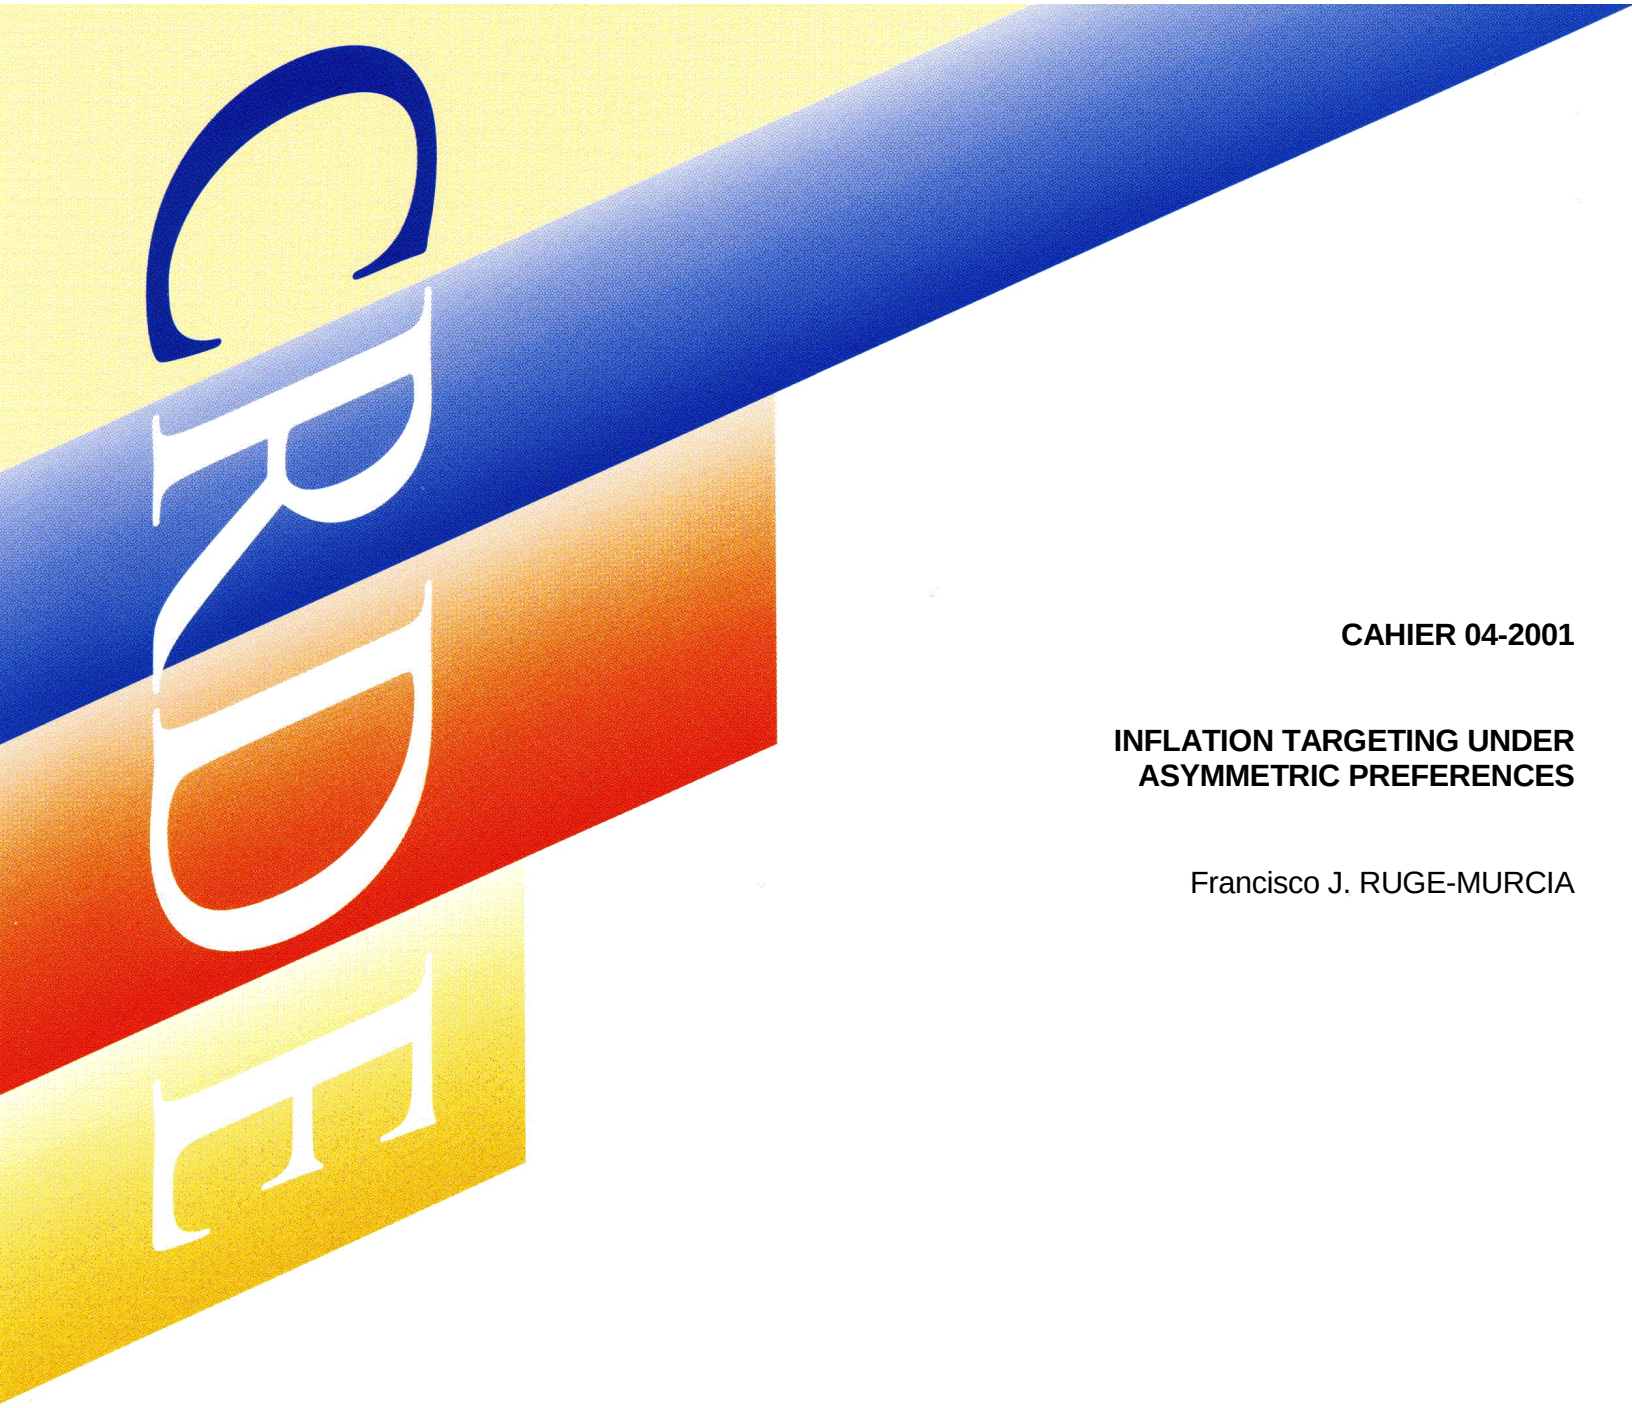

CAHIER 04-2001

INFLATION TARGETING UNDER ASYMMETRIC PREFERENCES

Francisco J. RUGE-MURCIA

Centre de recherche
et développement en économique

C.P. 6128, succursale Centre-ville
Montréal QC H3C 3J7

Téléphone : (514) 343-6557
Télécopieur : (514) 343-5831
crde@crde.umontreal.ca
http://www.crde.umontreal.ca/

Université de Montréal

CAHIER 04-2001

INFLATION TARGETING UNDER ASYMMETRIC PREFERENCES*

Francisco J. RUGE-MURCIA[1]

[1] Centre de recherche et développement en économique (C.R.D.E.) and Département de sciences économiques, Université de Montréal

February 2001

---

* This paper draws on preliminary research carried out jointly with Hashem Pesaran. The author is grateful to him for graciously allowing the use of that material here. A substantial part of this project was carried out while the author was visiting research fellow at the Bank of Spain in the spring of 2000. He has received helpful comments by Rui Castro, Larry Christiano, David Johnson, Athanasios Orphanides, and participants in seminars at Carleton University, the Bank of Spain, the Canadian Macroeconomics Study Group (London, November 1999), and the Winter Meetings of the Econometric Society (New Orleans, January 2001). All data and programs employed in this project are available from the author upon request. Financial support from the Social Sciences and Humanities Research Council and the Fonds pour la Formation de chercheurs et l'aide à la recherche is gratefully acknowledged.

RÉSUMÉ

Cet article développe et estime un modèle de la théorie des jeux qui a pour objet de cibler les taux d'inflation où les préférences de la banque centrale sont asymétriques autour du taux ciblé. Plus particulièrement, les déviations positives du taux ciblé peuvent avoir une plus grande ou une plus faible pondération dans la fonction de perte de la banque centrale que les déviations négatives. Il a été démontré que certains des résultats antérieurs, dérivés sous la présomption de la symétrie, ne sont pas robustes à la généralisation des préférences. Les estimations des paramètres de préférences pour les banques centrales du Canada, de la Suède et du Royaume-Uni sont différentes statistiquement de celles qui découlent de la fonction de perte quadratique utilisée habituellement. Les résultats économétriques sont robustes à différents modèles prévisionnels en ce qui a trait au taux de chômage, mais ne le sont pas face aux mesures d'inflation qui dépassent l'éventail ciblé.

Mots clés : cibles d'inflation, préférences asymétriques, crédibilité

ABSTRACT

