Economica (2008) **75**, 22–38
doi:10.1111/j.1468-0335.2007.00590.x

Measuring the Time Inconsistency of US Monetary Policy

By Paolo Surico

Bank of England and University of Bari

Final version received 2 April 2006.

This paper offers an alternative explanation for the great inflation of the 1970s by measuring a novel source of monetary policy time inconsistency. In the presence of asymmetric preferences, the monetary authorities generate a systematic inflation bias through the private-sector expectations of a larger policy response in recessions than in booms. The estimated Fed's implicit target for inflation has declined from the pre- to the post-Volcker regime. The average inflation bias was about 1% before 1979, but this has disappeared over the last two decades, because the preferences on output stabilization were large and asymmetric only in the former period.

Introduction

The behaviour 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. A more favourable macroeconomic environment and a better policy management in the recent past have been advocated as possible explanations of such great moderation. On the other hand, an important strand of the literature has investigated whether the time consistency problem of monetary policy can account for the movements in US inflation.

In a stimulating contribution, Ireland (1999) shows that Barro and Gordon's (1983) model of time consistency 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. This optimal plan, however, 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 is the time consistency problem of monetary policy.

As Persson and Tabellini (1999) make clear, the central bankers' ambition of attaining a level of unemployment below the natural rate is crucial to generating the inflation bias *à 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 practising the time inconsistency of their actions and eventually would revise their objective. Describing his experience as vice-chairman, Blinder (1998) claims that the Federal Reserve 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 history of US monetary policy.

Recent literature has in fact been focusing on the role of non-quadratic objective functions. This research programme was motivated by a growing concern among policy-makers about the hypothesis of symmetric central bank preferences. 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: he claims that '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 such as the ECB, and on the occasion of an interest rate cut of 50 basis points F. W. Duisenberg, in an introductory statement to a press conference on 8 November 2001, stated: 'the maintenance 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 empirical side, 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 has been targeting a level of real activity different from the natural rate is rejected, but the hypothesis that it has been weighting output contractions more severely than output expansions is not. This suggests the existence of a novel *average inflation bias* which, according to Cukierman (2002), comes from the private-sector expectations of a more vigorous policy response in recessions than in booms.

The average inflation bias is indeed a function of both the preferences of the central bank and the volatility of the output gap, which is dramatically declined over time in the US postwar data (see Table 1). 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. Bearing this observation in mind, 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 column (2) of Table 1 seems also consistent with this view.

This paper seeks to contribute to the literature on optimal monetary policy by proposing a measure of the average inflation bias that arises in a model of asymmetric

TABLE 1
DESCRIPTIVE STATISTICS[a]

Sample	Inflation mean (1)	Output gap standard deviation (2)
1960(I)–2005(II)	3.71	2.53
1960(I)–1982(III)	4.91	2.88
1982(IV)–2005(II)	2.51	2.04

[a]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.

central bank preferences. To this end, I develop a novel identification strategy that allows us 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 reveals that the actual inflation mean can be effectively decomposed into a target and a bias argument.

The paper also shows that the policy preferences of the Fed have significantly changed during the last forty years. In particular, while the implicit inflation target declined from 3.81% to roughly 2%, the average inflation bias, whose estimates were around 1% before 1979, has vanished over the last two decades. This finding 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.

On the theoretical side, a number of recent studies explore some 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. Galí *et al.* (2003) 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 United States have been historically large and asymmetric. Furthermore, 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), thereby suggesting that policy-makers, who aggregate over individual welfare, may be loss-averse.

Despite a growing interest on asymmetric preferences, only a few studies, developed independently, estimate asymmetric central bank reaction functions. Cukierman and Muscatelli (2003) and Martin and Milas (2004) provide some international evidence that supports the notion of nonlinear interest rate rules. Ruge-Murcia (2003), and Cukierman and Gerlach (2003) adopt an inflation rate reaction function and favour the hypothesis of an asymmetric objective for some OECD economies. To the best of my knowledge, this paper is the first study that quantifies the average inflation bias associated with asymmetric preferences and assesses its possible contribution to the great inflation of the 1960s and 1970s.

The paper is organized as follows. Section I sets up the model and solves for the optimal monetary policy under asymmetric preferences on the output gap. Section II shows that the results in Section I are robust to the specification of asymmetric preferences on inflation. Section III derives the reduced form of the model and reports the estimates of both the feedback rule coefficients and the average inflation bias. Section IV concludes.

I. The Model

This section describes the structure of the economy and the specification of the loss function, and derives solutions of the central bank optimization problem under commitment and discretion.

