Measuring the Time-Inconsistency of US Monetary Policy

Paolo Surico*
Bocconi University

April 2003

Abstract

This paper offers an alternative explanation for the rise and fall of postwar US inflation by measuring a novel source of monetary policy time-inconsistency due to Cukierman. 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 3.42% to 1.96%, the inflation bias drops from 1.01% to a value not statistically different from zero moving from the pre- to the post-Volcker regime. This result can be rationalized in terms of the preference on output stabilization, which is found to be asymmetric in the former but not in the latter period.

JEL Classification: E52, E58

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

*This paper has been written while the author was visiting the European Central bank whose kind hospitality is gratefully acknowledged. Address for correspondence: Istituto di Economia Politica, Università Bocconi, Via Gobbi 5, 20136 Milan, Italy. E-mail: paolo.surico@uni-bocconi.it

1 Introduction

The *Journal of Monetary Economics* has hosted an important debate on the behavior of postwar US inflation. In a stimulating contribution, Ireland (1999) shows that Barro and Gordon's (1983) model of time-consistency monetary policy imposes long-run restrictions on the time series properties of inflation and unemployment that cannot be 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 á 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 activities, 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 also 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 activities 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*, which according to Cukierman and Gerlach (2002) comes from the private sector expectations of a more vigorous policy response in recessions than in booms. More specifically, the inflation bias is a function of the central bank preferences and to the extent that a significant policy regime shift has occurred around 1979 with 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 subsamples 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 better interpreted as regime-specific.

This paper contributes to the literature on optimal monetary policy by constructing a measure of the average inflation bias within 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 solutions obtained under commitment and under discretion shows how the observed inflation mean can be successfully decomposed into a target and a bias argument. 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 the feedback rule coefficients and the average inflation bias. Section 4 concludes.

2 The model

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

$$y_t = \theta \left(\pi_t - \pi^e \right) + u_t, \;\; \theta > 0 \tag{1}$$

where y_t is the output gap measured as the difference between actual and potential output, π_t denotes inflation and π^e stands for the inflation expectation in period $t-1$ on the inflation rate in period t. The supply disturbance, u_t, obeys a potentially autoregressive process:

$$u_t = \rho u_{t-1} + \varepsilon_t$$

with $1 > \rho \geq 0$ and i.i.d. shock, ε_t. The private sector has rational expectations

$$\pi^e = E_{t-1}\pi_t \tag{2}$$

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 since it depends on u_t, and its variance is denoted by σ^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:

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

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

$$L_t = \frac{1}{2} \left(\pi_t - \pi^* \right)^2 + \lambda \left(\frac{\exp\left(\gamma y_t\right) - \gamma y_t - 1}{\gamma^2} \right) \tag{4}$$

where $\lambda > 0$ and γ represent the relative weight and the asymmetric preference on output stabilization, respectively. As in Ireland (1999), π^* is assumed stable enough to be approximated by a positive constant. Unlike in the Barro-Gordon model, the target level of real activities is not meant to overambitiously exceed potential. This is consistent with the empirical evidence reported by Ruge-Murcia (2003).

The objective function (4) tends to its minimum whenever both inflation and output gaps shrink and larger losses are associated with larger absolute values at an increasing rate. The linex specification, which has been originally proposed by Varian (1974) and Zellner (1986) in the context of Bayesian econometric analysis and introduced by Nobay and Peel (1998) in the optimal monetary policy literature, allows departures from the quadratic objective in that policy makers may treat differently output contractions and output expansions. Indeed, under an asymmetric loss function deviations of the same size but opposite sign yield different losses and a negative value of γ implies that negative gaps are weighted more severely than positive ones. To see this notice that whenever $y_t < 0$ the exponential component of the loss function dominates the linear component while the opposite is true for $y_t > 0$. The reasoning is reversed for positive values of γ.

