechanism and that policy makers have no strong *a priori* on the future state of the economy, the *thick* modelling of combining those prescriptions comes as a simple strategy for the design of a global optimal policy without requiring any restrictive decision about what model will provide the best description of the economy.

In practice, we specify a class of nested models for the structure of the economy and propose some *a priori* criterion to pool into a single robust *thick* policy rule the information that relate to the design of monetary policy. To this end, we estimate by OLS the dynamics generated by the relevant combinations of a base set of eight regressors for the AS and nine for the AD whose richest specification takes the following form:

$$\begin{aligned}\pi_{t+1} = \alpha_1\pi_t + \alpha_2\pi_{t-1} + \alpha_3\pi_{t-2} + \alpha_4\pi_{t-3} + \\ \alpha_5 y_t + \alpha_6 y_{t-1} + \alpha_7 y_{t-2} + \alpha_8 y_{t-3} + \xi_{t+1}\end{aligned} \tag{30}$$

$$\begin{aligned}y_{t+1} = \beta_1 y_t + \beta_2 y_{t-1} + \beta_3 y_{t-2} + \beta_4 y_{t-3} + \beta_5\pi_t + \\ \beta_6\pi_{t-1} + \beta_7\pi_{t-2} + \beta_8\pi_{t-3} + \beta_9\left(\bar{\imath}_t - \bar{\pi}_t\right) + \eta_{t+1}\end{aligned} \tag{31}$$

The selection of the relevant models is based on both empirical and theoretical arguments. First, we keep fixed across specifications the first lag of inflation and output gap in the AS and AD respectively. In so doing, we end up with those models displaying a fairly good empirical fit. Moreover, we discard the specifications that do not allow monetary policy to have a direct impact on the economy through both equations. In particular, we take the real interest rate, $\bar{\imath}_t - \bar{\pi}_t$, as a further fixed regressor and we constraint the AS to be dependent from, at least, one of the lagged values of the output gap. The latter amounts to cut off approximately the five percent of the 2^7x2^7 models specified in this class. Finally, we derive the optimal policy rules for each of the retained AD-AS specifications and we let policy makers implement, at each point in time, the average of the optimal rates associated to those specifications.

A number of alternative weighting schemes may be appropriated for computing the average optimal policy. Instead of using a simple statistical pooling, Granger and Jeon (2001) argues that a simple averaging may serve for the purpose at hand, corresponding to what in the literature is usually referred to as a non-informative prior with equal weights given to different monetary policies. An alternative somewhat in the spirit of Bayesian econometrics is to weight the OLS estimates across models by some statistical criterion corrected for the degrees of freedom. Doppelhofer, Miller and Sala-i-Martin

(2000) propose a weighting criterion analogous to the Schwarz in the context of the so-called Bayesian averaging of classical estimates (BACE), which has the advantage over the Bayesian model averaging of not requiring any specification of prior distributions for the model parameters.

These alternative weighting schemes describe the robust policy rules that we use in the next section to evaluate the ability of model uncertainty to account for the observed interest rate smoothing. Our *thick* strategy is in the spirit of Favero and Milani (2001), although we take three important departures. First, we analyze a different sample according to the reasoning that policy preferences are Chairman-specific. Second, we endogenously determine these preferences rather than simply imposing them. Lastly, we evaluate the robustness of our results to different weighting schemes for averaging the optimal policies obtained under the alternative policy scenarios.

3.4 The Fed policy preferences under model uncertainty

In this section, we use our identification method to recover the preference parameters for the Greenspan's tenure in the presence of model uncertainty. In order to gauge the merits of the robust *thick* policy rule we compare our results with those obtained under a multiplicative parameter uncertainty which a number of researchers have advocated as an important, although not exhaustive, source of policy attenuation (see Sack, 2000, Sack and Wieland, 2000, and Rudebusch, 2001 among others).

It is worthy to note that in contrast to the analysis in section 3, which considers a single specification of the economy and thus a single optimal rule, the calibration is based here on the distance between fitted and *thick* policy rates, where the latter are computed as some average of the optimal rules for each of the relevant models. In so doing, we incorporate model uncertainty into the identification of policy preferences. In other words, we investigate whether the Fed cares about model misspecification by assessing the ability of a *robust* rule to match the data without resorting to implausibly high values for the interest rate smoothing parameter.

3.4.1 The robust thick policy rule

The third row of table 2 reports some descriptive statistics of the optimal rule under model uncertainty as well as the corresponding calibrated policy parameters. The revealed preferences for the Greenspan's chairmanship write now $\lambda = 0.00$ and $\mu = 0.11$ while the first two moments of the associated optimal path are consistent with the historical policy (first row). Moreover, the average distance is still fairly low and the standard deviation of the interest rate changes, which actually defines interest rate smoothing, remains virtually identical moving from the historical rule to the robust *thick* rule. While the statistics and the following figures on model uncertainty refer to the simple average case, the picture does not change, both qualitatively and quantitatively, weighting each optimal policy with the relative *adjusted* R^2, *Akaike* and *Schwarz* criterion respectively. In the light of our trimming strategy, this result does not come as a surprise since the closer are the retained specifications the more the weighted average tends to the simple average, that is the grater is the likelihood that similar weights are attached to each specification.

Figure 2 compares the two optimal paths associated to the preferences $\lambda = 0.00$ and $\mu = 0.11$ in the absence and under model uncertainty respectively. The robust *thick* policy rule effectively describes the main features of funds rate movements throughout the sample, although there are some differences in magnitude. While this suggests that other source of uncertainty such as measurement errors for inflation and output gap may also be relevant, we find that by considering model misspecifications most of the *interest rate smoothing puzzle* seems to vanish, as the relative preference parameter take now only a modest value. Model uncertainty is eventually crucial because whenever neglected the optimal policy rule looses its ability to match the data. Hence, any identification method that did not take this form of uncertainty into account would miss an important part of the story, thereby delivering a set of policy preferences that cannot be sensibly interpreted.

The revealed policy preferences computed under model uncertainty show that the conduct of monetary policy in the US is successfully described by a *strict inflation targeting* as defined by Svensson (1999), and Rudebusch and Svensson (1999). According to it, the stabilization of output around potential has not been a final concern of the Federal Reserve (i.e. $\lambda = 0.00$). However, we do not mean that the output gap has been unimportant in policy actions. Indeed, as argued by Favero and Rovelli (2002) and Dennis (2001), it may well be that the output gap has been regarded as a leading indicator for future inflation rather than as a goal variable per sè (i.e. as an argument in the reaction function rather than in the loss). An alternative, in the spirit of the evidence in Smets (1999), Estrella and Mishkin (1999), and Wieland (1998) on output gap uncertainty, is that monetary authorities have placed less weight on the most poorly measured target, or yet, that the marked productivity growth of the 90s has drastically reduced any concern towards output stabilization.

3.4.2 Model uncertainty vs. parameter uncertainty

The result that uncertainty makes smoother an otherwise volatile path of policy rates does not come as new in the literature and a number of empirical studies have recently shown that multiplicative parameter uncertainty limits the responsiveness of the interest rate (see Sack, 2000 and the references therein). A relevant question at this point is the extent to which parameter uncertainty would be capable alone to replicate the observed path or rather there exists room for other forms of uncertainty. To this end, we bring together in the last two rows of table 2 some descriptive statistics for the robust policy rules obtained under model and parameter uncertainty respectively. We take as given the revealed policy preferences $\lambda = 0.00$ and $\mu = 0.11$, which assigns a very limited role to an interest rate smoothing goal, so that the performance of the robust rules can be readily compared. The computational difference between the two robust rules stems from the distribution of the AS-AD coefficients which only under parameter uncertainty are centered around our estimates of the Rudebusch and Svensson (1999) model and shaped by the relative estimated standard errors. By contrast, the robust *thick* approach does not impose any mean value to the parameter distributions whose support reflects a model specification uncertainty rather than the classical estimation uncertainty due to sampling.

The last row of table 2 shows that multiplicative parameter uncertainty attenuates the policy response of monetary authorities such that the relative robust descriptive statistics come closer to the data than the single specification counterparts. Nevertheless, the robust optimal policy seems to reduce but not to close the gap with the observed monetary policy confirming the conclusions in Sack (2000) and Rudebusch (2001). In addition, taking model uncertainty into account makes the robust *thick* policy rule more successful at describing the policy rate dynamics than the parameter uncertainty robust rule. This can be seen not only from the first two moments and the average distances but also, more importantly, from the standard deviations of the interest rate changes. Consistent with these findings, Figure 3 shows that the behavior of policy rates is considerably smoother under model uncertainty than under parameter uncertainty as the robust *thick* policy rule shows more limited deviations from the historical rule.