(a) Structure of the economy

Private-sector behaviour is characterized by an expectations-augmented Phillips curve:

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

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 at time $t-1$ on the inflation rate in period t. The supply disturbance u_t obeys an autoregressive process $u_t = \rho u_{t-1} + \varepsilon_t$, where $\rho \in [0, 1)$ and ε_t is an independent and identically distributed (i.i.d) shock with zero mean and variance σ_ε^2. The private sector has rational expectations:

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

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

Inflation expectations π_t^e may also be defined as $E_t\pi_{t+1}$. In the latter case equation (1) can be rearranged as $\pi_t = \pi_t^e + \frac{1}{\theta}y_t - \frac{1}{\theta}u_t$, which resembles the New Keynesian Phillips curve in Clarida *et al.* (1999) and King (2000). It is worth noticing that this alternative specification of π_t^e does not affect the derivation of optimal monetary policy: under discretion, expected inflation is taken as given; under commitment, expected inflation does not appear in the policy-makers' first-order conditions as it enters equation (1) linearly.

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.

Equation (1) may also be specified in a nonlinear form representing, for instance, a convex supply schedule. Appendix A shows that the nonlinearity in the policy rule induced by asymmetric preferences on the output gap is robust to specifying a nonlinear supply curve. Moreover, Dolado *et al.* (2005 and 2004) do not reject the hypothesis of a linear supply curve for the United States.

(b) The objective function of the monetary authorities

In the spirit of the Barro–Gordon literature, the central bank is assumed to have full and direct control over inflation, which is chosen to minimize the following intertemporal criterion:

$$\min_{\{\pi_t\}} E_{t-1} \sum_{\tau=0}^{\infty} \delta^\tau L_{t+\tau}, \tag{3}$$

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

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

where $\lambda > 0$ and γ represent the relative weight and the asymmetric preference on output stabilization. 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 exceed potential overambitiously. This is consistent with the empirical evidence in Ruge-Murcia (2003).

The linex specification departs from the quadratic objective in that policy-makers are allowed, but not required, to treat positive and negative deviations of output differently from the target. Varian (1974) and Zellner (1986) proposed this functional form in the context of Bayesian econometric analysis, while Nobay and Peel (2003) introduced it in the optimal monetary policy literature.[1]

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. 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 linex specification nests the quadratic form as a special case, and by means of L'Hôpital's rule it can be shown that if γ tends 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.

It is worth mentioning that according to the linex function the marginal loss from the output gap decreases as output rises further above equilibrium. Indeed, a zone targeting is another plausible way to model asymmetric preferences. In the latter case, the nonlinearity arises because the central bank attaches different weights to small and large deviations from the targets so that only large gaps require a policy response. Zone-like behaviours are associated with a *size* asymmetry, whereas the linex specification is associated with a *sign* asymmetry (see Boinet and Martin 2005).

The specification of an asymmetric loss with respect to the output gap alone is motivated by empirical as well as theoretical considerations. At the empirical level, Surico (2007) derives a general, nonlinear interest rate rule within a model of non-quadratic preferences over both inflation and output, and finds evidence of an asymmetric objective for the latter but not for the former variable.[2] At the theoretical level, Geraats (1999) shows that the labour market flows over the business cycle provide a natural micro foundation for an asymmetric welfare criterion as the firms' hiring–firing decisions are taken mainly along the extensive margin during recessions but along the intensive margin during booms. I show in Section II that the theoretical results derived in this paper are robust to the specification of asymmetric preferences over inflation.

(c) Optimal monetary policy under commitment

This section 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 (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 as follows:

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

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

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

$$(7) \qquad E(\pi_t) = \pi^*,$$

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 is independent of the output gap.

(d) Optimal monetary policy under discretion

If commitment is infeasible, the monetary authorities choose 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. The discretionary solution then reads:

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

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 that the optimal monetary policy becomes

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

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

The output gap is a zero mean, normally distributed process, and $\exp(\gamma y_t)$is characterized by a log normal distribution with mean $(\gamma^2/2)\sigma^2$. It follows that, by taking expectations of (8) and rearranging terms, it is possible to write the optimality condition as

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

To solve for average inflation, we take logs of both sides of (10):

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

The comparison between the expected rates under commitment (7) and discretion (11) illustrates the source of a novel *average inflation bias*. Like the Barro–Gordon model, the time inconsistency of monetary policy arises here because the policy-makers face an incentive to surprise inflation. The nature of the incentive in the two models is however very different. In Barro and Gordon (1983) the incentive comes from the central bank's desire to push the economy beyond its potential level. Here, the asymmetric concern about business cycle fluctuations requires more aggressive policy responses during periods of output contractions than during periods of 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 wish to stabilize output around its potential level. Moreover, the bias is higher the larger and the more asymmetric is the policy preference on output stabilization.