This paper develops and estimates a game-theoretical model of inflation targeting where the central banker's preferences are asymmetric around the targeted rate. In particular, positive deviations from the target can be weighted more, or less, severely than negative ones in the central banker's loss function. It is shown that some of the previous results derived under the assumption of symmetry are not robust to the generalization of preferences. Estimates of the central banker's preference parameters for Canada, Sweden, and the United Kingdom are statistically different from the ones implied by the commonly used quadratic loss function. Econometric results are robust to different forecasting models for the rate of unemployment but not to the use of measures of inflation broader than the one targeted.

Key words : inflation targets, asymmetric preferences, credibility

1 Introduction

This paper develops and estimates a game-theoretical model of monetary policy where the central banker's preferences are asymmetric around the targeted rate of price inflation. The preference specification permits different weights for positive and negative inflation deviations from the target, and includes as a special case the quadratic loss function employed by previous literature.[1] The symmetry embodied in the quadratic form means that the loss associated with an inflation deviation from the target depends solely on its magnitude and not on its sign. Instead, under asymmetric preferences both the magnitude and sign of a deviation matter to the policy maker, and uncertainty can induce a prudent behavior on the part of the central banker.

Arguments in favor of the quadratic loss function include that it is tractable, yields transparent analytical results, and might well provide a reasonable approximation to the central banker's preferences. On the other hand, recent anecdotal and empirical evidence appears consistent with the notion of asymmetric preferences. For example, Pesaran and Ruge-Murcia (1999) show that realignment probabilities in exchange rate target zones are asymmetric around the central parity. Goldfajn and Valdés (1999) examine more than 200 episodes of currency overvaluations and undervaluations, and document that while the former usually finish with sudden changes in the nominal exchange rate, the latter end by the smooth adjustment of prices and wages. Clarida and Gertler (1999) estimate a reaction function for the Bundesbank and find that it raises the day-to-day interest rate when inflation is above its steady-state trend value but barely responds when it is below. Ruge-Murcia (forthcoming) derives implicit bounds for the Canadian inflation target zone using data on market-determined nominal interest rates. Results indicate that while the public might perceive the band to be of approximately the same width as announced, it appears asymmetrically distributed around the official target.