We interpret these results as the evidence that model misspecification has been an important concern of the Fed such that its ability to limit the responsiveness of the fed funds rate goes beyond the ability of a multiplicative parameter uncertainty.[21]

---

[21] It should be noticed that we have modelled parameter uncertainty as the perturbation of the slope coefficient of inflation and the interest rate sensitivity on output only. While varying all parameters produces only limited changes, an alternative would be to consider a richer macroeconomics dynamics as the one in a VAR specification of the economy. However, Sack (2000) shows that even involving very

3.5 Conclusions

Actual policy rates appear to be smoother than optimal monetary policies predict. An obvious way to reconcile the historical evidence with an optimizing central bank behavior is to model the aversion to interest rate fluctuations as an independent argument in the central bank's loss function. However, the relative parameter should be imposed at values so high as they cannot be easily motivated by the theory, thereby making this choice alone unsatisfactory.

This paper contributes to the literature of optimal monetary policy by presenting a novel method to solve for a relevant form of uncertainty in practical policy making, namely uncertainty about the structure of the economy. While there may well be also other rationales such as data uncertainty or a minor goal to avoid interest rate variability, it is shown that the concern for potential misspecifications of the macroeconomic dynamics creates incentives for monetary authorities to move policy rates in a gradual manner. Indeed, a thick approach to model uncertainty appears to solve most of the observed *interest rate smoothing puzzle* as the preference calibration based on a robust policy rule returns values which are more readily interpretable. Moreover, the preference parameters show that the Greenspan's tenure as Fed chairman is effectively described by a *strict inflation targeting* policy according to which the stabilization of inflation around its target has been the only concern of monetary authorities.

We take these results as a promising deal for future research and the calibration exercise we propose proves these potentialities. Intriguing identification strategies for the preference parameters have returned unattractive results in that they display either implausible values for the inertial coefficient or extremely volatile paths for the policy rates whenever model uncertainty is neglected. By contrast, our revealed preferences move to sensible values when the calibration incorporates a wide number of possible

---

persistent interest rate movements, the optimal policy derived within a VAR dynamics is still more aggressive than the observed policy.

specifications. This seems to suggest that most of the observed policy inertia can be better interpreted as a consequence of monetary policy making under uncertainty rather than as an objective in itself and that omitted model uncertainty may lead to the spurious finding of an independent goal for interest rate smoothing.

Furthermore, our robust *thick* modeling can be extended to alternative formulations of the inflation dynamics and the output gap dynamics in order to evaluate the empirical relevance of model uncertainty within a class of non-nested specifications. Lansing and Trehan (2001), for instance, show that by introducing some degree of forward-looking behavior in output, the responses to inflation and output gap recommended by an optimizing Taylor rule are less pronounced. In particular, they show that private sector expectations may be an important channel through which monetary policy can be effectively conducted by means of small interest rate changes (see also Levin, Wieland and Williams, 1999, and Sack and Wieland, 2000). However, Söderlind, Söderström and Vredin (2002), who calibrate the preferences of the Fed within a New-Keynesian model of output and inflation, still find a large value for the policy parameter on interest rate smoothing. This suggests that model uncertainty about the relevant macroeconomic dynamics may turn out to be an issue also in such a framework and therefore further work can be usefully done along these lines.

Appendix: the optimal control problem

For a discount factor δ, $0 < \delta < 1$, the central bank faces an intertemporal optimization problem of the form:

$$E_t \sum_{\tau=0}^{\infty} \delta^{\tau} LOSS_{t+\tau} \tag{32}$$

according to which it minimizes the expected discounted sum of future loss values. In particular, the objective function reads in each period:

$$LOSS_t = \bar{\pi}_t^2 + \lambda y_t^2 + \mu \left(i_t - i_{t-1} \right)^2 \tag{33}$$

The loss function is quadratic in the deviations of output and inflation from their target values and embodies an additional term that is meant to penalize for an excessive volatility of the policy instrument, i_t. The parameters λ and μ represent the relative policy preferences of the central bank towards output stabilization and interest rate smoothing respectively. The inflation stabilization weight in the objective function is normalized to one.

When the discount factor, δ, approaches unity, the intertemporal loss function in (32) approaches the unconditional mean of the period loss function:

$$E\left[LOSS_t\right] = Var\left[\bar{\pi}_t\right] + \lambda Var\left[y_t\right] + \mu Var\left[\Delta i_t\right] \tag{34}$$

The constraints of the optimization problem describe the structure of the economy, and they are specified by the AD-AS system in (22) and (23). This has a convenient state-space representation of the form:

$$X_{t+1} = AX_t + Bi_t + \eta_{t+1} \tag{35}$$

where the elements of (35) are given by:

$$X_t' = \left[\begin{array}{ccccccccc} \pi_t & \pi_{t-1} & \pi_{t-2} & \pi_{t-3} & y_t & y_{t-1} & i_{t-1} & i_{t-2} & i_{t-3} \end{array}\right] \tag{36}$$

$$A = \left[\begin{array}{ccccccccc} \alpha_1 & \alpha_2 & \alpha_3 & \alpha_4 & \alpha_5 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{-\beta_3}{4} & \frac{-\beta_3}{4} & \frac{-\beta_3}{4} & \frac{-\beta_3}{4} & \beta_1 & \beta_2 & \frac{\beta_3}{4} & \frac{\beta_3}{4} & \frac{\beta_3}{4} \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{array}\right], \quad B = \left[\begin{array}{c} 0 \\ 0 \\ 0 \\ 0 \\ \frac{\beta_3}{4} \\ 0 \\ 1 \\ 0 \\ 0 \end{array}\right] \tag{37}$$

$$\eta_t' = \left[\begin{array}{ccccccccc} \varepsilon_t & 0 & 0 & 0 & u_t & 0 & 0 & 0 & 0 \end{array}\right] \tag{38}$$

X_{t+1} is the $9x1$ vector of state variables, i_t is the policy control (i.e. the federal funds rate) and η_{t+1} is a $9x1$ vector of supply and demand i.i.d. normally distributed shocks with mean vector zero and covariance matrix $E\eta_t\eta_t' = \Omega$. Lastly, A and B are the matrices of behavioral parameters.

The loss function in (33) can be represented in a more compact form by defining the $3x1$ vector Y_t of goal variables. This vector reads:

$$Y_t = CX_t + Di_t \tag{39}$$

where the elements of (39) are given by:

$$Y_t = \begin{bmatrix} \bar{\pi}_t \\ y_t \\ i_i - i_{t-1} \end{bmatrix}, \quad C = \begin{bmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \end{bmatrix}, \quad D = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \tag{40}$$

Accordingly, the loss function can be rewritten as:

$$LOSS_t = Y_t' R Y_t \tag{41}$$

where R is a negative semidefinite symmetric $3x3$ matrix characterized by the weight 1, λ and μ on the main diagonal and zeros elsewhere. Then, the central bank optimal control problem is to minimize over choice of $\{i_t\}_{t=0}^{\infty}$ the criterion:

$$\sum_{\tau=0}^{\infty} \delta^{\tau} \left\{ Y_{t+\tau}' R Y_{t+\tau} \right\} \tag{42}$$

subject to the dynamic evolution of the economy described in (35) and given the current state of the economy X_t.

The quadratic objective function, the linear transition equation and the property $E\left(\eta_{t+1} \mid X_t\right) = 0$ are convenient forms for the stochastic optimal linear regulator problem (see Ljungqvist and Sargent, Ch. 4, 2000). It follows that the feedback rule that solves the optimization is linear and independent from the problem's noise statistics, Ω, as the certainty equivalence holds. Then, the first-order necessary condition turns out to be:

$$(S + \delta B' P B) i = -(V' + \delta B' P A) X \tag{43}$$

This implies the following feedback rule for the policy instrument

$$i = fX \tag{44}$$

where f is given by:

$$f = -\left(S + \delta B'PB\right)^{-1}\left(V' + \delta B'PA\right)$$

The $9x9$ matrix P is the solution of the algebraic Riccati equation:

$$P = Q + \delta\left(A + Bf\right)' P\left(A + Bf\right) + f'Sf + Vf + f'V' \tag{45}$$

where Q, V and S are defined as $C'RC$, $C'RD$ and $D'RD$ respectively.

The reaction function (44) resembles an augmented Taylor's rule according to which monetary authorities set the federal funds rate in every period as the optimal response to movements in the current and lagged values of the state variables as well as lagged values of the fed funds rate itself.

Given this optimal feedback rule, the transition law of the economy can be rewritten as $X_{t+1} = MX_t + \eta_{t+1}$ where the $9x9$ matrix M is equal to $A + Bf$.

References

Andrews, D.W., 1993. Tests for parameter instability and structural change with unknown change point, Econometrica, 61, 821-856.

Ball, L., 1999. Efficient rules for monetary policy. International Finance, 63-83

Bank of International Settlements, 1995. Financial structure and the monetary policy transmission mechanism. Basle, Switzerland.

Barro, R.J. and D. Gordon, 1983, A Positive Theory of Monetary Policy in a Natural Rate Model, Journal of Political Economy 91, 589-610.

Bernanke, B. and I. Mihov, 1998, Measuring Monetary Policy, Quarterly Journal of Economics 63, 869-902.

Blinder, A., 1998, Central Banking in Theory and Practice, (Mit press).