Possible improvements to the discretionary solution include the appointment of a more conservative central banker, i.e. one endowed with a lower relative weight λ in the spirit of Rogoff (1985) or with a lower inflation target than society. Alternatively, the appointment of a more symmetric policy-maker, endowed with a smaller absolute value of γ, can also enhance welfare. 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 goes to zero, as it does in equation (9), the marginal benefit becomes linear and the average inflation bias disappears together with the precautionary motive. This feature resembles the precautionary motive result in the theory of consumption according to which non-quadratic preferences and labour income risks generate above-average saving rates in periods of high uncertainty.

II. Robustness of the Time Inconsistency Measure

This section derives the optimal monetary policy for a more general loss function in which the policy preferences over inflation are also asymmetric. Interestingly, the result that asymmetric preferences over the output gap generate an inflation *bias*, defined as the difference between the discretion and the commitment solutions, is robust to such generalization. Furthermore, the presence of asymmetric preferences over inflation is found to augment the first-order conditions with an inflation *tendency* term that would not otherwise emerge. However, this term appears in both the discretion and the commitment solutions, and therefore it cannot be interpreted as a bias because it does not generate a time inconsistency problem.[3]

(a) A more general loss function

Consider the following objective function:

$$(12)\qquad L_t = \left(\frac{\exp[\gamma_\pi(\pi_t - \pi^*)] - \gamma_\pi(\pi_t - \pi^*) - 1}{\gamma_\pi^2}\right) + \lambda\left(\frac{\exp(\gamma_y y_t) - \gamma_y y_t - 1}{\gamma_y^2}\right)$$

and the aggregate supply curve (1). In contrast to the analysis in the previous section, the central bank is now assumed to have only imperfect control over inflation. The reason for this choice is that asymmetric preferences imply that the variances of the target variables enter policy-makers' first-order condition. Under the hypothesis of perfect inflation control, however, the model of Section I predicts that the variance of inflation is zero, which implies that the asymmetry on inflation stabilization will have no implications for the solution of the model.

(b) Commitment

The solution under commitment can be derived from the first-order conditions of (12) with respect to π_t and π_t^e:

$$(13)\qquad \frac{1}{\gamma_\pi}E_{t-1}\{\exp[\gamma_\pi(\pi_t - \pi^*)] - 1\} + \frac{\lambda\theta}{\gamma_y}E_{t-1}[\exp(\gamma y_t) - 1] - \mu = 0,$$

$$(14)\qquad -\frac{\lambda\theta}{\gamma_y}E_{t-1}[\exp(\gamma y_t) - 1] + \mu = 0.$$

Substituting (14) into (13), we obtain

$$E_{t-1}\{\exp[\gamma_\pi(\pi_t - \pi^*)] - 1\} = 0,$$

which can be rewritten as

$$\text{(15)} \qquad \exp\left(\gamma_\pi E(\pi_t - \pi^*) + \frac{\gamma_\pi^2}{2}\sigma_\pi^2\right) = 1,$$

where we have used the fact that $\exp(\pi_t - \pi^*)$ has a log normal distribution. Taking logs of both sides of (15), we obtain

$$\text{(16)} \qquad E(\pi_t) \simeq \pi^* - \frac{\gamma_\pi}{2}\sigma_\pi^2.$$

When policy preferences are asymmetric with respect to both inflation and the output gap, the average value of π_t has a *tendency* to fall short of target. If preferences on inflation are symmetric (i.e. if $\gamma_\pi = 0$) the inflation mean equals the target. I refer to the inflation aversion term, $-\gamma_\pi/2\sigma_\pi^2$, in (16) as a *tendency* because I show below that it also characterizes the discretionary solution. In line with the rest of the literature, I reserve the expression inflation *bias* for the difference between the commitment and the discretionary solutions.

(c) Discretion

When the monetary authorities take the private sector's expectations as given, the optimal control problem implies a sole first-order condition, with respect to inflation:

$$\text{(17)} \qquad \frac{1}{\gamma_\pi}E_{t-1}\{\exp[\gamma_\pi(\pi_t - \pi^*)] - 1\} + \frac{\lambda\theta}{\gamma_y}E_{t-1}[\exp(\gamma y_t) - 1] = 0.$$

Taking expectations of (17) and linearising the resulting expression using a first-order Taylor approximation, we have

$$\text{(18)} \qquad E(\pi_t) \simeq \pi^* - \frac{\gamma_\pi}{2}\sigma_\pi^2 - \frac{\lambda\theta\gamma_y}{2}\sigma_y^2.$$

Two points are worth noticing about this expression. First, the presence of asymmetric preferences over inflation implies that the discretionary solution has now an additional term, namely $-\gamma_\pi/2\sigma_\pi^2$, relative to solution (11) derived under the assumption of *symmetric* inflation preferences. Second, the inflation bias, which is the difference between the solutions under commitment and under discretion, is still $(\lambda\theta\gamma_y/2)\sigma_y^2$, implying that the theoretical results derived in Section I are robust to such generalization of the objective function.

