

# Measuring the Time-Inconsistency of US Monetary Policy

Paolo Surico\*  
Bocconi University

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

---

\*I wish to thank Alberto Alesina, Filippo Altissimo, Efrem Castelnuovo, Sheng-Hung Chung, Günter Coenen, Carlo Favero, Jordi Galí, Tommaso Monacelli, Anton Muscatelli, Jorge 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. 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 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 The model

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

$$y_t = \theta (\pi_t - \pi_t^e) + u_t, \quad \theta > 0 \quad (1)$$

where  $y_t$  is the output gap measured as the difference between actual and potential output,  $\pi_t$  denotes inflation and  $\pi_t^e$  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  $\varepsilon_t$  is an i.i.d. shock with zero mean and variance  $\sigma_\varepsilon^2$ . The private sector has rational expectations

$$\pi_t^e = E_{t-1} \pi_t \quad (2)$$

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 as it depends on  $u_t$ , and its variance, which is a positive function of both  $\rho$  and  $\sigma_\varepsilon^2$ , is denoted by  $\sigma_y^2$ .

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:

$$\underset{\{\pi_t\}}{\text{Min}} \quad E_{t-1} \sum_{\tau=0}^{\infty} \delta^\tau L_{t+\tau} \quad (3)$$

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

$$L_t = \frac{1}{2} (\pi_t - \pi^*)^2 + \frac{\lambda}{2} \left( y_t^2 + \frac{\gamma}{3} y_t^3 \right) \quad (4)$$

where  $\lambda > 0$  and  $\gamma$  represent the relative weight and the asymmetric preference on output stabilization, respectively. The inflation target,  $\pi^*$ , 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 (4) departs 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  $\gamma$  implies that negative output gaps are weighted more severely than positive ones. To see this notice that whenever  $y_t < 0$  the cubic term,  $\gamma y_t^3$ , 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  $\gamma$  is equal to zero the central bank objective function (4) 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  $\gamma$  that is consistent with the estimates reported below.<sup>1</sup>

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 (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 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,  $\pi_t$ , and expected inflation,  $\pi_t^e$ , 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:

$$(\pi_t - \pi^*) + E_{t-1} \left\{ \lambda \theta \left[ y_t + \frac{\gamma}{2} y_t^2 \right] - \mu \right\} = 0 \quad (5)$$

$$-E_{t-1} \left\{ \lambda \theta \left[ y_t + \frac{\gamma}{2} y_t^2 \right] \right\} + \mu = 0 \quad (6)$$

with  $\mu$  being the Lagrange multiplier associated to the rational expectation constraint. Combining the optimality conditions (5) and (6) to eliminate  $\mu$ , and taking expectations of the resulting expression produce

$$E(\pi_t) = \pi^* \quad (7)$$

---

<sup>1</sup>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  $\gamma$  reported below, the cubic and the linex function behave very similarly.

where we have used the law of iterated expectations to 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  $\pi_t$  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

$$(\pi_t - \pi^*) + E_{t-1} \left\{ \lambda \theta \left[ y_t + \frac{\gamma}{2} y_t^2 \right] \right\} = 0 \quad (8)$$

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

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

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 (8), and using the fact that the unconditional mean of the output gap is zero, we obtain the following expression:

$$E(\pi_t) = \pi^* - \frac{\lambda \theta \gamma}{2} \sigma_y^2 \quad (10)$$

The comparison between the expected rates under commitment (7) and under discretion (10) 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  $\lambda$  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  $\gamma$ . 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  $\gamma$  is equal to zero as it is in equation (9), 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.

### **3 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 (10) 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.

### 3.1 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 (9), 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 (9). 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.

### 3.2 The reduced-form

We solve equation (8) for  $\pi_t$  and prior to estimation we replace expected output gaps with actual values. The empirical version of the feedback rule is given by:

$$\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t \quad (11)$$

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

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

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 (11) 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 (7) and the solution under discretion (10).