The hypothesis of asymmetric central bank preferences is examined here in the situation where the targeted rate of inflation is publicly announced.[2] Publicly-announced inflation targets have been adopted by a number of countries as a framework for the conduct of monetary policy (for example, New Zealand, Australia, Canada, the United Kingdom, and Sweden). Under this arrangement, the central bank commits itself to gear monetary policy towards keeping a measure of annual inflation close to an explicit rate. Bernanke *et al.* (1999) review the basic features of this framework and the recent experience of various inflation-targeting countries. Not only is inflation targeting interesting on its own right, but it has the advantage that the value around which the central banker's loss function is defined is observable by the econometrician. This permits the comparison between the stated policy goal and the inflation realizations, simplifies the estimation strategy, reduces the number of parameters to be estimated, and, as we will below, means that the model generates testable empirical implications, regardless of the whether the target is the socially optimal or not.

[1] See, among others, Kydland and Prescott (1977), Barro and Gordon (1983), Rogoff (1985), Walsh (1995), Svensson (1999), Clarida, Galí, and Gertler (1999), and Eijffinger, Hoeberichts, and Schaling (2000). In the inflation targeting literature, see Green (1996), Svensson (1997), Beetsma and Jensen (1998), and Muscatelli (1999).

[2] For a model where the inflation target is unobserved and possibly time-varying, see Tetlow (1999).

The analysis is carried out in a standard game-theoretical framework. However, in contrast to previous literature, this paper derives and tests the model's empirical predictions[3] and obtains estimates of the central banker's preference parameters. The functional form of the policy maker's loss function is based on the linex model proposed by Varian (1974) and examined by Zellner (1992), Granger and Pesaran (1996), and Christoffersen and Diebold (1997) in the context of optimal forecasting.. This function is attractive because it is analytically tractable, yields a closed-form solution, nests the usual quadratic loss function as a special case, and generates sharp empirical predictions.

In related research, Nobay and Peel (1998) use the linex loss function to study optimal commitment and discretion in monetary policy. This paper extends their analysis by fully characterizing the theoretical and empirical implications of the model. In particular, the properties of the central banker's reaction are derived, conditions for the existence and uniqueness of the Nash equilibrium are established, and the government's problem of optimal delegation is formulated and solved. Rather than using simulations (as Nobay and Peel), this paper formulates and tests the empirical predictions of a model with asymmetric preferences, and examines to what extend they might be a feature in real-world policy making.

As related models of monetary policy [*e.g.,* Barro and Gordon (1983), Rogoff (1985), Svensson (1997), Clarida *et al.* (1999)], the asymmetric model predicts a positive relation between inflation and unemployment (or a negative relation between inflation and output). However, in contrast to the quadratic model, this relation is nonlinear and concave.

While the quadratic model predicts that inflation should be on average above its target [see Green (1996) and Svensson (1997)], under asymmetric preferences, inflation can be on average above or below the target depending on the central banker's preferences. To understand the intuition of this result, recall that relaxing the assumption of quadratic preferences means that certainty equivalence no longer holds. Then, the expected marginal cost of departing from the inflation target is nonlinear in inflation. In particular, when the central banker associates a larger loss to positive than negation inflation deviations from the target, uncertainty raises the expected marginal cost and induces prudent behavior on the part of the monetary authority. This the prudence motive can be large enough to deliver a below-target average rate of inflation.

The asymmetric model also implies that the conditional variance of inflation is helpful in forecasting its mean. The preference parameter that measures the asymmetry is statistically identified as the coefficient of the conditional variance of inflation and from the nonlinear term on unemployment. Because the quadratic model corresponds to the special case where this parameter is zero, it is possible to test the null hypothesis of quadratic preferences against the well-defined alternative of asymmetric preferences. In an empirical application to Canada, Sweden, and the United Kingdom, results support the notion of asymmetric preferences in the form a positive and statistically significant estimate of the asymmetry preference parameter. These results appear robust to modeling the rate of unemployment as an stationary or unit-root process, but not to the use of measures of inflation broader than the one targeted.

The paper is organized as follows: section 2 describes the economy and central banker's

[3] An exception is Ireland (1999) who tests the predictions of the Barro-Gordon model using U.S. data.

preferences, derives conditions for the existence and uniqueness of the Nash equilibrium, and outlines the empirical implications of the model; section 3 constructs a reduced form version of the model, reports empirical estimates, examines the robustness of the results, and explores other theories that might account for the econometric results; and section 4 concludes.

2 The Model

The model consists of (*i*) a central banker that implements monetary policy, and (*ii*) a continuum of identical individuals assumed to construct their expectations rationally. In what follows, this continuum will be referred to as the public.

2.1 The Economy

The economic behavior of the private sector is summarized by an expectations-augmented Phillips curve:[4]

$$u_t = u_t^n - \lambda(\pi_t - \pi_t^e) + \eta_t, \qquad \lambda > 0, \tag{1}$$

where u_t is the rate of unemployment, u_t^n is the natural rate of unemployment, π_t is the rate of inflation, π_t^e is the public's forecast of inflation at time t constructed at time $t-1$, and η_t is a supply disturbance. The choice of unemployment as measure of real economic activity has no effect on the analytical derivations below, but it does allow the comparison of empirical results in section 3 with the theoretical predictions in earlier literature.

Under the assumption of rational expectations, the public's (subjective) inflation forecast is the objective conditional expectation:

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

where E_{t-1} is the expectation conditional on all information available at time $t-1$. The public's information set is denoted by I_t and is assumed to contain all model parameters, current and past realizations of the variables, and current and next-period inflation targets. The assumption that future inflation targets form part of I_t follows from the fact that in actual practice, the inflation targets are preannounced.

The natural rate of unemployment is assumed to evolve over time according to:

$$\Delta u_t^n = \psi - (1-\delta)u_{t-1}^n + \sum_{i=1}^{q-1} \theta_i \Delta u_{t-i}^n + \zeta_t, \tag{3}$$

where ζ_t denotes the unpredictable component of the natural rate. This specification includes as special cases the stationary ($0 \leq \delta < 1$) and unit root ($\delta = 1$) models of the natural rate employed in earlier literature. For example, the model in Barro and Gordon (1983) can be obtained by setting $\theta_i = 0$ for $i = 1, \ldots, q-1$, the one in Ireland (1999) by setting $\delta = 1$, $\psi = 0$, and $\theta_i = 0$ for $i = 2, \ldots, q-1$, and the one in Cukierman (2000) by setting $\delta = 0$,

[4] For the microfoundations of this view of aggregate supply, see Lucas (1972, 1973).

and $\theta_i = 0$ for $i = 1, \ldots, q-1$. The intention here is to adopt a very general time series specification for the natural rate and then examine to what extend results are robust to the use of different forecasting models for u_t^n.[5] The constant term in (3) could be written as $\psi(1-\delta)$ in order to rule out the presence of a drift when unemployment is $I(1)$. However, for the estimation of the model, it is desirable to allow a nonzero intercept. Standard procedures can then be used to test whether this term is significantly different from zero or not. Finally, in order to exclude the possibility of u_t^n being an $I(2)$ process, all roots of the polynomial $1 - \sum_{i=1}^{q-1} \theta_i z^i$ are assumed to lay outside the unit circle.