The intuition for having an asymmetric loss function with respect to the output gap comes from the labor market asymmetry over the business cycle between the extensive and the intensive margin. Indeed, whenever output is at its potential level the economy experiences full employment and production can only be expanded along the intensive margin, namely by increasing the number of worked hours per employee. By contrast, during recessions also the extensive margin becomes available and production can be lowered through a reduction of both the number of workers and the number of worked hours per employee. This introduces a natural asymmetry in the cost of business fluctuations that policy makers are likely to suffer. A formal microfoundation for an objective function featuring the properties of (4) can be found in Geraats (1999).

Lastly, the linex specification nests the quadratic form as a special case and by means of L'Hôpital's rule on (4) it can be shown that whenever γ tends to zero the central bank objective function reduces to the conventional symmetric parametrization $L_t = \frac{1}{2}\left[(\pi_t - \pi^*)^2 + \lambda y_t^2\right]$. As argued by Ruge-Murcia (2003), this feature is attractive as it allows to test whether the relevant preference parameter is statistically different from zero.

2.1 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^e, to minimize the asymmetric loss function (4) subject to the augmented Phillips curve (1) and to the additional constraint (2) imposed by the rational expectations hypothesis. The corresponding first order conditions are, respectively:

$$(\pi_t - \pi^*) + E_{t-1}\left\{\frac{\lambda\theta}{\gamma}\left[\exp(\gamma y_t) - 1\right] - \mu\right\} = 0 \tag{5}$$

$$-E_{t-1}\left\{\frac{\lambda\theta}{\gamma}\left[\exp(\gamma y_t) - 1\right]\right\} + \mu = 0 \tag{6}$$

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

$$E(\pi_t) = \pi^* \tag{7}$$

where we have used the law of iterated expectations get rid of E_{t-1}. Equation (7) states that the planned inflation rate equals on average the socially desirable inflation rate and therefore it is independent of the output gap.

2.2 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

$$(\pi_t - \pi^*) + E_{t-1}\left\{\frac{\lambda\theta}{\gamma}\left[\exp(\gamma y_t) - 1\right]\right\} = 0 \tag{8}$$

It is instructive at this point to compare the solution obtained under asymmetric preferences with the solution obtained under the standard quadratic case. Whenever γ tends to zero, it is possible to show using L'Hôpital's rule that the optimal monetary policy becomes

$$(\pi_t - \pi^*) = -\lambda\theta E_{t-1}(y_t) \tag{9}$$

This implies that under quadratic preferences the output gap conditional mean is a sufficient statistics for the inflation bias at each point in time. 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.

Turning back to equation (8), we notice that if the output gap is a zero mean, normally distributed process, then $\exp(\gamma y_t)$ is distributed log normal with mean $\exp\left(\frac{\gamma^2}{2}\sigma^2\right)$. It follows that by taking expectations of (8) and rearranging terms, it is possible to write the optimality condition as:

$$1 - \frac{\gamma}{\lambda\theta}E(\pi_t - \pi^*) = \exp\left(\frac{\gamma^2}{2}\sigma^2\right) \tag{10}$$

To compute the *average inflation bias*, we use a simple transformation of the model that confronts directly the time-inconsistency of monetary policy. This amounts to take logs of both side of (10) and gives the following expression:

$$E(\pi_t) \simeq \pi^* - \frac{\lambda\theta\gamma}{2}\sigma^2 \tag{11}$$

A comparison between the expected rates under commitment (7) and under discretion (11) illustrates the source of a novel *average inflation bias*. The time-inconsistency of monetary policy arises here because policy preferences are asymmetric rather than because the desired level of output is above potential like in the Barro-Gordon model. As the private sector correctly anticipates the monetary authorities' incentive to respond more aggressively to output contractions than to output expansions (i.e. $\gamma < 0$), the inflation rate exceeds the first-best solution attainable under commitment. Hence, policy makers end up generating a systematic boost in inflation expectations, which is higher the larger and the more asymmetric the policy preference on output stabilization is.