Blinder, A., 1997, Distinguished Lecture on Economics and Government: What Central Bankers Could Learn from Academics and Viceversa, Journal of Economic Perspective 11, 3-19.

Brainard, W., 1967. Uncertainty and the effectiveness of policy. American Economic Review, 57, papers and proceeding, 411-425.

Castelnuovo, E., 2003, Squeezing the Interest Rate Smoothing Weight with an Hybrid Expectations Model, FEEM working paper No. 06.

Cecchetti, S.G. and M. Ehrmann, 1999. Does inflation targeting increase output volatility? An international comparison of policymakers' preferences and outcomes. NBER working paper No 7426.

Chadha, J.S. and P. Schellekens, 1999, Monetary policy loss functions: two cheers for the quadratic, DAE Working Paper 99/20.

Christiano L.J., M. Eichenbaum and C. Evans, 1998. Monetary policy shocks: what have we learned and to what end? In: M. Woodford and J.B. Taylor (Eds). Handbook of monetary economics, North-Holland.

Clarida, R., J. Galí and M. Gertler, 2000, Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory, Quarterly Journal of Economics 115, 147-180.

Clarida, R., J. Galí and M. Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, Journal of Economic Literature 37, 1661-1707.

Cukierman, A, and S. Gerlach, 2003, The Inflation Bias Revisited: Theory and Some International Evidence, The Manchester School 71, 541-565.

Cukierman, A, and A. Muscatelli, 2003, Do Central Banks have Precautionary Demands for Expansions and for Price Stability? Theory and Evidence, mimeo, Tel-Aviv University.

Cukierman, A., 2002, Are Contemporary Central Banks Transparent about Economic Models and Objectives and What Difference Does it Make?, Federal Reserve Bank of St. Louis Review 84, 15-45.

De Long, J.B., 1997, America's Peacetime Inflation: the 1970s, in: Romer, C.D. and D.H. Romer, eds., Reducing Inflation: Motivation and Strategy, (Chicago University press).

Dennis, R., 2001. The policy preferences of the US federal reserve. Federal Reserve of San Francisco working paper No. 2001-08.

Dennis, R., 2000. Steps toward identifying central bank policy preferences. Federal Reserve of San Francisco working paper No. 2000-13.

Diebold, F.X., and R.S. Mariano, 1995, Comparing Predictive Accuracy, Journal of Business and Economic Statistics 13, 253-263.

Dolado, J.J., R. Maria-Dolores and M. Naveira, 2004, Are Monetary-Policy Reaction Functions Asymmetric? The Role of Nonlinearity in the Phillips Curve, European Economic Review, forthcoming.

Dolado, J.J., R. Maria-Dolores and F.J. Ruge-Murcia, 2004, Nonlinear Monetary Policy Rules: Some New Evidence for the US, Studies in Nonlinear Dynamics and Econometrics, forthcoming.

Doppelhofer, G.P., R.I. Miller and X. Sala-i-Martin, 2000. Determinants of long-term growth: a Bayesian averaging of classical estimates (BACE) approach, NBER Working paper No. 7750.

Duisenberg, W.F., 2001, Introductory Statement to the Press Conference on 8 November.

Erosa, A. and G. Ventura, 2002, On Inflation as a Regressive Consumption Tax, Journal of Monetary Economics 49, 761-795.

Estrella, A. and F.S. Mishkin, 1999. Rethinking the role of Nairu in monetary policy: implications of model formulation and uncertainty. In: John B. Taylor (Ed.), Monetary policy rules. Chicago University press.

Favero, C.A. and F. Milani, 2001. Parameters' instability, model uncertainty and optimal monetary policy. IGIER working paper No. 196.

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, 545-556.

Fuhrer, J.C., 1997. The (un)importance of forward-looking behavior in price specifications. Journal of Money, Credit and Banking, 29, 338-350.

Fuhrer, J.C., G.R. Moore and S.D. Schuh, 1995, Estimating the Linear-Quadratic Inventory Model: Maximum Likelihood versus Generalized Method of Moments, Journal of Monetary Economics 35, 115-157.

Galí, J., M. Gertler and J.D. Lopez-Salido, 2003a, Markups, Gaps, and the Welfare Costs of Business Fluctuations, mimeo, Universitat Pompeu Fabra.

Galí, J., M. Gertler and J.D. Lopez-Salido, 2003b, Robustness of the Estimates of the Hybrid New Keynesian Phillips Curve, mimeo, Universitat Pompeu Fabra.

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

Geraats, P., 1999, Inflation and Its Variation: An Alternative Explanation, CIDER Working Paper C99-105.

Goodfriend, M., 1991. Interest rate and the conduct of monetary policy. Carnegie-Rochester conference series on public policy, 34, 7-30.

Goodhart, C.A.E., 1997. Why do the monetary authorities smooth interest rates? In: Stefan Collignon (Ed). European monetary policies. Chapter 8, Pinter: London.

Granger, C. and Jeon, Y., 2001. Thick modeling. Mimeo, University of California San Diego.

Hamilton, J.D., 1996, That is what Happened to the Oil Price-Macroeconomy Relationship, Journal of Monetary Economics 38, 215-220.

Hansen, L.P., 1982, Large Sample Properties of Generalized Method of Moments Estimators. Econometrica 50, 1029-1054.

Hansen, L.P.and T.J. Sargent, 2001. Elements of robust control and filtering for macroeconomics. Unpublished monograph, Stanford University.

Ireland, P., 2001, Sticky-Price Models of the Business Cycle: Specification and Stability, Journal of Monetary Economics 47, 3-18.

Ireland, P.N., 1999, Does the Time-Consistency Problem Explain the Behavior of US Inflation?, Journal of Monetary Economics 44, 279-292.

Kydland, F. and E. Prescott, 1977, Rules Rather than Discretion: the Inconsistency of Optimal Plans, Journal of Political Economy 85, 473-490.

Judd, J.P. and G.D. Rudebusch, 1998. Taylor's rule and the fed: 1970-1997. Federal Reserve Bank of San Francisco, Economic Review, 3, 3-16.

Kahneman, D. and A., Tversky, 1979, Prospect Theory: An Analysis of Decision under Risk, Econometrica 47, 263-292.

Kim, D.H., D.R. Osborne and M. Sensier, 2004, Nonlinearity in the Fed's Monetary Policy Rule, Journal of Applied Econometrics, forthcoming.

Kimball, M.S., 1990, Precautionary Saving in the Small and the Large, Econometrica 58, 53-73.

Kuha, J. and J. Temple, 2003, Covariate Measurement Error in Quadratic Regression, International Statistical Review 71, 131-150.

Lansing, K.J. and B. Trehan, 2001. Forward-looking behavior and the optimality of the Taylor rule. Federal Reserve of San Francisco working paper No. 2001-03.

Levin, A., V. Wieland and J.C. Williams, 2001. The performance of forecast-based monetary policy rules under model uncertainty. European Central Bank working paper No 68.

Levin, A., V. Wieland and J.C. Williams., 1999. Robustness of simple monetary policy rules under model uncertainty. In: J.B. Taylor (Ed). Monetary policy rules. 263-299, Chicago University Press.

Ljungqvist, L. and T. Sargent, 2000. Recursive macroeconomic theory. MIT press.

Lowe, P. and L. Ellis, 1997. The smoothing of official interest rates. In: Reserve Bank of Australia (Ed). Monetary policy and inflation targeting: proceedings of a conference.

Lucas, R.E., 1976. Econometric policy evaluation: a critique. Carnegie-Rochester conference series on public policy, 1, 19-46.

Martin, C. and C. Milas, 2004, Modelling Monetary Policy: Inflation Targeting in Practice, Economica, forthcoming.

McCallum, B.T., 1998. Robustness property of a rule for monetary policy. Carnegie-Rochester conference series on public policy, 29, 173-204.

McCallum, B.T., 1997, Crucial Issues Concerning Central Bank Independence, Journal of Monetary Economics 39, 99-112.

McCallum, B.T., and E. Nelson, 1999. Performance of operational policy rules in an estimated semi-classical structural model. In: John B. Taylor (Ed.), Monetary policy rules. Chicago University press..

Meyer, L.H., Swanson, E.T., and V. Wieland, 2001, NAIRU Uncertainty and Nonlinear Policy Rules, American Economic Review 91, 226-231.

Muscatelli, V.A. and C. Trecroci, 2001. Central bank goals, institutional change and monetary policy: evidence from the US and UK. In: Mahadeva, L., P. Sinclair and G. Stern (Eds). Monetary transmission in diverse economies, Cambridge University Press.

Neiss K.S. and E. Nelson, 2001. The real interest rate gap as an inflation indicator, Bank of England working paper No 130.

Nobay, R. and D. Peel, 2003, Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences, Economic Journal 113, 657-665.

Onatski, A., 2000. Minimax analysis of monetary policy under model uncertainty. Mimeo, Harvard University.

Onatski, A. and J.H. Stock, 2002. Robust monetary policy under model uncertainty in a small model of the US economy. Macroeconomic Dynamics 6.