An empirically plausible characterization of asymmetric preferences is that positive (negative) deviations of inflation (output) from target are associated with larger losses than negative (positive) deviations, corresponding to a positive value for γ_π—and a negative value for γ_y. Under this scenario, the last two terms in (18) move in opposite directions, implying that the relative weight of the asymmetries on inflation and output determines whether average inflation is below or above target.

(d) Empirical issues

As inflation-asymmetric preferences influence average inflation without affecting the time inconsistency measure, they do not generate an inflation bias, though they might explain part of the reason for the lower inflation of the last two decades. If asymmetric preferences had indeed contributed to lower inflation, then the Fed's implicit target π^*

should have been significantly higher than the inflation mean over the most recent period. However, the estimates of the Fed's inflation target reported by Ireland (2006) do not support this view. Furthermore, in Surico (2007) we find no evidence of asymmetric policy responses of the Fed to movements in inflation.

It is worth emphasizing that, even if the assumption of symmetric preferences on inflation was incorrect, the sign of the coefficient on the variance of the output gap in the policy rule would allow us to identify empirically the relative importance of asymmetric preferences over inflation and the output gap. To see this, it is useful to calculate the variance of the output gap on the basis of the expectations-augmented Phillips curve:

$$\text{(19)} \qquad \sigma_y^2 = \theta^2 \sigma_\pi^2 + \sigma_u^2,$$

where we have used the facts that expectations are taken as given under discretion, that the error term is not correlated with inflation and that σ_π^2 reflects the variance of policy control errors. Using (19) to substitute for σ_π^2 in (18), we obtain the following expression for average inflation:

$$\text{(20)} \qquad E(\pi_t) \simeq \pi^* - \frac{1}{2}\left(\frac{\gamma_\pi}{\theta^2} + \lambda\theta\gamma_y\right)\sigma_y^2 + \frac{\gamma_\pi}{2\theta^2}\sigma_u^2,$$

where σ_y^2 is now a function of the variances of the policy control errors and the supply shocks.

To the extent that output recessions are more costly than output expansions, $\gamma_y < 0$. On the other hand, if policy-makers are more concerned about overshooting than undershooting the inflation target, $\gamma_\pi > 0$. The sign of the coefficient on σ_y^2 can thus be informative on the relative importance of the two asymmetries on policy preferences. In particular, according to equation (20), a positive coefficient on σ_y^2, as in this paper, unambiguously identifies the fact that output asymmetries are relatively more important than inflation asymmetries: the time inconsistency of monetary policy still generates an inflation bias. Appendix B further discusses the identification issues involved with asymmetric preferences on inflation.

III. Evidence

This section investigates the empirical merits of the asymmetric preference model to account for the behaviour of postwar US inflation. The analysis considers the period 1960:(I)–2005:(II) and is conducted on quarterly, seasonally adjusted data obtained in October 2005 from the website 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 the results comparable with those reported in Ruge-Murcia (2003), I first considered the whole sample. I then employed my identification strategy to estimate asymmetric preferences and obtain a measure of the inflation bias for both the pre- and the post-Volcker regimes. Indeed, equation (11) makes it clear that the inflation bias is a function of policy-makers' preferences and therefore 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, this identification scheme provides a sharper evaluation of the model because it measures the time inconsistency across two eras.

(a) The reduced form

The parameter γ and the exponential function in (8) govern the asymmetric response of the policy rate to positive and negative deviations of output from potential. I estimate a nonlinear reaction function and assess whether the asymmetric preference parameter is significantly different from zero. This amounts to testing linearity against a nonlinear specification, which is complicated by the fact that it is not possible to recover all structural parameters of the model from the reduced-form estimates. To overcome the issue and identify both γ and the inflation bias, I follow Ruge-Murcia (2003) and 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, and produces

$$(21) \qquad (\pi_t - \pi^*) + \lambda\theta E_{t-1}(y_t) + \frac{\lambda\theta\gamma}{2} E_{t-1}(y_t^2) + e_t = 0,$$

with e_t being the remainder of the approximation.

This condition relates the inflation rate to the expected values of the level and the square of the output gap conditional upon the information available at time $t-1$. We can solve (21) 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

$$(22) \qquad \pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t,$$

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

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

The term in braces is a linear combination of forecast disturbances and an approximation error, and therefore v_t is orthogonal to any variable in the information set available at time $t-1$.

Equation (22) reveals that, under the assumption of an optimizing central bank behaviour the reaction function parameters can be interpreted only 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 is $\beta\sigma_y^2$. The latter is obtained as the difference between the solution of the central bank optimization under commitment (7) and the solution under discretion (11).