Possible improvements to the discretionary solution would require the appointment of either 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 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 in (8) is convex in the output gap. When γ tends to zero, such a marginal benefit becomes linear and the inflation bias disappears.

3 The evidence

This section investigates the empirical merits of the asymmetric preference model to account for the dynamics of the postwar US inflation. The analysis spans the period 1960:1-2002:3 and it is conducted on quarterly, seasonally adjusted data which have been obtained from the web site of the Federal Reserve Bank of St. Louis. Inflation is measured as the annualized change in the (log) GDP chain-weighted price index, whereas the output gap is constructed as the difference between the (log) real GDP and the (log) real potential output that is 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 develop a measure of the inflation bias over both the pre- and post-Volcker regimes. We also address the issue of subsample stability by re-estimating the model over Greenspan's tenure, which begins in the third quarter of 1987. Indeed, equation (11) 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 problem across the two eras.

3.1 Preliminary analysis

As a way to provide a preliminary evidence before turning to the estimates of the nonlinear optimal monetary policy (8), we evaluate the performance of the symmetric quadratic paradigm upon the behavior of the inflation bias that this specification predicts. According to equation (9), the conditional mean of the output gap is a sufficient statistics for the difference between realized inflation and the inflation target. In particular, 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 1 displays the nonparametric kernel estimates of the reciprocal of the output gap conditional mean 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 (9). Rather, we are computing from the bivariate 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 number of interesting results emerges from Figure 1. First, the third quarter of 1982 appears to indicate the beginning of a new era as witnessed 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 reserve targeting that Bernanke and Mihov (1998) date in 1982:3. Moreover, the measure of the inflation bias is not statistically different from zero only after 1982, implying that the model of quadratic preferences and potential output targeting is rejected by the data over the earlier regime unless one is willing to accept a highly persistent process for the supply shocks before but not after the structural break. Nevertheless, the reciprocal of the output gap conditional mean and therefore the inflation bias appears to be *on average* statistically different from zero over the first subsample, thereby calling for an extension of the theory to account for the short-run inflation dynamics.

3.2 The reduced-form

The parameter γ and the exponential function in (8) govern the asymmetric response of policy rates to positive and negative deviations of output from potential. Our task consists in estimating a nonlinear reaction function in order to evaluate whether the asymmetric preference is significantly different from zero. This amounts to test linearity against a nonlinear model, which is complicated by the fact that in small samples the estimation criterion is insensitive to the so-called smoothness coefficients as there exists a large set of γ-values yielding almost the same inflation rate behavior (see Granger and Teräsvirta, 1993, Ch. 7). Moreover, it cannot be possible to recover all structural parameters, and in particular neither γ nor the inflation bias are identified. To overcome the issue, we follow Ruge-Murcia (2003) and we take a simple transformation of the model. This involves the linearization of the exponential terms in (8) by means of a first-order Taylor series expansion around $\gamma = 0$, which produces:

$$(\pi_t - \pi^*) + \lambda\theta E_{t-1}(y_t) + \frac{\lambda\theta\gamma}{2} E_{t-1}\left(y_t^2\right) + e_t = 0 \tag{12}$$

with e_t being the remainder of the Taylor series approximation.

This condition relates the inflation rate to the expected values of the level and the squared of the output gap conditional upon the information available at time $t-1$. We solve equation (12) 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:

$$\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t \tag{13}$$

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

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

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 (13) reveals that by assuming an optimizing central bank behavior 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. Nevertheless, the reduced-form parameters allow now to recover the asymmetric preferences, $\gamma = 2\beta/\alpha$, and most importantly the measure of the bias, which is given by the difference between equations (7) and (11), namely $\beta\sigma^2$.

3.3 Empirical results

To the extent that the monetary authorities are more concerned with output contractions than output expansions, the model predicts $\gamma < 0$, $\alpha < 0$ (since λ, $\theta > 0$), and $\beta > 0$. In particular, a negative value of α in (13) together with the aggregate supply curve (1) imply that during recessions the central bank faces an incentive to generate unexpected inflation in an effort to move the economy from negative to positive values of the output gap. Moreover, asymmetric preferences imply a precautionary demand for output expansions that predicts a positive relation between average inflation and the variance of the output gap.