Orphanides, A., 2004, Monetary Policy Rules, Macroeconomic Stability and Inflation: a View from the Trenches, Journal of Money, Credit and Banking, forthcoming.

Orphanides, A., 1998. Monetary policy evaluation with noisy information. Finance and Economics Discussion series No 50, Board of Governors of the Federal Reserve.

Ozlale, U., 2001. Price stability vs. output stability: tales from three Federal reserve administrations. Mimeo, Boston College.

Peersman, G. and F. Smets, 1999. The Taylor rule: a useful monetary policy benchmark for the Euro area? International Finance, 2, 85-116.

Persson, T., and G. Tabellini, 1999, Political Economics and Macroeconomic Policy, in: Taylor, J.B. and M. Woodford, eds., Handbook of Macroeconomics 1, (North-Holland).

Rogoff, K., 1985, The Optimal Degree of Commitment to a Monetary Target, Quarterly Journal of Economics 100, 1169-1190.

Rotemberg, J.J. and M. Woodford, M., 1999, Interest Rate Rules in an Estimated Sticky Price Model, in: Taylor, J.B., ed., Monetary Policy Rules, (Chicago University press).

Rudebusch, G.D., 2002, Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia, Journal of Monetary Economics 49, 1161-1187.

Rudebusch, G.D., 2001. Is the fed too timid?. Monetary policy in an uncertain world. Review of Economics and Statistics, 83, 203-217.

Rudebusch, G.D., 1995. Federal reserve interest rate targeting, rational expectations and the term structure. Journal of Monetary Economics, 35, 245-274.

Rudebusch, G.D. and L.E.O. Svensson, 1999. Policy rules for inflation targeting. In: J.B. Taylor (Ed.), Monetary policy rules. Chicago University press.

Ruge-Murcia, F.J., 2004, The Inflation Bias when the Central Banker Targets the Natural Rate of Unemployment, European Economic Review 48, 91-107.

Ruge-Murcia, F.J., 2003, Inflation Targeting under Asymmetric Preferences, Journal of Money, Credit and Banking 35, 763-785.

Sack, B., 2000. Does the fed act gradually? A VAR analysis. Journal of Monetary Economics, 46, 229-256.

Sack, B. and V. Wieland, 2000. Interest-rate smoothing and optimal monetary policy: a review of recent empirical evidence. Journal of Economics and Business, 52, 205-228.

Salmon, C. and B. Martin, 1999. Should uncertain monetary policymakers do less? In: Hunt, B. and A. Orr (Eds.), Monetary policy under uncertainty. Reserve Bank of New Zealand.

Schaling, E., 1999, The Non-Linear Phillips Curve and Inflation Forecast Targeting, Bank of England working paper No 98.

Smets, F., 1999. Output gap uncertainty. Does it matter for the Taylor rule? In: Hunt, Benjamin and Adrian Orr (Eds.), Monetary policy under uncertainty. Reserve Bank of New Zealand.

Smets, F. and R. Wouters, 2003a, Shocks and Frictions in US Business Cycle: a Bayesian DSGE Approach, mimeo, European Central Bank.

Smets, F. and R. Wouters, 2003b, Output Gaps: Theory versus Practice, mimeo, European Central Bank.

Söderlind, P., U. Söderström and A. Vredin, 2002. Can calibrated New-Keynesian models of monetary policy fit the facts? Mimeo, Sveriges Riksbank.

Söderström, U., 1999. Should central banks be more aggressive? Stockholm school of economics, working paper in economics and finance No 309.

Staiger, D., and J. Stock, 1997, Instrumental Variables Regression with Weak Instruments, Econometrica 65, 557-586.

Stock, J.H., 1999. Making policies robust to model uncertainty: comment on 'Policy rules for inflation targeting' by Rudebusch, G.D. and L.E.O. Svensson. In: J.B. Taylor (Ed), Monetary policy rules. Chicago University press.

Surico, P., 2003a, Measuring the Time-Inconsistency of US Monetary Policy, European Central Bank working paper No. 291.

Surico, P., 2003b, Inflation Targeting and Nonlinear Policy Rules: the Case of Asymmetric Preferences, mimeo, Bocconi University.

Svensson, L.E.O., 1999. Inflation targeting as a monetary policy rule. Journal of Monetary Economics, 43, 607-654.

Svensson, L.E.O., 1999, Inflation Targeting as a Monetary Policy Rule, Journal of Monetary Economics, 43, 607-654.

Svensson, L.E.O., 1997, Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts, American Economic Review 87, 98-114.

Taylor, J.B., 1999. The robustness and efficiency of monetary policy rules as guidelines for interest rate setting by the European Central Bank. Journal of Monetary Economics, 43, 655-679.

Taylor, J.B., 1993, Discretion versus Policy Rules in Practice, Carnegie-Rochester conference series on public policy 39, 195-214.

Tetlow, R.J. and P. von zur Muehlen, 2001. Robust monetary policy with misspecified models: does model uncertainty always call for attenuated policy? Journal of Economic Dynamics and Control, 25, 911-949.

Wieland, V., 1998. Monetary policy and uncertainty about the natural unemployment rate. Finance and Economics Discussion series No 22, Board of Governors of the Federal Reserve System.

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

Woodford, M., 1999, Optimal Monetary Policy Inertia, NBER working paper No 7261.

Yun, T., 1996, Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles, Journal of Monetary Economics 37, 345-370.

Table 1: Reduced-form Estimates

- baseline measures of inflation and output gap -

	1960:1 – 1979:2	*1982:4 2003:2*	*1987:3 – 2003:2*
c1	0.80**	1.45**	2.74**
	(0.06)	(0.22)	(0.34)
c2	0.79**	0.95**	2.15**
	(0.11)	(0.17)	(0.23)
c3	0.01	0.198	-0.19
	(0.01)	(0.101)	(0.10)
c4	-0.11**	-0.041	-0.06
	(0.02)	(0.023)	(0.08)
ρ	0.63**	0.80**	0.85**
	(0.04)	(0.02)	(0.02)
i^*	5.4	6.0	5.3
	-	-	-
π^*	4.0	2.9	2.4
	-	-	-
α	0.01	0.25	-0.13
	(0.03)	(0.13)	(0.07)
γ	-0.29**	-0.09	-0.05
	(0.03)	(0.05)	(0.08)
W(2) p-value	0.000	0.078	0.121
J(19) p-value	0.960	0.874	0.963

Specification: $i_t = (1-\rho)\left[i^* + c_1(\pi_t - \pi^*) + c_2 y_t + c_3(\pi_t - \pi^*)^2 + c_4 y_t^2\right] + \rho i_{t-1} + v_t$

Notes: Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as the change in the consumer price index (cpi) and the output gap is obtained using the CBO potential output. The instrument set includes four lags of cpi inflation, squared cpi inflation, cbo output gap, squared cbo output gap, the fed funds rate and the rate of change in the gdp deflator. The asymmetric preference parameters are computed as $\alpha=2c_3/c_1$ and $\gamma=2c_4/c_2$ while the standard errors are obtained using the delta method. *W*(*n*) refers to the Wald-type statistics of the test for *n* parameter restrictions, which is distributed as a $\chi^2(n)$ under the joint null hypothesis c3=c4=0. The latter is equivalent to the original null of symmetric central bank preferences, $\alpha=\gamma=0$. *J*(*m*) refers to the statistics of Hansen's test for *m* overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 2: Reduced-form Estimates

- alternative measure of inflation -

	1960:1 – 1979:2	*1982:4 2003:2*
c1	0.81**	3.29**
	(0.08)	(0.63)
c2	1.07**	1.13**
	(0.13)	(0.39)
c3	0.03	0.76
	(0.02)	(0.55)
c4	-0.18**	-0.25
	(0.02)	(0.14)
ρ	0.65**	0.89**
	(0.04)	(0.04)
i^*	5.4	6.0
	-	-
π^*	3.7	2.6
	-	-
α	0.08	0.46
	(0.06)	(0.31)
γ	-0.34**	-0.45
	(0.02)	(0.30)
W(2) p-value	0.000	0.194
J(19) p-value	0.959	0.985

Specification: $i_t = (1-\rho)\left[i^* + c_1(\pi_t - \pi^*) + c_2 y_t + c_3(\pi_t - \pi^*)^2 + c_4 y_t^2\right] + \rho i_{t-1} + v_t$

Notes: Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as the rate change in the gdp deflator and the output gap is obtained using the CBO potential output. The instrument set includes four lags of gdp inflation, squared gdp inflation, cbo output gap, squared cbo output gap, the fed funds rate and cpi inflation. The asymmetric preference parameters are computed as $\alpha=2c_3/c_1$ and $\gamma=2c_4/c_2$ while the standard errors are obtained using the delta method. *W*(*n*) refers to the Wald-type statistics of the test for *n* parameter restrictions, which is distributed as a $\chi^2(n)$ under the joint null hypothesis c3=c4=0. The latter is equivalent to the original null of symmetric central bank preferences, $\alpha=\gamma=0$. *J*(*m*) refers to the statistics of Hansen's test for *m* overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 3: Reduced-form Estimates