### 3.3 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  $\lambda$ ,  $\theta > 0$ ), 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 (11). 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 (11) for the full sample. Each row corresponds to a different set of instruments. The parameter on the output gap,  $\alpha$ , is not statistically different from zero whereas the parameter on the squared output gap,  $\beta$ , 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  $\alpha$  can only be rationalized by a decline in  $\lambda$ , 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  $\gamma$  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 (10), 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.

## 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 (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.

## 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., 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.
- Cukierman, A, and S. Gerlach, 2003, The Inflation Bias Revisited: Theory and Some International Evidence, *The Manchester School* 71, 541-565.
- Geraats, P., 1999, Inflation and Its Variation: An Alternative Explanation, *CIDER Working Paper* C99-105.
- 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, 2003, Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences, *Economic Journal* 113, 657-665.
- 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* 50, 1375-1390.
- Staiger, D., and J. Stock, 1997, Instrumental Variables Regression with Weak Instruments, *Econometrica* 65, 557-586.
- Surico, P., 2003a, Measuring the Time-Inconsistency of US Monetary Policy, ECB 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., 1997, Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts, *American Economic Review* 87, 98-114.

**Table 1: Descriptive Statistics**

<i>Sample</i>	<i>Inflation mean</i>	<i>Output gap standard deviation</i>
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 -

<i>Instruments</i>	$\mathbf{p}^*$	$\mathbf{a}$	$\mathbf{b}$	<i>p-values</i>
Sample 1960:1 2002:3				
(1)	2.34** (0.24)	0.09 (0.11)	0.04** (0.01)	<i>F-stat: .00/.00</i> <i>J(7): .13</i>
(2)	2.33** (0.24)	0.10 (0.12)	0.04** (0.02)	<i>F-stat: .00/.00</i> <i>J(7): .14</i>

Specification:  $\mathbf{p}_t = \mathbf{p}^* + \mathbf{a}y_t + \mathbf{b}y_t^2 + v_t$

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: Reaction Function and Policy Preference Estimates**  
- sub samples -

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

Specification:  $p_t = p^* + a y_t + b y_t^2 + v_t$

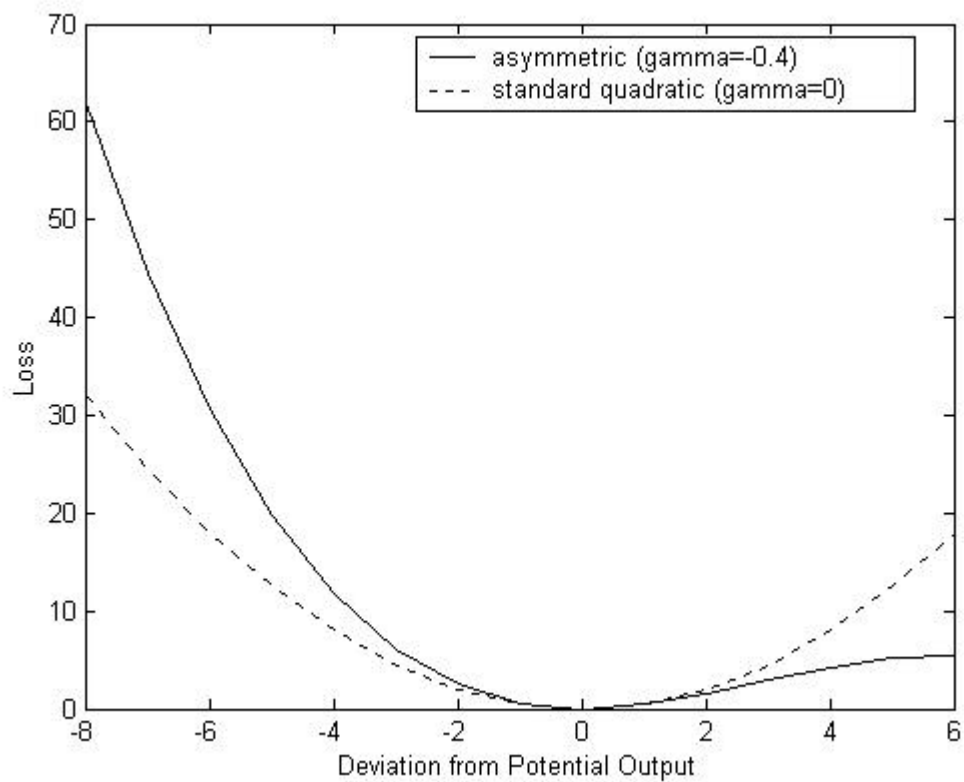
Standard errors using a four lag Newey-West covariance matrix are reported in brackets. Inflation is measured as changes in the GDP chain-type price index and output gap is obtained from the CBO. The instrument set (1) includes a constant, three lags of inflation, output gap and squared output gap. The 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 4: The Average Inflation Bias**