(b) Empirical results

When the penalty associated with output contractions is larger than the penalty associated with output expansions of the same size, the model predicts $\gamma < 0$, $\alpha < 0$ (since λ, $\theta > 0$), and $\beta > 0$. Moreover, while persistent supply shocks also 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 in forecasting inflation.

The orthogonality conditions implied by the rational expectations hypothesis makes the generalized method of moments (GMM) a natural candidate for the estimation of

TABLE 2

REACTION FUNCTION AND POLICY PREFERENCE ESTIMATES, FULL SAMPLE, 1960(I)–2005(II)[a,b]

Instruments	π^* (1)	α (2)	β (3)	p-values (4)
(1)	2.36**	0.15	0.04**	F-stat: 0.00/0.00
	(0.24)	(0.12)	(0.017)	J(7): 0.06
(2)	2.40**	0.15	0.05**	F-stat: 0.00/0.00
	(0.19)	(0.11)	(0.02)	J(7): 0.09

[a]Specification: $\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t$.
[b]Standard errors using a four-lag Newey–West covariance matrix are reported in parantheses. Inflation is measured as changes in the GDP chain-type price index and output gap is obtained from the CBO. Instrument set (1) includes a constant, three lags of inflation, an output gap and a squared output gap. 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 over-identifying restrictions, which is distributed as a $\chi^2(m)$ under the null hypothesis of valid over-identifying restrictions.
Superscripts ** and *denote rejection of the null hypothesis that the true coefficient is zero at the 5% and 10% significance levels, respectively.

(22). Instrumental variable estimators also have the advantage that no arbitrary restrictions need to be imposed on the information set used by private agents to form expectations. Three lags—inflation, output gap and squared output gap—are used as instruments corresponding to a set of seven over-identifying 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 over-identifying restrictions are imposed. In addition, I check the robustness of these results to changes in the instrument set; in particular, I 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.[4]

An important empirical issue that must be confronted is the fact that the theoretical model may lack important dynamics. Although serial correlation in the error terms influences the efficiency it does not affect the consistency of GMM estimates. To control for the possibility of an efficiency problem, I follow Hansen (1982) and use an optimal weighting estimate of the covariance matrix that accounts for both serial correlation and heteroscedasticity in the error terms. Hence all tables report robust standard errors. The qualitative results are unchanged if the empirical reaction function is augmented with lagged values of inflation. (The estimates are not reported and are available upon request.)

Table 2 reports the estimates of the feedback rule (22) 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. Estimates of the slope coefficients as well as those of the implicit inflation target are robust to the instrument selection, and the hypothesis of valid overidentifying restrictions is never rejected. These results show that the finding in Ruge-Murcia (2003) of a significant role for asymmetric preference is robust to using a different method of estimation and a different measure of real activity.

TABLE 3
REACTION FUNCTION AND POLICY PREFERENCE ESTIMATES, SUB-SAMPLES[a,b]

	Instruments				
	π^* (1)	α (2)	β (3)	γ (4)	p-values (5)
1960(I)–1979(III)					
(1)	3.81**	− 0.56**	0.13**	− 0.45**	F-stat: 0.00/0.00
	(0.64)	(0.18)	(0.06)	(0.19)	J(7): 0.13
(2)	3.81**	− 0.60**	0.16**	− 0.54**	F-stat: 0.00/0.00
	(0.43)	(0.14)	(0.05)	(0.14)	J(7): 0.31
Sample 1982(IV)–2005(II)					
(1)	1.99**	− 0.14*	0.01	− 0.13	F-stat: 0.00/0.00
	(0.12)	(0.07)	(0.01)	(0.20)	J(7): 0.11
(2)	1.89**	− 0.16*	0.01	− 0.13	F-stat: 0.00/0.00
	(0.11)	(0.09)	(0.01)	(0.23)	J(7): 0.11

[a]Specification: $\pi_t = \pi^* + \alpha y_t + \beta y_t^2 + v_t$.
[b]Standard errors using a four-lag Newey–West covariance matrix are reported in Parentheses. Inflation is measured as changes in the GDP chain-type price index and output gap is obtained from the CBO. Instrument set (1) includes a constant, three lags of inflation, an output gap and a squared output gap. 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 over-identifying restrictions, which is distributed as a $\chi^2(m)$ under the null hypothesis of valid over-identifying restrictions.
Superscripts ** and * denote rejection of the null hypothesis that the true coefficient is zero at the 5% and 10% significance levels, respectively.

Table 3 reports the estimates for the pre- and post-Volcker regimes. I removed from the second sub-sample the period 1979(III)–1982(III), when the temporary switch in the Fed operating procedure documented by Bernanke and Mihov (1998) can be held responsible for the failure to gain control over inflation. Results are robust to letting the first sub-sample begin in the first quarter of 1970, which corresponds to the starting date chosen by Ruge-Murcia (2003).