Modeling the natural rate of unemployment as time-varying is important for two reasons. *First*, it seems plausible that changes in technology, labor force demographics, unionization rates, and welfare benefits could affect the labor market and generate movements in the natural rate. For the United States, Weiner (1993), Tootell (1994), and Staiger, Stock, and Watson (1997) report evidence that the natural rate has changed during the postwar period. Shimer (1998) argues that in the absence of the baby boom, the rate of unemployment would neither have increased in the 1960's and 1970's, nor fallen afterwards. *Second*, although the assumption that u_t^n is constant can be innocuous for certain theoretical results, it does lead to the empirical prediction that realized unemployment is white noise. However, for most countries and sample periods the rate of unemployment is serially correlated and substantially persistent. As an example, table 1 presents the first 10 autocorrelations of the quarterly rate of unemployment in three inflation-targeting countries, namely Canada, Sweden, and the United Kingdom. In all cases, the correlation between current and previous unemployment (*i.e.,* the first autocorrelation) is above 0.90, and the correlation between current unemployment and its realization two years before (*i.e.,* the eight autocorrelation) is above 0.55.

As in Barro and Gordon (1983) and Ireland (1999), the process of u_t^n is not affected by the monetary policy instrument (see below) or lagged unemployment rates. This assumption reflects the view that the natural rate is primarily determined by factors outside the scope of monetary policy (for example, exogenous demographic or technological variables). In this sense, this paper adopts Friedman's interpretation of the natural unemployment rate as "the level that would be ground out by the Walrasian system of general equilibrium equations, provided there is imbedded in them the actual structural characteristics of the labor and commodity markets, including market imperfections, stochastic variability in demands and supplies, the cost of gathering information about job vacancies and labor availability, the costs of mobility, and so on" [Friedman (1968, p. 8)].

An alternative process for the natural rate is proposed by Lockwood and Philippopoulos (1994). Their specification also seeks to capture the observed persistence in the unemployment rate, and consists of a strictly stationary process for the natural rate where the first unemployment lag is one of the regressors. Unfortunately, when coupled with the more general preference specification employed here, the model has no closed-form solution, the value function needs to be approximated by a finite-order polynomial, and the uniqueness of the Nash equilibrium cannot be insured. Since (3) also captures the serial correlation in un-

[5] In principle, one could also consider a nonlinear formulation of the natural rate. However, the construction of the reduced-form version of the model in section 3.3 does make use of the assumption of linearity.

employment but in a more tractable manner, and permits both stationary and nonstationary specifications, I did not pursue their approach further.

The policy maker is assumed to affect the rate of inflation through a policy instrument (as in Walsh (1995, 1998 ch. 8)). The instrument, whether a monetary aggregate or a short-term nominal interest rate, is imperfect, in the sense that in a stochastic world it cannot determine inflation completely and its effect takes place with a one-period lag:

$$\pi_t = f(i_t) + \epsilon_t, \tag{4}$$

where $f(\cdot)$ is a monotonic, continuous, and differentiable function, i_t is the policy instrument chosen at time $t-1$, and ϵ_t is a control error that represents imperfections in the conduct of monetary policy. This simple specification serves two purposes. *First*, it relaxes the usual assumption that the monetary authority chooses directly the rate of inflation after observing (before the public does) the random shocks. In contrast, the policy maker here has no informational advantage over the public since neither of them observe at time $t-1$ the realization of the disturbances at time t. Because there is no private information in the model, the central banker's information set coincides with the public's and is also given by I_t. *Second*, (4) introduces an additional structural disturbance to the model and permits the derivation of time-series predictions on the joint probability distribution function of inflation and unemployment.[6]

Finally, to complete the description of the economy, define $\boldsymbol{\xi}_t$ to be the 3×1 vector that contains all the model's structural disturbances at time t. It is assumed that $\boldsymbol{\xi}_t$ is serially uncorrelated and normally distributed with zero mean and constant conditional variance-covariance matrix:

$$\boldsymbol{\xi}_t | I_{t-1} = \left[\begin{array}{c} \eta_t \\ \zeta_t \\ \epsilon_t \end{array} \right] \Bigg| \; I_{t-1} \quad \sim N(\mathbf{0}, \Omega),$$

where Ω is a 3×3 positive-definite variance-covariance matrix.[7] The assumption of normality is crucial for obtaining an analytical solution of the model, but section 3.5 examines the robustness of empirical results to shocks drawn from a different distribution. Note that for the moment, no restrictions are imposed on the off-diagonal elements of Ω.

2.2 The Central Banker

Asymmetric disappointment-aversion preferences whereby individuals or households treat differently expected gains and losses have been proposed in the literature by Epstein and Zin (1989) and Gul (1991). In the case of a government or central bank, their different attitude *vis a vis* booms and recessions, and the more downward- than upward-rigidity in prices might

[6] Previous literature frequently assumes that the only random shock is the aggregate supply disturbance and, consequently, the dynamics of π_t and u_t are driven by exactly the same random term. As an alternative to (4), one could postulate an aggregate demand relation [see, for example, Orphanides and Wilcox (1996)]. In this case, the model solution is unchanged but the structural interpretation of the reduced-form disturbances is slightly different.