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

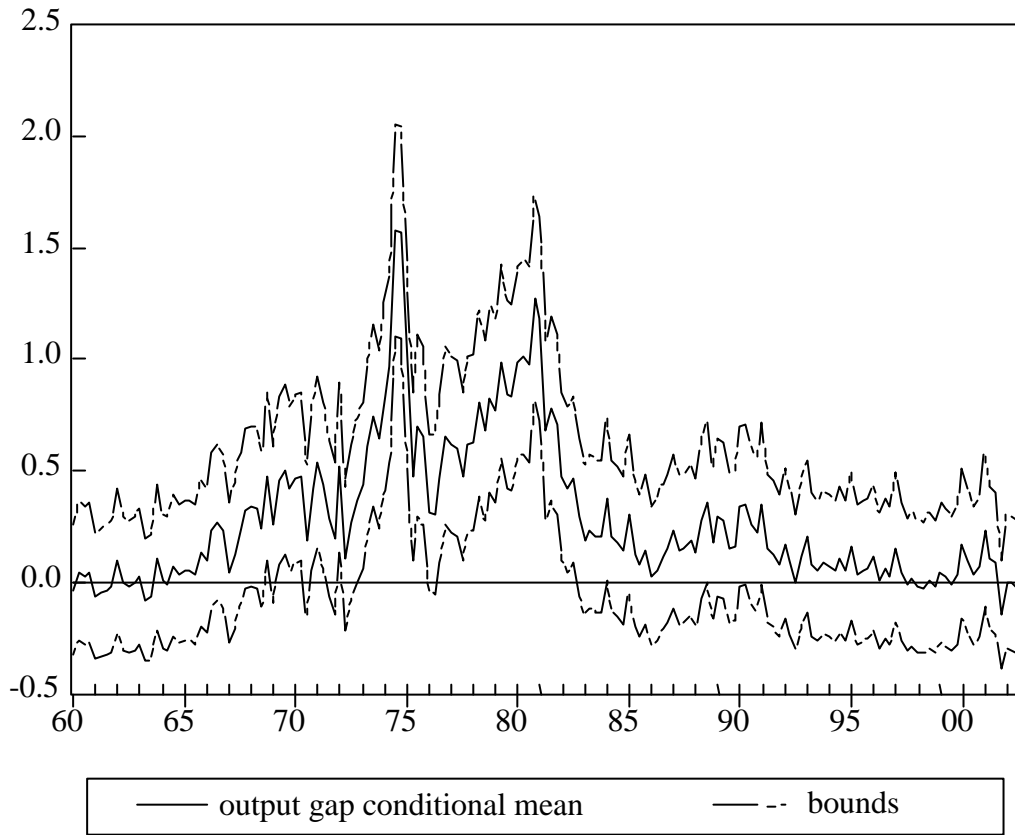
Standard errors in parenthesis. The instrument set (1) includes a constant, three lags of inflation, output gap and squared output gap. The instrument set (2) includes a constant, five lags of inflation, and two lags of output gap and squared output gap. The superscript \*\* and \* denote the rejection 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  $bs_y^2$ .

**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.