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 together with positive and significant parameters for y_t^2. The point estimate of the implicit inflation target is 3.81% while the asymmetric preference parameter is negative and statistically significant. These results sharply contrast with the post-1979 values, which are displayed in the bottom rows of Table 3. Indeed, not only do the implicit inflation target decline significantly to values around 2%, but also the impact of the output gap level on inflation appears weaker, although it is still statistically different from zero.

To the extent that the parameters representing the structure of the economy have remained stable during the last forty years, a smaller value of α can be rationalized only with a decline in λ, corresponding to a more conservative monetary policy stance.[5] The most dramatic difference between the two regimes, however, emerges on the squared output gap, which loses explanatory power over the post-1979 period for both sets of instruments. The latter result translates into values of the policy parameter γ not statistically different from zero.

Turning to the measure of *asymmetric preference-induced time inconsistency*, Table 4 reports the estimates of the average inflation bias. According to equation (11), the bias is

TABLE 4
THE AVERAGE INFLATION BIAS[a]

Instruments	Inflation bias (1)	Inflation target (2)	Inflation bias + inflation target (3)	Inflation mean (4)
1960(I)–1979(III)				4.47
(1)	0.82**	3.81**	4.63**	
	(0.40)	(0.64)	(0.48)	
(2)	1.07**	3.81**	4.88**	
	(0.32)	(0.43)	(0.32)	
Sample 1982(IV)–2005(II)				2.51
(1)	0.04	1.99**	2.03**	
	(0.04)	(0.12)	(0.12)	
(2)	0.04	1.89**	1.93**	
	(0.05)	(0.11)	(0.13)	

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

a convolution of the structural parameters of the model and the variance of the output gap. Given the decline in the latter reported in column (2) of Table 1, we would also expect a reduction in the inflation bias moving from the pre- to the post-Volcker period.[6]

It is worth emphasizing that the figures in Table 4 are based on the discretionary feedback rule estimated in Table 3. The difference between the commitment and discretion solutions implied by the theoretical model is used in Table 4 as a simple metric to quantify any possible time inconsistency of US monetary policy.

Column (1) of Table 4 shows the measure of the average inflation bias. The first block reports the pre-Volcker values, with point estimates ranging from 0.82% in the baseline case to 1.07% for the alternative instrument set. In contrast, the inflation bias is not found to be 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 period but not during the latter. Empirical support for this form of regime shift can also be found in the cross-country evidence reported by Cukierman and Gerlach (2003).

The actual 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 the inflation mean during the post-Volcker era appears 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 existing over the last two decades.

IV. CONCLUDING REMARKS

This paper has developed a method to measure the time inconsistency of monetary policy when the preferences of the central bank are asymmetric. As shown 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 its potential level. 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, we show that the 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 implicit target is 3.81% and the bias is about 1% during the pre-1979 policy regime. In contrast, over the last two decades the inflation target has declined to values close to 2% while the average inflation bias has disappeared. This result can be rationalized by the fact that the policy preference on output stabilization is large and asymmetric before but not after the appointment of Paul Volcker as Fed chairman. Although other factors, including unfavourable supply shocks, are also likely to have played a role, this paper provides empirical support and quantitative measures for a new, additional explanation of the great inflation of the 1960s and 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 a state-contingent bias as well as a stabilization bias. As shown by Svensson (1997) and Cukierman (2002), however, the average inflation bias would then be larger than it would be if a standard expectations-augmented Phillips curve were used. This suggests not only that the present estimates are better interpreted as a lower bound, but also that a richer specification of private agents' behaviour may account for the gap between the model-based average inflation and the actual inflation mean 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.

APPENDIX A: OPTIMAL MONETARY POLICY WHEN THE SUPPLY CURVE IS NONLINEAR

Consider the linex objective function in the main text and the following form for the supply curve:

$$y_t = \frac{\theta(\pi_t - \pi_t^e)}{1 - \theta\tau(\pi_t - \pi_t^e)} + u_t. \tag{A1}$$

A similar specification is employed by Dolado *et al.* (2005). The parameter τ governs any potential nonlinearity in the structure of the economy as it allows the slope of the aggregate supply curve to be steeper at a higher level of inflation and output gap. Linearity is recovered when $\tau = 0$, which yields the linear supply curve relation in the main text.

Commitment

The first-order conditions with respect to π_t and π_t^e are, respectively,

$$\frac{1}{\gamma} E_{t-1}\{\exp[\gamma_\pi(\pi_t - \pi^*)] - 1\} + \frac{\lambda\theta}{\gamma}\left[E_{t-1}\left(\frac{(\exp(\gamma y_t) - 1)}{(1 - \theta\tau(\pi_t - \pi_t^e))^2}\right) - \mu\right] = 0, \tag{A2}$$

$$-\frac{\lambda\theta}{\gamma} E_{t-1}\left(\frac{(\exp(\gamma y_t) - 1)}{(1 - \theta\tau(\pi_t - \pi_t^e))^2}\right) + \mu = 0. \tag{A3}$$

Substituting (A3) into (A2), and following the steps in Section I(a), we obtain that the inflation mean equals the target:

$$E(\pi_t) = \pi^*,$$

implying that nonlinearity in the supply curve does not affect the solution under commitment.