[7] Some evidence supporting the assumption that Ω is conditionally homoskedastic is presented in section 3.4.

plausibly induce asymmetries in their loss function. Consistent with this idea are empirical findings by Goldfajn and Valdés (1999), Ruge-Murcia (forthcoming), and Clarida and Gertler (1997). Goldfajn and Valdés examine more than 200 episodes of currency overvaluations and undervaluations, and document that while the former usually finish with sudden changes in the nominal exchange rate, the latter end by the smooth adjustment of prices and wages. Ruge-Murcia derives implicit bounds for the Canadian inflation target zone using market-determined nominal interest rates. Results indicate that while the public might perceive the band to be of approximately the same width as announced, it appears asymmetrically distributed around the official target. Clarida and Gertler estimate a reaction function for the Bundesbank and find that it raises the day-to-day interest rate when inflation is above its steady-state trend value but barely responds when it is below.[8]

In order to examine formally the hypothesis that the central banker's preferences are asymmetric over inflation outcomes, generalize the usual quadratic loss function as:

$$C(\pi_t, u_t) = (\exp(\alpha(\pi_t - \tilde{\pi}_t)) - \alpha(\pi_t - \tilde{\pi}_t) - 1)/\alpha^2 + (\phi/2)(u_t - \tilde{u}_t)^2, \qquad \alpha \neq 0, \quad (5)$$

where $C : \Re^2 \rightarrow \Re$ is a continuous and differentiable function that represents the central banker's preferences, $\tilde{\pi}_t$ is the inflation target, $\tilde{u}_t$ is the unemployment target, and ϕ is a positive coefficient that measures the relative importance of unemployment stabilization. In the specific institutional setup considered here, the inflation target is assigned to the monetary authority by the government prior to time t. Since the target is announced publicly beforehand, it follows that $\tilde{\pi}_t \in I_{t-1}$. A general model for $\tilde{u}_t$ is presented below. In contrast to previous literature where both components of the policy maker's loss function are quadratic, the inflation component in (5) is given by the linex function [Varian (1975)]. Zellner (1992), Granger and Pesaran (1996), and Christoffersen and Diebold (1997) examine this function in the context of optimal forecasting.

In order to develop some intuition, the linex function is plotted in figure 1(a) for the special case where $\alpha > 0$. Notice that for rates of inflation above the target (that is, $\pi_t > \tilde{\pi}_t$), the exponential term eventually dominates and the loss associated with a positive deviation rises exponentially. For $\pi_t < \tilde{\pi}_t$, it is the linear term that becomes progressively more important as unemployment decreases and, consequently, the loss rises linearly. This asymmetry can be easily seen by considering, for example, the loss associated with a ± 1 unemployment deviation from $\tilde{\pi}_t$. It is apparent that even though their magnitudes are the same, the -1 deviation delivers a smaller loss than the $+1$ deviation. Hence, positive inflation deviations are weighted more severely than negative ones in the central banker's loss function. The converse holds when $\alpha < 0$.

It is useful to compare the linex function with the very-familiar quadratic loss function in figure 1(b). The quadratic function is symmetric around zero, meaning that positive and negative deviations of the same size yield exactly the same loss. Hence, only the magnitude, and not the sign, of the deviation is important for the policy maker. Applying L'Hôpital's

[8] Although this research is suggestive, it is clearly not a substitute for rigorous microfundations. In future research, I seek to provide such foundations by considering the properties of the social welfare function in an economy where individuals have disappointment-aversion preferences.

rule, it is possible to show that the quadratic function can be obtained as a special case of the linex function when the asymmetry parameter, α, tends to zero. This result is important because it formally proves that the preference specification employed here nests the one assumed in earlier literature as a special case. It also suggests that the hypothesis that the central banker's preferences are quadratic over inflation could be tested econometrically by evaluating whether α is significantly different from zero or not.

The functional form in (5) is attractive because it is analytically tractable, yields a closed-form solution, and generates sharp empirical predictions. In principle, one could extend it to allow asymmetries regarding the rate of unemployment. This possibility is discussed in section 3.5, where it is also shown that the basic predictions of the model are unchanged by such formulation. Empirical evidence by Bai and Ng (1998) suggests that asymmetries in unemployment might be less pronounced than in inflation. Bai and Ng propose a test of conditional symmetry that compares the empirical distribution function of positive and negative shocks. Their results using monthly data indicate that the hypothesis of symmetry around the mean cannot be rejected for the US unemployment rate at the 5% percent level but it is rejected for CPI inflation at the same significance level.[9] In related research, Ruge-Murcia (2000) estimates a model with asymmetric unemployment preferences for the G7 countries. However, estimates of the reduced-form parameters do not support the hypothesis as asymmetric unemployment preferences for Canada and the United Kingdom.

The targeted rate of unemployment is assumed to be proportional to the expected natural rate:

$$\tilde{u}_t = kE_{t-1}u_t^n, \qquad 0 < k \leq 1. \tag{6}$$

Previous literature usually assumes that $\tilde{u}_t$ is strictly less than the natural rate (*i.e.,* $0 < k < 1$). This assumption is based on the notion that distortions in goods and labor markets (for example, monopolistic competition) render the natural rate of unemployment higher than socially optimal. Persson and Tabellini (2000) note that the premise $0 < k < 1$ is crucial in generating an inflation bias in the linear-quadratic framework of Kydland and Prescott (1977) and Barro and Gordon (1983). On the other hand, King (1996) and Blinder (1998) suggests on the basis of institutional evidence, that central bankers actually target the expected natural rate of unemployment (*i.e.,* $k = 1$). Both views are accommodated here by allowing $0 < k \leq 1$. Analytical results obtained under both assumptions are compared below.

2.3 Nash Equilibrium

Consider the problem of a central banker that must choose the sequence of instruments that minimizes the present discounted value of her loss function:

$$\underset{\{i_{t+s}\}_{s=0}^{\infty}}{Min} \qquad E_{t-1}\sum_{s=0}^{\infty}\beta^s C\left(\pi_{t+s}, u_{t+s}\right),$$

[9] Note, however, that the hypothesis of symmetry cannot be rejected at the 5% level for the percentage change in the GDP deflator.

where $0 < \beta < 1$ represents the discount rate. The optimization is made subject to the expectations-augmented Phillips curve [eq. (1)], and takes as given the public's inflation forecasts and the inflation targets. Recall that the natural rate of unemployment is assumed to be determined primarily by factors outside the scope of monetary policy. Hence, the policy instrument does not affect the path of u_t^n and the central banker's objective function can be break down into a sequence of one-period problems. This decomposition simplifies the solution of the model and, as it will be shown below, delivers a unique Nash equilibrium. In contrast, Lockwood and Philippopoulos (1994) show that in a model where the natural rate depends on lagged realizations of the unemployment rate, two stable equilibria are possible with different comparative-statics properties. Only one of the equilibria is intuitive in its economic predictions and corresponds to the one obtained in the static case. However, it is not possible to rule out the second, nonintuitive equilibrium on based theoretical considerations alone.