- alternative measure of output gap -

	1960:1 – 1979:2	*1982:4 2003:2*
c1	0.67**	2.63**
	(0.09)	(0.34)
c2	1.45**	2.17**
	(0.31)	(0.35)
c3	-0.02	0.07
	(0.02)	(0.18)
c4	-0.17**	-0.19
	(0.04)	(0.10)
ρ	0.72**	0.83**
	(0.05)	(0.02)
i*	5.4	6.0
	-	-
π*	4.1	2.9
	-	-
α	-0.06	0.06
	(0.04)	(0.13)
γ	-0.25**	-0.18
	(0.07)	(0.095)
W(2) p-value	0.000	0.161
J(19) p-value	0.969	0.895

Specification: $i_t = (1-\rho)\left[i^* + c_1(\pi_t - \pi^*) + c_2 y_t + c_3(\pi_t - \pi^*)^2 + c_4 y_t^2\right] + \rho i_{t-1} + v_t$

Notes: Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as changes in the cpi and the output gap is obtained with the Hodrick-Prescott (H-P) filter (smoothing parameter = 1600). The instrument set includes four lags of cpi inflation, squared cpi inflation, H-P output gap, squared H-P output gap, the fed funds rate and gdp inflation. The asymmetric preference parameters are computed as $\alpha=2c_3/c_1$ and $\gamma=2c_4/c_2$ while the standard errors are obtained using the delta method. *W*(*n*) refers to the Wald-type statistics of the test for *n* parameter restrictions, which is distributed as a $\chi^2(n)$ under the joint null hypothesis c3=c4=0. The latter is equivalent to the original null of symmetric central bank preferences, $\alpha=\gamma=0$. *J*(*m*) refers to the statistics of Hansen's test for *m* overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 4: Structural Estimates

- baseline normalization of the orthogonality conditions -

	1960:1 – 1979:2	*1982:4 2003:2*
c1	0.96**	1.35**
	(0.03)	(0.04)
c2	1.15**	0.34**
	(0.16)	(0.08)
c3	0.03	0.13
	(0.02)	(0.11)
c4	-0.17**	-0.01
	(0.03)	(0.01)
ρ	0.70**	0.73**
	(0.04)	(0.02)
i*	5.4	6.0
	-	-
π*	4.2	2.5
	-	-
α	0.05	0.20
	(0.04)	(0.18)
γ	-0.29**	-0.06
	(0.02)	(0.10)
t-type statistics p-value	0.000	0.000
W(2) p-value	0.000	0.530
J(19) p-value	0.950	0.895

Specification:

$$E_{t-1}\left\{\left[-\frac{\mu}{\kappa\varphi}\left(i_t-i^*\right)+(1-\rho)\left(\left(\pi_t-\pi^*\right)+\frac{\lambda}{\kappa}y_t+\frac{\alpha}{2}\left(\pi_t-\pi^*\right)^2+\frac{\gamma\lambda}{2\kappa}y_t^2\right)+\rho\left(i_{t-1}-i^*\right)\right]z_{t-1}\right\}=0$$

Notes: This table reports the nonlinear GMM estimates of the structural parameters α and γ. The estimates of the reduced-form coefficients are recovered from the estimates of the structural parameters while the standard errors are computed using the delta method. Inflation, output gap and the instrument set z_{t-1} correspond to the baseline measures described in the notes to Table 1. The t-type test refers to the null hypothesis $(\mu/\kappa\varphi)=0$. $W(n)$ refers to the Wald-type statistics of the test for n parameter restrictions, which is distributed as a $\chi^2(n)$ under the joint null hypothesis $\alpha=\gamma=0$. $J(m)$ refers to the statistics of Hansen's test for m overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 5: Structural Estimates

- alternative normalization of the orthogonality conditions -

	1960:1 – 1979:2	*1982:4 2003:2*
c1	0.79*	3.64**
	(0.08)	(0.77)
c2	1.11**	1.08**
	(0.06)	(0.03)
c3	0.01	0.53
	(0.02)	(0.40)
c4	-0.17**	-0.04
	(0.01)	(0.05)
ρ	0.65**	0.90**
	(0.02)	(0.02)
i^*	5.4	6.0
	-	-
π^*	4.1	2.5
	-	-
α	0.03	0.29
	(0.04)	(0.22)
γ	-0.31**	-0.06
	(0.02)	(0.10)
t-type statistics p-value	0.000	0.000
W(2) p-value	0.000	0.423
J(19) p-value	0.949	0.876

Specification:

$$E_{t-1}\left\{\left[-\frac{\mu}{\lambda\varphi}\left(i_t-i^*\right)+(1-\rho)\left(\frac{\kappa}{\lambda}\left(\pi_t-\pi^*\right)+y_t+\frac{\alpha\kappa}{2\lambda}\left(\pi_t-\pi^*\right)^2+\frac{\gamma}{2}y_t^2\right)+\rho\left(i_{t-1}-i^*\right)\right]z_{t-1}\right\}=0$$

Notes: This table reports the nonlinear GMM estimates of the structural parameters α and γ. The estimates of the reduced-form coefficients are recovered from the estimates of the structural parameters while the standard errors are computed using the delta method. Inflation, output gap and the instrument set z_{t-1} correspond to the baseline measures described in the notes to Table 1. The t-type test refers to the null hypothesis $(\mu/\lambda\varphi)=0$. $W(n)$ refers to the Wald-type statistics of the test for n parameter restrictions, which is distributed as a $\chi^2(n)$ under the joint null hypothesis $\alpha=\gamma=0$. $J(m)$ refers to the statistics of Hansen's test for m overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 6: Inferring the Inflation Target

- baseline measures of inflation and output gap -

	1960:1 – 1979:2	*1982:4 2003:2*
c1	0.82** (0.06)	2.60** (0.72)
c2	0.84** (0.19)	0.97* (0.48)
c3	0.00 -	0.00 -
c4	-0.13** (0.04)	-0.07 (0.09)
ρ	0.68** (0.06)	0.89** (0.04)
*i**	5.4 -	6.0 -
π*	3.61** (0.31)	2.77** (0.25)
α	0.00 -	0.00 -
γ	-0.32** (0.03)	-0.14 (0.21)
J(16) p-value	0.897	0.793

Specification: $i_t = (1 - \rho)\left[\left(i^* - c_1\pi^*\right) + c_1\pi_t + c_2 y_t + c_4 y_t^2\right] + \rho i_{t-1} + v_t$

Notes: Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as the change in the consumer price index (cpi) and the output gap is obtained using the CBO potential output. The instrument set includes four lags of cpi inflation, cbo output gap, squared cbo output gap, the fed funds rate and the rate of change in the gdp deflator. The asymmetric preference parameter on inflation is restricted to zero while the one on the output gap is computed as $\gamma=2c_4/c_2$. The standard errors are obtained using the delta method. *J*(*m*) refers to the statistics of Hansen's test for *m* overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Table 7: Inflation Mean and its Components

	1960:1 – 1979:2	*1982:4 2003:2*
Average Inflation Bias	1.11** (0.12)	0.32 (0.48)
Inflation Target	3.61** (0.31)	2.77** (0.25)
Model-Based Inflation Mean	4.72** (0.26)	3.09** (0.50)
Actual Inflation Mean	4.5 -	2.8 -
Standard Deviation of the Output Gap	2.7 -	2.1 -

Model-based inflation mean: $E(\pi_t) = \pi^* - \frac{c_4}{c_1}\sigma_y^2 = \pi^* - \frac{\gamma\lambda}{2\kappa}\sigma_y^2$

Notes: Standard errors in brackets. The average inflation bias, which is defined as the difference between the model-based average inflation and the inflation target, is recovered from the estimates of the interest rate reaction function reported in Table 6 as $(-c_4\sigma_y^2/c_1)$. The standard errors are obtained using the delta method. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 1 percent and 5 percent significance levels, respectively.