Discretion

When the monetary authorities take private sector's expectations as given, the optimal control problem implies a sole first-order condition with respect to inflation:

$$\text{(A4)} \qquad \frac{1}{\gamma} E_{t-1}\{\exp[\gamma_\pi(\pi_t - \pi^*)] - 1\} + \frac{\lambda\theta}{\gamma}\left[E_{t-1}\left(\frac{(\exp(\gamma y_t) - 1)}{(1 - \theta\tau(\pi_t - \pi_t^e))^2}\right)\right] = 0$$

Notice that, under the assumption that the monetary authorities control inflation perfectly, the equilibrium of the model is characterized by $\pi_t = \pi_t^e$. Therefore implies that the solution under discretion is independent from τ, and therefore a nonlinear supply curve such as (23) is not capable of generating a relationship between the level of inflation and the square of the output gap.

APPENDIX B: IDENTIFICATION

This section illustrates identification problems arising from asymmetric preferences on inflation. To this end, it is helpful to linearize the exponential terms in (17):

$$\text{(A5)} \qquad (\pi_t - \pi^*) + \frac{\gamma_\pi}{2} E_{t-1}\,(\pi_t - \pi^*)^2 + \lambda\theta E_{t-1}(y_t) + \frac{\lambda\theta\gamma_y}{2} E_{t-1}\,(y_t)^2 + e_t = 0,$$

with e_t being the remainder of the approximation. Solving (A5) for current inflation and rearranging terms, we obtain the following *nonlinear* reaction function in the asymmetric preferences parameters:

$$\text{(A6)} \qquad \pi_t = \left(\frac{1}{1 - \gamma_\pi\pi^*}\right)\left[\pi^*\left(1 - \frac{\gamma_\pi\pi^*}{2}\right) - \lambda\theta E_{t-1}(y_t) - \frac{\gamma_\pi}{2} E_{t-1}\,(\pi_t)^2 - \frac{\lambda\theta\gamma_y}{2} E_{t-1}\,(y_t)^2 - e_t\right].$$

The latter expression shows that, if the feedback rule is augmented with squared inflation, the asymmetric preference parameters and the inflation target cannot be identified separately unless additional restrictions are imposed (see Boinet and Martin 2005; Surico 2007). Equation (20) nevertheless implies that, even if γ_π is non-zero, the sign of the coefficient on output gap-squared could identify the relative importance of the asymmetric preferences on inflation and output. In particular, a positive sign implies that the asymmetries on the output gap are dominant.

ACKNOWLEDGEMENTS

Many thanks to the editor Alan Manning and two anonymous referees for very helpful suggestions. The views expressed in this paper are my own, and do not necessarily reflect those of the Bank of England.

NOTES

1. Additional references include Chadha and Schellekens (1999), Ruge-Murcia (2003) and Surico (2003).
2. Interested readers are referred to Surico (2007) for estimates of a policy rule with statistically insignificant coefficients associated with asymmetric preferences on inflation.
3. It should be noted that under this more general specification of the objective function the inflation target and the asymmetric preference parameters cannot be identified separately. I show in Appendix B that it is nevertheless possible to infer the relative importance of output and inflation asymmetries from the sign of the coefficient on the squared output gap.
4. As an additional check for weak instruments, I also perform the multiple endogenous regressor analog of the *F*-test from the first-stage regression. The statistics exceed the critical values, tabulated in Stock and

Yogo (2003), which ensures that the bias is no more than 10% of the inconsistency of OLS and that a 5% hypothesis test is rejected no more than 15% of the time. Given the conservative nature of the test, it appears that the instruments can be deemed strong in this application.

5. Ahmed *et al*. (2004), Primiceri (2005), and Sims and Zha (2006) find in fact that most of the decline in the volatility of output since the mid-1980s can be attributed to a decline in the volatility of the shocks rather than to a change in the parameters of the structure of the economy.
6. This would be consistent with the historical decline in the volatility of supply shocks documented by Hamilton (1996), Ahmed *et al*. (2004), Primiceri (2005) and Sims and Zha (2006), among others.

REFERENCES

AHMED, S., LEVIN, A. T. and WILSON, B. A. (2004). Recent US macroeconomic stability: good policies, good practices, or good luck? *Review of Economics and Statistics*, **86**, 824–32.

BARRO, R. J. and GORDON, D. (1983). A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, **91**, 589–610.