The first-order necessary condition is

$$E_{t-1}((\partial C/\partial \pi_t)(\partial \pi_t/\partial i_t) + (\partial C/\partial u_t)(\partial u_t/\partial \pi_t)(\partial \pi_t/\partial i_t)) = 0,$$

and is satisfied by the value of i_t that equates the marginal cost of higher unemployment with the marginal benefit lower inflation. Since the objective function is globally convex, this value is unique. Computing the partial derivatives and dividing through by $\partial \pi_t/\partial i_t = \partial f(\cdot)/\partial i_t \neq 0$, yield

$$E_{t-1}((\exp(\alpha(\pi_t - \tilde{\pi}_t)) - 1)/\alpha - \lambda\phi(u_t - kE_{t-1}u_t^n)) = 0.$$

In order to find this conditional expectation two intermediate results are useful. *First*, as it will be shown below, the assumption that the model shocks are normal implies that, conditional on the information set, inflation is also normally distributed. Then, $\exp(\alpha(\pi_t - \tilde{\pi}_t))$ is log normal with mean $\exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t + (\alpha/2)\sigma_\pi^2))$, where σ_π^2 denotes the conditional variance of the rate of inflation. (See below for the derivation of σ_π^2 in terms of the elements of $\boldsymbol{\xi}_t$.) *Second*, the conditional expectation of unemployment can be found by taking E_{t-1} in both sides of (1) to obtain:

$$E_{t-1}u_t = E_{t-1}u_t^n - \lambda(E_{t-1}\pi_t - \pi_t^e).$$

With these results, the first-order condition can be written as

$$(\exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t) + \alpha^2\sigma_\pi^2/2) - 1)/\alpha + \lambda^2\phi(E_{t-1}\pi_t - \pi_t^e) - \lambda\phi(1-k)E_{t-1}u_t^n = 0. \quad (7)$$

In the quadratic model, the first-order condition of the central banker's minimization problem is linear and can be solved explicitly to obtain her reaction function in terms of the public's inflation forecast, π_t^e.[10] To see this, take the limit of (7) as $\alpha \to 0$ and rearrange to obtain

$$E_{t-1}\pi_t = (\tilde{\pi}_t + \lambda^2\phi\pi_t^e + \lambda\phi(1-k)E_{t-1}u_t^n)/(1+\lambda^2\phi). \quad (8)$$

[10] Strictly speaking, the reaction function relates the policy instrument, i_t, and π_t^e, both of which are determined in the previous period. However, in what follows it will be convenient to work with $E_{t-1}\pi_t$ rather than i_t. Since these two variables are monotonically related by the function $f(\cdot)$, this approach entails no loss of generality.

Since $\lambda^2\phi/(1+\lambda^2\phi) \in (0,1)$, the central banker's reaction is linear and monotonically increasing on the public's inflation forecast. It is trivial to show that in this case, the Nash equilibrium always exists and is unique.

In contrast, under asymmetric preferences, the first-order condition only defines the reaction function implicitly:

$$\begin{aligned} & g(E_{t-1}\pi_t, \pi_t^e) \\ & \quad = (\exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t) + \alpha^2\sigma_\pi^2/2) - 1)/\alpha + \lambda^2\phi(E_{t-1}\pi_t - \pi_t^e) - \lambda\phi(1-k)E_{t-1}u_t^n), \quad (9) \\ & \quad = 0. \end{aligned}$$

However, using the implicit function theorem, it is possible to show that

$$\partial E_{t-1}\pi_t/\partial\pi_t^e = \lambda^2\phi/(\lambda^2\phi + \exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t) + \alpha^2\sigma_\pi^2/2)) \in (0,1),$$

for all values of α. Hence, as in the quadratic model, the central banker's reaction is a monotonically increasing function of the public's inflation forecast. Also

$$\begin{aligned} & \partial^2 E_{t-1}\pi_t/\partial(\pi_t^e)^2 = \\ & \qquad -\alpha\lambda^4\phi^2 \exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t) + \alpha^2\sigma_\pi^2/2)/(\lambda^2\phi + \exp(\alpha(E_{t-1}\pi_t - \tilde{\pi}_t) + \alpha^2\sigma_\pi^2/2))^3, \end{aligned}$$

that is less than zero for $\alpha > 0$, equal to zero for $\alpha \to 0$, and larger than zero for $\alpha < 0$. Hence, for $\alpha > 0$ ($\alpha < 0$) the central banker's reaction is a concave (convex) function of π_t^e.

In order to develop further the reader's intuition and to illustrate future theoretical results, it is useful to plot the central banker's reaction function for different values of the preference parameter α. This is done in figure 2 under the assumption that the inflation target is $\tilde{\pi} = 0$, and remaining parameters are $\lambda = 2$, $\phi = 0.5$, $k = 0.8$, $u^n = 5$, and $\sigma_\pi^2 = 2.5^2$. Treating all parameters as fixed, the central banker's reaction was computed by solving numerically the implicit function (9) for 80 equally-spaced values of π^e between -8 and 8. The figure also includes the reaction function of the quadratic central banker in (8), and the public's reaction function that is summarized by the rational expectations relation (2). Graphically, the Nash equilibrium is the point where (9) (or (8)) and (2) intersect.

Notice that although in all cases the central banker's reaction is an increasing function of the public's inflation forecast, her willingness to accommodate its expectations depends on the preference parameter α. Consider first the case where $\alpha < 0$. Note that in this situation the central banker responds to π_t^e at an increasing rate and the inflation rate will be larger than under quadratic preferences. For values of $\alpha \leq -1/(\lambda\phi(1-k)E_{t-1}u_t^n)$, there is no finite rate of inflation at which (9) and (2) intersect, and the Nash equilibrium will not exist [see proposition 1 below]. On the other hand, when $\alpha > 0$ the central banker accommodates the public's inflation forecast at a decreasing rate and the inflation rate will be always smaller than under quadratic preferences. For a large-enough value of α, a deflation bias, whereby inflation systematically undershoots its targeted value, can arise in equilibrium. This result is important because it shows that asymmetric preferences can provide a theoretical foundation for Stanley Fischer's observation that a deflation bias is a possible outcome in the practice of monetary policy [Fischer (1994)].[11]

[11] To my knowledge, this point was first made by Nobay and Peel (1998).