Figure 1: Preference over Output Stabilization

- cubic vs. quadratic –

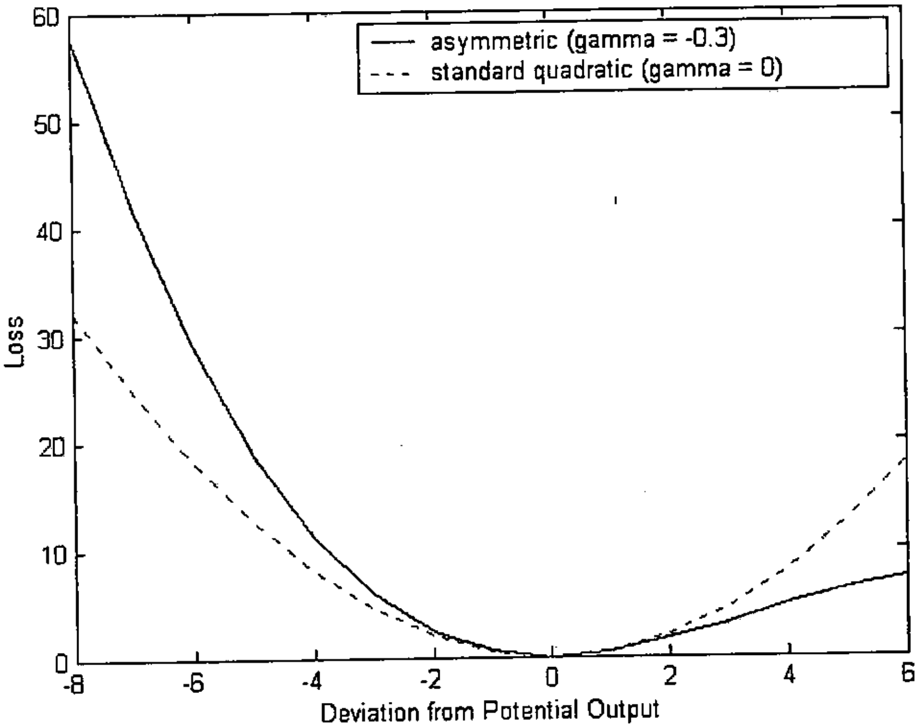

The horizontal axis spans the range of historical values for the CBO output gap during the sample 1960:1 – 2003:2 while the value of gamma in the asymmetric specification is consistent with the estimates reported below.

Figure 2: Federal Funds Rate, CPI Inflation and CBO Output Gap

- full sample: 1960:1 2003:2 -

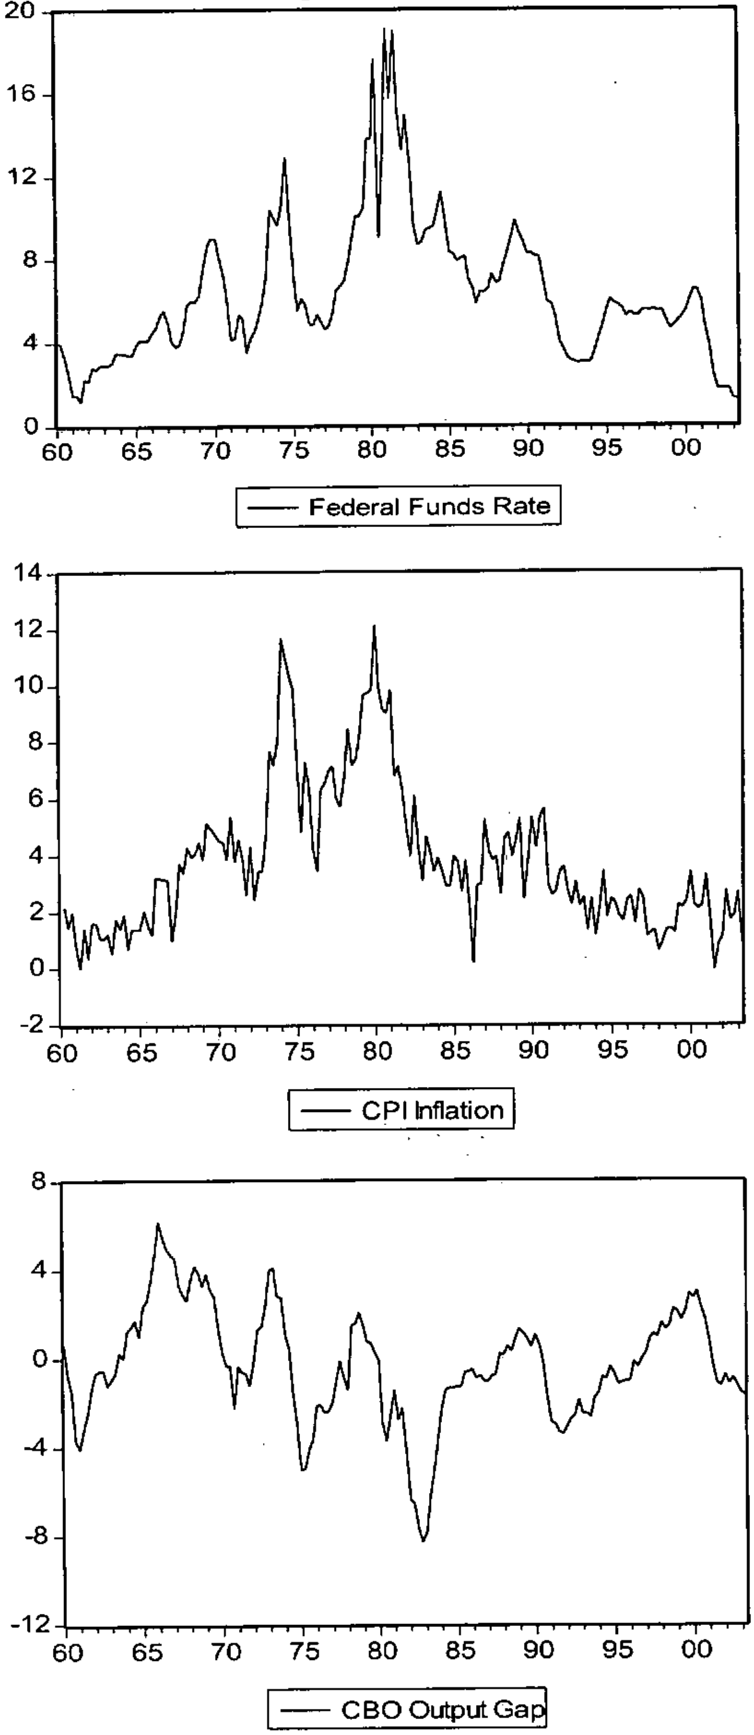

Table 1: Descriptive Statistics

Sample	*Inflation mean*	*Output gap standard deviation*
1960 – 2002	3.78	2.61
1960 – 1982	4.87	3.03
1983 - 2002	2.51	1.98

US quarterly data. Inflation is measured as the changes in the log of the GDP chain-type price index and the output gap is the difference between the log of real GDP and the log of the CBO potential output.

Table 2: Reaction Function and Policy Preference Estimates

- full sample -

Instruments	π^*	α	β	*p-values*
Sample 1960:1 2002:3				
(1)	2.34**	0.09	0.04**	*F-stat: .00/.00*
	(0.24)	(0.11)	(0.01)	*J(7): .13*
(2)	2.33**	0.10	0.04**	*F-stat: .00/.00*
	(0.24)	(0.12)	(0.02)	*J(7): .14*

Specification: $\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t$

Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as changes in the GDP chain-type price index and output gap is obtained from the CBO. The instrument set (1) includes a constant, three lags of inflation, output gap and squared output gap. The instrument set (2) includes a constant, five lags of inflation, and two lags of output gap and squared output gap. F-stat refers to the statistics of the hypothesis testing for weak instruments relative to output gap and squared output gap, respectively. J(*m*) refers to the statistics of Hansen's test for *m* overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 5 percent and 10 percent significance levels, respectively.

Table 3: Reaction Function and Policy Preference Estimates
- sub samples -

Instruments	π^*	α	β	γ	*p-values*
Sample 1960:1-1979:2					
(1)	3.42**	-0.63**	0.14**	-0.46**	*F-stat: .00/.00*
	(0.58)	(0.19)	(0.06)	(0.15)	*J(7): .35*
(2)	3.69**	-0.84**	0.19**	-0.46**	*F-stat: .00/.00*
	(0.67)	(0.27)	(0.08)	(0.13)	*J(7): .37*
Sample 1982:4-2002:3					
(1)	1.96**	-0.18**	0.01	-0.07	*F-stat: .00/.00*
	(0.13)	(0.08)	(0.01)	(0.17)	*J(7): .51*
(2)	1.94**	-0.16*	0.01	-0.10	*F-stat: .00/.00*
	(0.14)	(0.09)	(0.02)	(0.24)	*J(7): .47*
Sample 1987:3-2002:3					
(1)	1.76**	-0.13**	0.04	-0.79	*F-stat: .00/.00*
	(0.19)	(0.06)	(0.04)	(0.83)	*J(7): .73*
(2)	1.96**	-0.17**	-0.01	-0.03	*F-stat: .00/.00*
	(0.18)	(0.08)	(0.04)	(0.49)	*J(7): .38*

Specification: $\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t$

Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as changes in the GDP chain-type price index and output gap is obtained from the CBO. The instrument set (1) includes a constant, three lags of inflation, output gap and squared output gap. The instrument set (2) includes a constant, five lags of inflation, and two lags of output gap and squared output gap. F-stat refers to the statistics of the hypothesis testing for weak instruments relative to output gap and squared output gap, respectively. J(m) refers to the statistics of Hansen's test for m overidentifying restrictions which is distributed as a $\chi^2(m)$ under the null hypothesis of valid overidentifying restrictions. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 5 percent and 10 percent significance levels, respectively.