The orthogonality conditions implied by the rational expectation hypothesis makes the Generalized Method of Moments (GMM) a natural candidate to estimate equation (13). 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 instrument sets.

Table 1 reports the estimates of the feedback rule (13) for the whole 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 cannot be rejected. These results are similar to those reported by Ruge-Murcia (2003) and Surico (2002) as they confirm the presence of asymmetric preference using a different method of estimation and a different measure of real activities.

Table 2 reports the estimates for the pre- and post-Volcker regimes. We remove from the second subsample 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 2 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 2. 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. Nevertheless, the most dramatic difference between the two regimes emerges on the squared output gap, which loses any explanatory power for both set of instruments as well as for both post-1979 subsamples. This translates into values of the asymmetric preference parameter, γ, that are not statistically different from zero.

Turning to the measure of the time-inconsistency, Table 3 reports the estimates of the average inflation bias, which according to equation (11) is a convolution of the 'deep' parameters of the model and the variance of the output gap. Given the decline in the volatility of the supply shocks between the 1970s and the 1980s that is documented by Hamilton (1996) and it is reflected by the change in the output gap standard deviations in the very right column, we would expect also the inflation bias to decline moving from the pre- to the post-Volcker period.

The first column of Table 3 shows the measure of the average inflation bias implied by the reduced-form estimates of Table 2. 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. This result mirrors the cross-country evidence for a group of 21 OECD countries reported by Cukierman and Gerlach (2002). Lastly, while 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, its post-Volcker counterparts appear to be higher than the model predicts. This suggests that while 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, it needs to be extended to account more fully for the gap that appears in the data over the last two decades.

4 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 and Gerlach (2002), if policy makers are more concerned about output contractions than output expansions, an inflation bias can emerge on average even though the level of economic activities 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. In contrast, the inflation target declines to 1.96% while the inflation bias disappears over the last two decades. 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. While also other factors such as better policy making and more favorable supply shocks are likely to have played a role, this paper provides empirical support and quantitative measures of a new, additional explanation for the behavior of postwar US inflation.

While suggestive, the results reported in this paper are based on a static model of the economy, and adding explicitly output persistence in the Phillips curve is likely to produce also a state-contingent bias as well as a stabilization bias. However, as shown by Svensson (1997) and the reference therein, the average inflation bias is then larger than it would be without dynamics, thereby suggesting not only that our estimates are better interpreted as a lower bound but also that some notion of persistence may account for the gap between the model-based and the observed average inflation over 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.

References

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).

Cukierman, A, and S. Gerlach, 2002, The Inflation Bias Revisited: Theory and Some International Evidence, mimeo, Tel-Aviv University.

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

Granger, C.W.J. and T. Teräsvirta 1993, Modelling Non-linear Economic Relationship, (Oxford University Press).

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.

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.

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

Nobay, R. and D. Peel, 1998, Optimal Monetary Policy in a Model of Asymmetric Central Bank Preferences, mimeo, London School of Economics.

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

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

Ruge-Murcia, F.J., 2003, Does the Barro-Gordon Model Explain the Behavior of US Inflation? A Reexamination of the Empirical Evidence, Journal of Monetary Economics, forthcoming.

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

Surico, P., 2002, Uncovering Policy Makers' Loss Function, Econometrics and Applied Economics Series, Bocconi University working paper No 9.

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

Varian, H., 1974, A Bayesian Approach to Real Estate Assessment, in: Feinberg, S.E., and A. Zellner, eds., Studies in Bayesian Economics in Honour of L.J. Savage (North Holland).

Zellner, A., 1986, Bayesian Estimation and Prediction Using Asymmetric Loss Functions, Journal of the American Statistical Association 81, 446-451.