Conditions for the existence and uniqueness of the Nash equilibrium are presented in the following proposition:

Proposition 1. *Provided* $1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n > 0$, *there exists a unique* $\pi_t^e = E_{t-1}\pi_t$, *such that* $g(E_{t-1}\pi_t, \pi_t^e) = 0$.
Proof. To prove existence, construct a

$$\pi_t^e = E_{t-1}\pi_t = \tilde{\pi}_t - (\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n). \tag{10}$$

Plugging (10) into (9) and using $\pi_t^e = E_{t-1}\pi_t$ delivers $g(E_{t-1}\pi_t, \pi^e) = 0$. To show uniqueness, assume there exists a second inflation forecast, say $\hat{\pi}^e = \tilde{\pi}_t - (\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n) + x$, that also lies on the 45° line on the plane $(\pi_t^e, E_{t-1}\pi_t)$ and satisfies $g(E_{t-1}\pi_t, \pi_t^e) = 0$. It will be shown that the only way this can happen is if $x = 0$. Replace $\hat{\pi}_t^e$ in (9). Then,

$$((1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n)\exp(\alpha x) - 1 - \alpha\lambda\phi(1-k)u_t^n))/\alpha = 0.$$

Simplifying,

$$(1 + \alpha\lambda\phi(1-k)E_{t-1}u^n)(\exp(\alpha x) - 1)/\alpha = 0.$$

Since $1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n > 0$ and $\alpha \neq 0$, then it must be the case that $x = 0$.¶

From this proposition follows:

Corollary 1. *In the special case where the central banker targets the expected natural rate of unemployment, there always exists a unique* $\pi_t^e = E_{t-1}\pi_t$, *such that* $g(E_{t-1}\pi_t, \pi_t^e) = 0$.
Proof. The result follows from noting that when $k = 1$, the condition $1+\alpha\lambda\phi(1-k)E_{t-1}u_t^n > 0$ is always satisfied.¶

Notice in (10) that depending on the sign of the preference parameter α, inflation is an increasing or decreasing function of its conditional variance. In order to understand this result, it is useful to recall that when the loss function is quadratic, certainty equivalence holds. Hence, the model solution is the same regardless of whether there is uncertainty or not. On the other hand, when the loss function is asymmetric on inflation, the marginal cost of departing from $\tilde{\pi}_t$ is not linear in inflation, but convex (when $\alpha > 0$) or concave (when $\alpha < 0$). Thus, when $\alpha > 0$, an increase in uncertainty increases the expected marginal cost of deviating from $\tilde{\pi}_t$. The converse result is obtained when $\alpha < 0$.

A comparable result can be found in the literature on precautionary savings. When the assumption of quadratic utility is relaxed and labor-income risk in nondiversifiable, then uncertainty increases the expected marginal utility of future consumption. To satisfy the Euler condition, prudent households decrease current consumption compared to future consumption and increase their savings. For a textbook treatment of precautionary savings, see Romer (1996, ch. 7.6).

Conditional on the inflation target, I now derive the stochastic processes of the rates of inflation and unemployment and show that, conditional on the information set, both are normally distributed. Taking conditional expectations of both sides of (4), noting that i_t forms part of the public's information set, substituting back into (4), and using (10) yields

$$\pi_t(\alpha) = \tilde{\pi}_t - (\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n) + \epsilon_t. \tag{11}$$

The notation $\pi_t(\alpha)$ makes explicit the dependence of the inflation process on the preference parameter that measures the asymmetry in the central banker's loss function. Because $\tilde{\pi}_t$ and $E_{t-1}u_t^n$ are included in the public's information set at time $t-1$:

$$E(\pi_t|I_{t-1}) = E_{t-1}\pi_t,$$

and

$$Var(\pi_t|I_{t-1}) = \sigma_\pi^2 = \sigma_\epsilon^2 = \mathbf{A}\Omega\mathbf{A}',$$

where $\mathbf{A} = (0, 0, 1)$. Since the linear combination $\mathbf{A}\boldsymbol{\xi}_t$ is normally distributed, then

$$\pi_t|I_{t-1} \sim N(E_{t-1}\pi_t, \mathbf{A}\Omega\mathbf{A}').$$

Turning now to the rate of unemployment, plug (11) into (1) and use the assumption of rational expectations to obtain:

$$u_t = u_t^n - \lambda\epsilon_t + \eta_t.$$

Notice that, as usual, the rate of unemployment realized in equilibrium does not differ systematically from the natural rate. However, since u_t^n is serially correlated, u_t will be serially correlated as well. Decompose the process of the natural rate as $u_t^n = E_{t-1}u_t^n + \zeta_t$ to write

$$u_t = E_{t-1}u_t^n + \zeta_t - \lambda\epsilon_t + \eta_t. \tag{12}$$

Since $E_{t-1}u_t^n$ forms part of I_{t-1} :

$$E(u_t|I_{t-1}) = E_{t-1}u_t^n,$$

and

$$Var(u_t|I_{t-1}) = \sigma_u^2 = \mathbf{B}\Omega\mathbf{B}',$$

where B= $(1, 1, -\lambda)$. Since the linear combination $\mathbf{B}\boldsymbol{\xi}_t$ is normally distributed, then

$$u_t|I_{t-1} \sim N(E_{t-1}u_t^n, \sigma_u^2),$$

as claimed above. Note that in contrast to models where the supply disturbance is the only shock, in this case the error term is a linear combination of the structural disturbances, rather than only a scaled version of the supply shock, η_t.