Table 4: The Average Inflation Bias

Instruments	***Inflation Bias***	*Inflation Target*	*Inflation Bias + Inflation Target*	*Inflation Mean*
Sample 1960:1-1979:2				4.39
(1)	1.01**	3.42**	4.43**	
	(0.39)	(0.58)	(0.52)	
(2)	1.36**	3.69**	5.05**	
	(0.54)	(0.57)	(0.68)	
Sample 1982:4-2002:3				2.53
(1)	0.03	1.96**	1.99**	
	(0.06)	(0.13)	(0.14)	
(2)	0.04	1.94**	1.98**	
	(0.07)	(0.14)	(0.14)	
Sample 1987:3-2002:3				2.36
(1)	0.16	1.76**	1.92**	
	(0.11)	(0.19)	(0.12)	
(2)	-0.01	1.96**	1.95**	
	(0.13)	(0.18)	(0.13)	

Standard errors in parenthesis. The instrument set (1) includes a constant, three lags of inflation, output gap and squared output gap. The instrument set (2) includes a constant, five lags of inflation, and two lags of output gap and squared output gap. The superscript ** and * denote the rejection of the null hypothesis that the true coefficient is zero at the 5 percent and 10 percent significance levels, respectively. The inflation bias is computed as $\beta\sigma_y^2$.

Figure 1: Preferences over Output Stabilization

- cubic vs. quadratic -

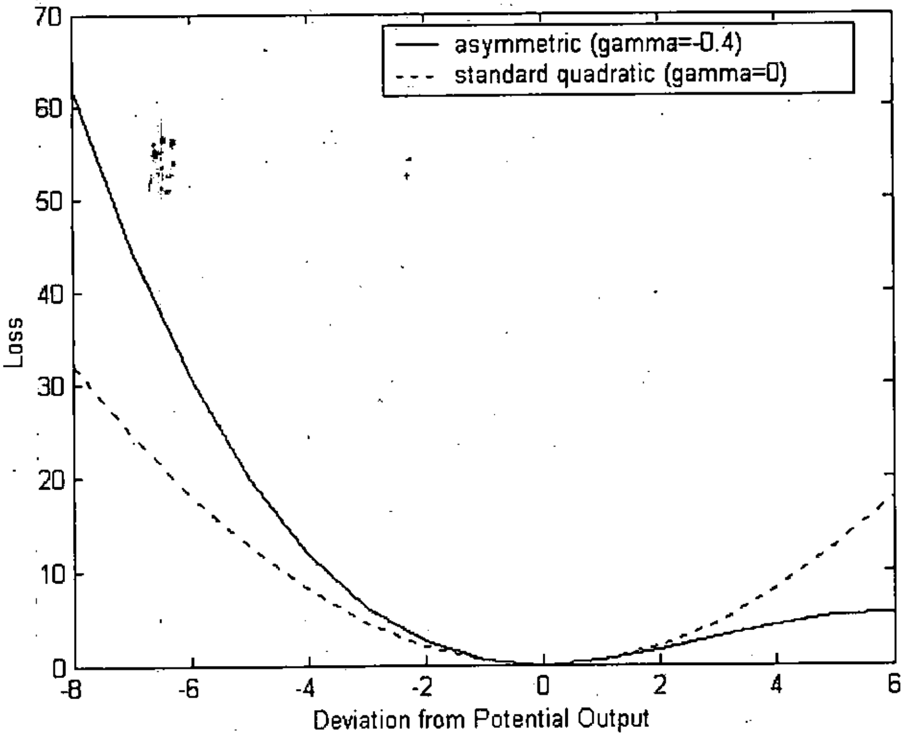

The horizontal axis spans the range of historical values for the CBO output gap during the sample 1960:1 – 2002:3 while the value of gamma in the asymmetric specification is consistent with the estimates reported below.

Figure 2: The Evolution of the Inflation Bias over Time

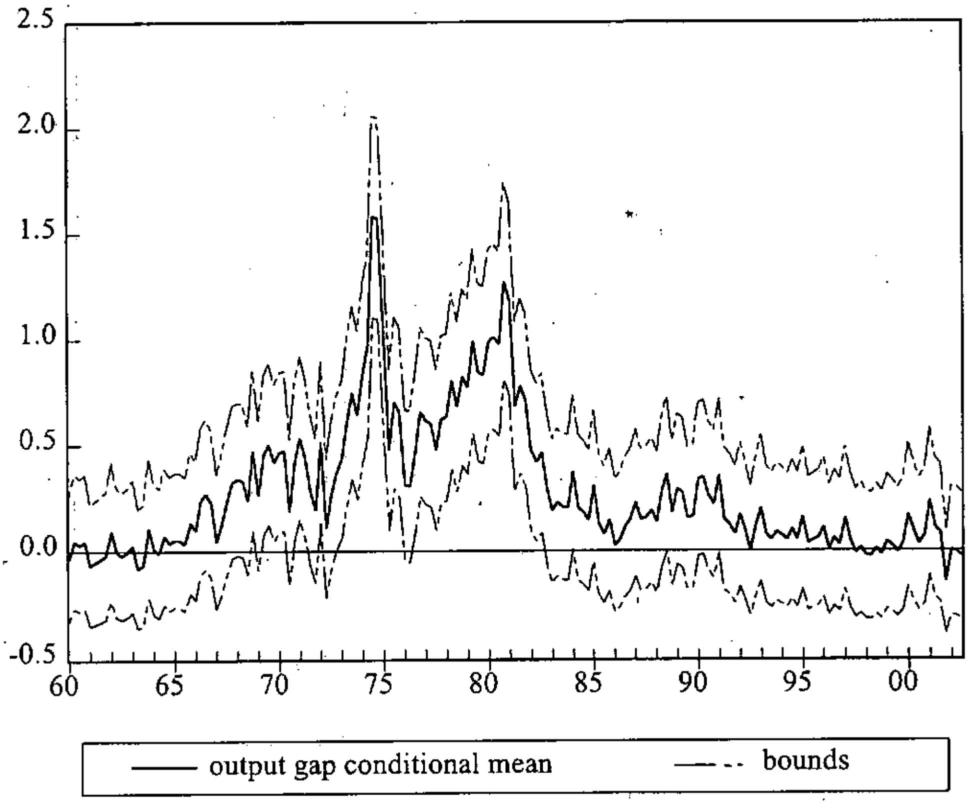

Sample: 1960:1 – 2002:3, US quarterly data. Inflation is measured as the changes in the log of the GDP chain-type price index and the output gap is the difference between the log of real GDP and the log of the CBO potential output. The kernel estimates of the output gap conditional mean on inflation are obtained using the Nadaraya-Watson method, a second order Gaussian kernel and the likelihood cross validation procedure to get a value for the fixed bandwidth parameter. Dashed lines represent upper and lower bounds of the 95% confidence interval.

Table 1 - Historical policy rule vs. optimal policy rules: a quantitative comparison of empirical evidence

Panel A: Descriptive statistics of the fitted policy rule, 1987:3 – 2001:1

Mean	*Standard deviation*
0.000	1.7307

Panel B: Descriptive statistics, policy preferences and average distance of the optimal rules

Author/s	*Estimates*	*Mean*	*Standard deviation*	*Average distance*
Castelnuovo and Surico (present paper)	$\lambda = 1.000$ $\mu = 8.000$	0.4913	1.9100	1.4459
Dennis (2001)	$\lambda = 0.815$ $\mu = 6.181$	0.4888	1.9797	1.4894
*Favero and Rovelli (2002)**	$\lambda = 0.00125$ $\mu = 0.00850$	0.3564	16.9932	41.5373
Ozlale (2001)	$\lambda = 0.525$ $\mu = 0.975$	0.5563	2.4752	2.8621

* The estimates in Favero and Rovelli are based on the Volcker-Greenspan period, 1980:3 1998:3, rather than on the Greenspan tenure only, from the 1987:3 onwards. As discussed in the main text, this does not affect our conclusions.
Note: the preference parameter on inflation stabilization is normalized to one. The parameter on output stabilization is denoted by λ while the one on interest rate smoothing is μ. The average distance is measured as the mean of the sum of the squared deviations between optimal and fitted policy rates at each point in time.

Table 2 – Optimal monetary policy rules and uncertainty: descriptive statistics

Optimal Rules	*Estimates*	*Mean*	*Standard deviation of interest rate levels*	*Standard deviation of interest rate changes*	*Average distance*
Fitted policy rule	-	0.000	1.7307	0.5207	-
Thin policy rule	$\lambda = 0.000$ $\mu = 0.111$	0.4635	4.2493	1.2980	11.4717
Thick model uncertainty robust policy rule	$\lambda = 0.000$ $\mu = 0.111$	0.0087	1.8024	0.5165	2.0385
Parameter uncertainty robust policy rule	$\lambda = 0.000$ $\mu = 0.111$	0.3051	2.9353	0.8439	3.5341

Note: the preference parameter on inflation stabilization is normalized to one. The parameter on output stabilization is denoted by λ while the one on interest rate smoothing is μ. The average distance is measured as the mean of sum of the squared deviations between optimal and fitted policy rates at each point in time. The thick robust policy rule is computed as the simple average at each point in time of the optimal rates for each of the possible specifications. The parameter uncertainty robust policy rule is computed as multiplicative uncertainty on the key coefficients α_5 (slope of the Phillips curve, equation (1) in the main text) and β_3 (semi-elasticity of the output-gap with respect to the real interest rate, equation (2) in the main text). The uncertainty is determined upon the Variance-Covariance matrix of the OLS estimators.