Figure 1: The Evolution of the Inflation Bias over Time

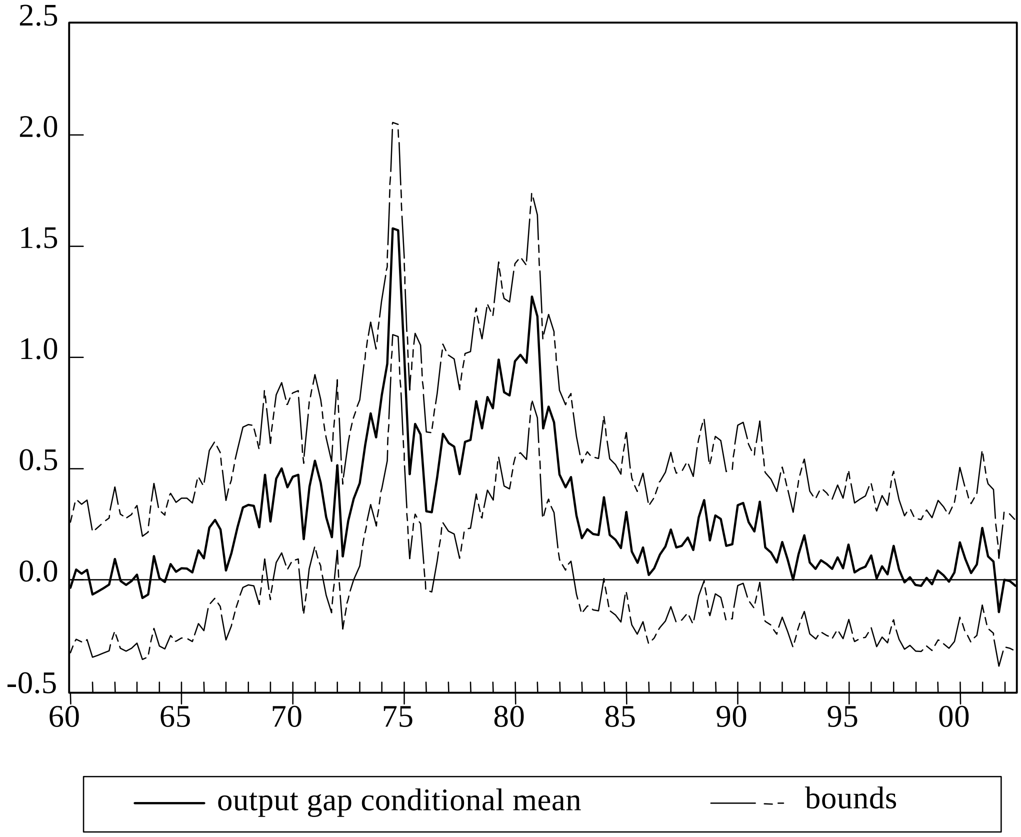

Sample: 1960:1 – 2002:3, US quarterly data. The kernel estimates of the output gap conditional mean on inflation are obtained using the Nadaraya-Watson method, a second order Gaussian kernel and the likelihood cross validation procedure to get a value for the fixed bandwidth parameter. Dashed lines represent upper and lower bounds of the 95% confidence interval.

Table 1: Reaction Function and Policy Preferences Estimates

- full sample -

Instruments	***p****	***a***	***b***	*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: $p_t = p^* + ay_t + by_t^2 + v_t$

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 2: Reaction Function and Policy Preferences Estimates

- sub samples -

Instruments	***p****	***a***	***b***	***g***	*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)	(0.10)	(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: $p_t = p^* + a y_t + b y_t^2 + v_t$

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: The Average Inflation Bias

Instruments	***Inflation Bias***	*Inflation Target*	*Inflation Bias + Inflation Target*	*Inflation Mean*	*Standard Deviation of Output Gap*
Sample 1960:1-1979:2				4.39	2.63
(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	2.15
(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	1.70
(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 bs^2.