BERNANKE, B. and MIHOV, I. (1998). Measuring monetary policy. *Quarterly Journal of Economics*, **63**, 869–902.

BLINDER, A. (1997). Distinguished lecture on economics and government: what central bankers could learn from academics and vice versa. *Journal of Economic Perspectives*, **11**, 3–19.

——— (1998). *Central Banking in Theory and Practice*. Cambridge, Mass.: MIT Press.

BOINET, V. and MARTIN, C. (2005). Targets, zones and asymmetries: A flexible nonlinear model of recent UK monetary policy. Mimeo: Brunel University.

CHADHA, J. S. and SCHELLEKENS, P. (1999). Monetary policy loss functions: two cheers for the quadratic. DAE Working Paper no. 99/20.

CLARIDA, R., GALÍ, J. and GERTLER, M. (1999). The science of monetary policy: a New Keynesian perspective. *Journal of Economic Literature*, **27**, 1661–707.

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.

——— and GERLACH, S. (2003). The Inflation Bias Revisited: Theory and Some International Evidence. The Manchester School, **71**, 541–565.

——— and MUSCATELLI, A. (2003). Do Central Banks have Precautionary Demands for Expansions and for Price Stability? Theory and Evidence. Mimeo: Tel-Aviv University. April.

DOLADO, J. J., MARÍA-DOLORES, R. and RUGE-MURCIA, F. J. (2004). Nonlinear monetary policy rules: some new evidence for the US. *Studies in Nonlinear Dynamics and Econometrics*, **8**(3).

———, MARÍA-DOLORES, R. and NAVEIRA, M. (2005). Are monetary-policy reaction functions asymmetric? The role of nonlinearity in the Phillips curve. *European Economic Review*, **49**, 485–503.

GALÍ, J., GERTLER, M. and LOPEZ-SALIDO, J. D. (2003). Markups, gaps, and the welfare costs of business fluctuations. Mimeo: Universitat Pompeu Fabra, October.

GERAATS, P. (1999). Inflation and its variation: an alternative explanation. CIDER Working Paper no. C99–105.

HAMILTON, J. D. (1996). That is what happened to the oil price–macroeconomy relationship. *Journal of Monetary Economics*, **38**, 215–20.

HANSEN, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, **50**, 1029–54.

IRELAND, P. N. (1999). Does the time consistency problem explain the behavior of US inflation? *Journal of Monetary Economics*, **44**, 279–92.

——— (2006). Changes in the Federal Reserve's inflation target: causes and consequences. Mimeo: Boston College.

KAHNEMAN, D. and TVERSKY, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, **47**, 263–92.

KING, R. G. (2000). The new IS–LM model: language, logic, and limits. *Federal Reserve Bank of Richmond Economic Quarterly*, **86** (3), 45–103.

KYDLAND, F. and PRESCOTT, E. (1977). Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy*, **85**, 473–90.

MARTIN, C. and MILAS, C. (2004). Modelling monetary policy: inflation targeting in practice. *Economica*, **71**, 209–21.

MCCALLUM, B. T. (1997). Crucial issues concerning central bank independence. *Journal of Monetary Economics*, **39**, 99–112.

NOBAY, R. and PEEL, D. (2003). Optimal discretionary monetary policy in a model of asymmetric central bank preferences. *Economic Journal*, **113**, 657–65.

PERSSON, T. and TABELLINI, G. (1999). Political economics and macroeconomic policy. In J. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*. Amsterdam: North-Holland.

PRIMICERI, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, **72**, 821–52.

ROGOFF, K. (1985). The optimal degree of commitment to a monetary target. *Quarterly Journal of Economics*, **100**, 1169–90.

RUGE-MURCIA, F. J. (2003). Does the Barro–Gordon model explain the behavior of US inflation? A re-examination of the empirical evidence. *Journal of Monetary Economics*, **50**, 1375–90.

SIMS, C. A. and ZHA, T. (2006). Were there regime switches in US monetary policy? *American Economic Review*, **96**, 54–81.

STAIGER, D. and STOCK, J. (1997). Instrumental variables regression with weak instruments. *Econometrica*, **65**, 557–86.

STOCK, J. and YOGO, M. (2003). Testing for weak instruments in linear IV regression. Mimeo: Harvard University, February.

SURICO, P. (2003). Measuring the time inconsistency of US monetary policy. ECB Working Paper no. 291, December.

——— (2007). The Fed's monetary policy rule and US inflation: the case of asymmetric preferences. *Journal of Economic Dynamics and Control*, **31**, 305–324.

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 S. E. Feinberg and A. Zellner (eds.), *Studies in Bayesian Economics in Honour of L. J. Savage*. Amsterdam: North Holland.

ZELLNER, A. (1986). Bayesian estimation and prediction using asymmetric loss functions. *Journal of the American Statistical Association*, **81**, 446–51.