Figure 1 - Historical policy rule vs. optimal policy rules: a graphical comparison of empirical evidence

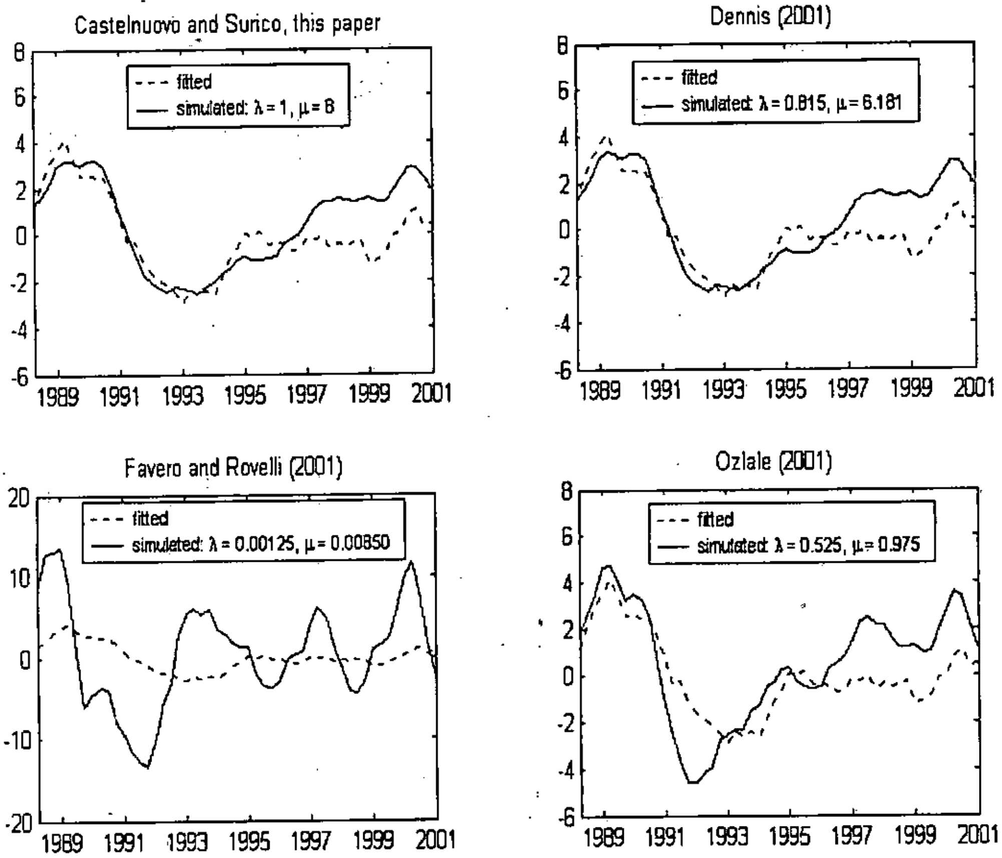

Note: the preference parameter on inflation stabilization is normalized to one. The parameter on output stabilization is denoted by λ while the one on interest rate smoothing is μ. Each optimal path shows the values that the funds rate would have taken if the Fed had historically implemented that optimal policy rule. Demeaned values of the federal funds rate are on the vertical axis.

Figure 2 - Thick robust policy rule vs. thin policy rule

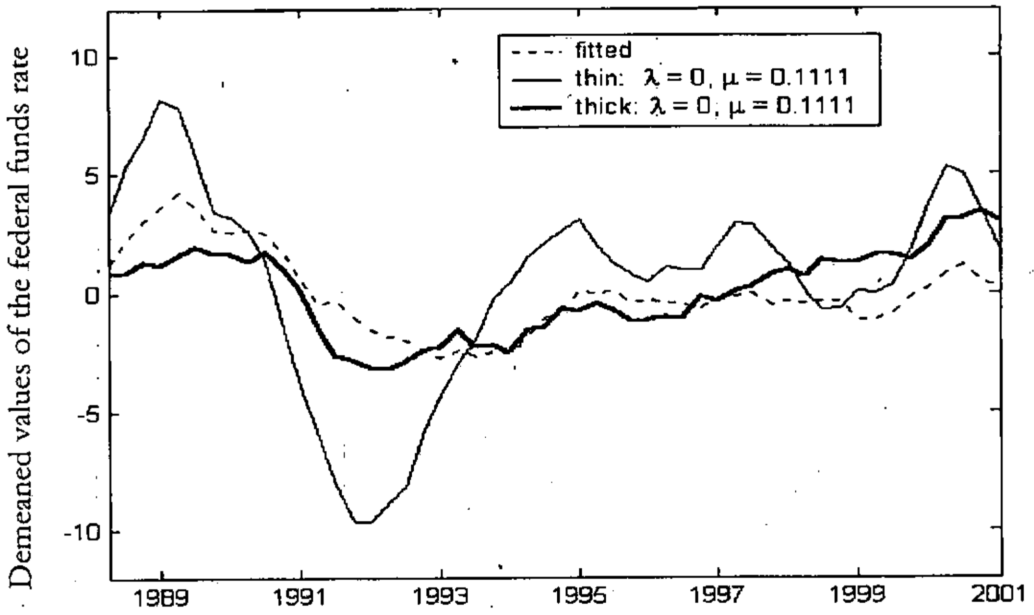

Note: The preference parameter on inflation stabilization is normalized to one. The parameter on output stabilization is denoted by λ while the one on interest rate smoothing is μ. The optimal paths show the values that the funds rate would have taken if the Fed had historically implemented the optimal policy rule. The thick robust policy rule is computed as the simple average at each point in time of the optimal federal funds rates for each of the possible specifications.

Figure 3 - Model vs. parameter uncertainty

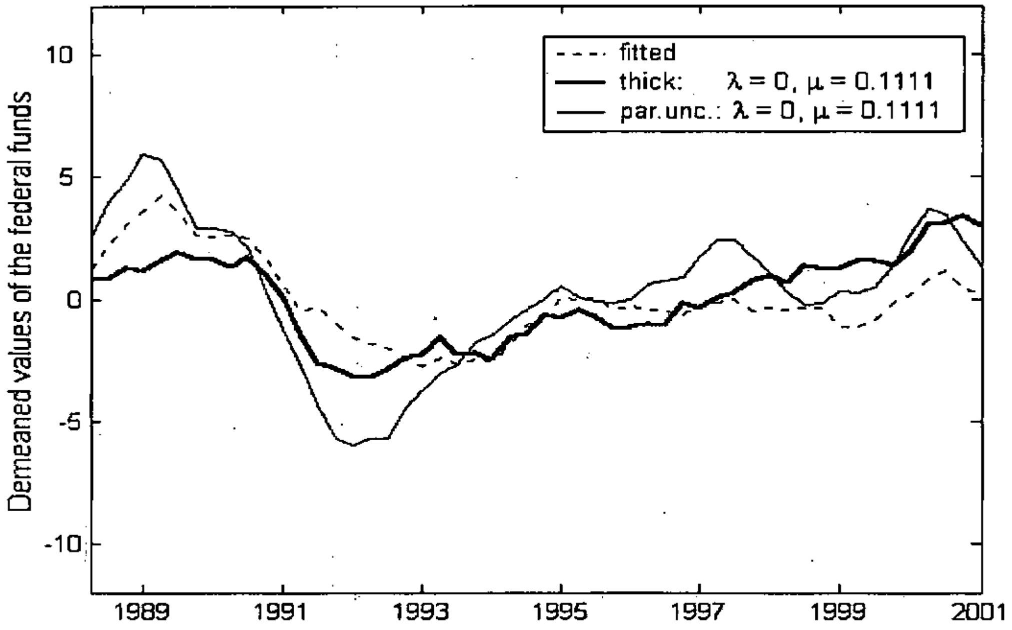

Note: The preference parameter on inflation stabilization is normalized to one. The parameter on output stabilization is denoted by λ while the one on interest rate smoothing is μ. The optimal paths show the values that the funds rate would have taken if the Fed had historically implemented the optimal policy rule. The thick model uncertainty robust policy rule is computed as the simple average at each point in time of the optimal federal funds rates for each of the possible specifications. The parameter uncertainty robust policy rule is computed as multiplicative uncertainty on the key coefficients α_5 (slope of the Phillips curve, equation (1) in the main text) and β_3 (semi-elasticity of the output-gap with respect to the real interest rate, equation (2) in the main text). The uncertainty is determined upon the Variance-Covariance matrix of the OLS estimators.