Under current inflation targeting arrangements, the targets are publicly announced and, consequently, are observable by the econometrician. Although the natural rate of unemployment is not directly observed, it is possible to write the model in terms of inflation and unemployment alone, for which data is available (this is done below in section 3.3). All this

suggests that the inflation process in (11) generates testable empirical implications regardless of whether $\tilde{\pi}_t$ is socially optimal or not. Readers interested on these predictions are invited to go directly to section 2.5. However, from the normative point of view, it is still important to derive the optimal inflation target in the more general setup where preferences are asymmetric and compare results with the ones obtained using a quadratic loss function. This task is undertaken in the following section.

2.4 Digression: The Optimal Inflation Target

In order to make results comparable with the ones is earlier literature , assume that society's or, interpreted more narrowly, the government's preferences are described by the quadratic loss function:

$$G(\pi_t, u_t) = (1/2)(\pi_t - \pi_t^*)^2 + (\phi/2)(u_t - u_t^*)^2, \tag{13}$$

where π_t^* and u_t^* denote the socially-optimal rates of inflation and unemployment, and remaining notation is as previously defined. The possibly nonzero rate of inflation π_t^* could be interpreted as the one associated with the optimal inflation tax, that might reduce the need for distorting income taxes or other type of non-lump-sum taxes.

In order to derive the optimal inflation target, consider the problem of a government that chooses the sequence $\{\tilde{\pi}_t\}_{t=0}^{\infty}$ that maximizes the present discounted value of the social welfare function. It is important to note that in actual practice, targets are sometimes jointly agreed by the government and the monetary authority. For example, section 9 of the Reserve Bank of New Zealand Act requires the Finance Minister *and* the Governor of the Reserve Bank to negotiate and make public the Policy Target Agreements (PTA's) that state the inflation targets [Bernanke *et al.* (1999, p. 88)]. Similarly, in Canada the targets are jointly determined and announced by the government and central bank. However, since in this model there is no private information and preferences are observable prior to delegation, the simplifying assumption that the choice of targets rests with the government is innocuous.[12]

The government's problem is

$$\underset{\{\tilde{\pi}_{t+s}\}_{s=0}^{\infty}}{Min} \quad E_{t-1} \sum_{s=0}^{\infty} \beta^s G\left(\pi_{t+s}(\tilde{\pi}_{t+s}), u_{t+s}\right),$$

where the minimization is made subject to (11). Without loss of generality, it is assumed that the government discounts the future at the same rate, β, as the central banker. Because the choice of the inflation target for time t does not affect the inflation process beyond

[12]Beestma and Jensen (1998) and Muscatelli (1999) develop static models where the central banker's preferences are unobserved by the government at the moment of delegation. The uncertainty pertains only to the inflation/unemployment weights that are stochastic, in an otherwise standard quadratic loss function. Note however that the rigorous modeling of preference uncertainty in a dynamic setup would also require the specification of a learning mechanism on the part of the public/government. Thus, in order to keep the scope of this project manageable, I have adopted the simplifying assumption that preferences are public information. To motivate this assumption, it could be argued that the track record and the previously-expressed opinions of individuals might provide information about their preferences when in office.

time t, the government's objective function can also be decomposed into a sequence of one-period problems. Noting that the rate of unemployment does not depend on $\tilde{\pi}_t$ and that $\partial\pi_t/\partial\tilde{\pi}_t = 1$, the first-order condition is

$$E_{t-1}((\partial G/\partial\pi_t)(\partial\pi_t/\partial\tilde{\pi}_t)) = E_{t-1}\pi_t - \pi_t^* = 0.$$

Then, using (10) and solving for the optimal inflation target yields

$$\tilde{\pi}_t(\alpha) = \pi_t^* + (\alpha/2)\sigma_\pi^2 - (1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n). \tag{14}$$

The notation $\tilde{\pi}_t(\alpha)$ underscores the dependence of the optimal target on the asymmetry of central banker's loss function. As in Svensson (1997), the optimal target is time-varying, rather than the constant value typically observed in actual practice. However, in this case the relation between the target and unemployment is nonlinear, depends on the conditional variance of inflation, and can be smaller or larger than the optimal inflation rate.

The optimal inflation target under quadratic preferences can be obtained by taking the limit of (14) as $\alpha \to 0$. This manipulation yields

$$\tilde{\pi}_t(0) = \pi_t^* - \lambda\phi(1-k)E_{t-1}u_t^n < \pi_t^*.$$

Notice that when $\lambda\phi(1-k)E_{t-1}u_t^n > 0$, the optimal target under quadratic preferences is strictly lower than the optimal inflation rate. The result that $\tilde{\pi}_t(0) < \pi_t^*$ means that the welfare-maximizing government entrusts the "quadratic" central banker with a target that is low enough *vis a vis* π_t^* that, even if the latter is subject to an inflation bias, the marginal costs and benefits of inflation to the central banker are equalized at the socially-optimal inflation rate. This result was first reported by Svensson (1997). In the special case where the central banker targets the natural rate of unemployment ($k = 1$), the optimal inflation target corresponds exactly to the socially-optimal rate of inflation.

That the optimal inflation target can deliver the optimal monetary policy can be seen by replacing (14) into (11) to obtain

$$\pi_t = \pi_t^* + \epsilon_t. \tag{15}$$

Hence, the unconditional mean of inflation is the socially-optimal rate, and inflation realizations differ from π_t^* only by a mean-zero and serially uncorrelated random term. Since the disturbance ϵ_t is outside the control of the monetary authority, the model highlights the fact that inflation is not perfectly controllable by the central banker and predicts that, in practice, announced inflation targets are likely to be specified in terms of a tolerance range rather than as a sole numerical value. This implication of the model appears consistent with empirical evidence: of the eight inflation targeting countries surveyed by Bernanke and Mishkin (1997), only Finland defines its targets as a number alone. Israel has defined its inflation target as range in 1993 and from 1996 onwards, and a as number in 1994 and 1995. Since June 1995, the United Kingdom defines its target as a point rather than as a range. However, since June 1997 there is a 1% threshold around the target rate that, if breached by the inflation rate, prompts and explanatory letter from the Monetary Policy Committee to the Chancellor of the Exchequer.

2.5 Implications

As noted above, the fact that inflation targets are observable means that (11) generates testable predictions regardless of whether the targets are optimal or not. These predictions can be formulated in terms of either the rate of inflation or its deviation from the announced target. For the general case where the loss function is asymmetric, it is trivial to rewrite (11) as:

$$\pi_t(\alpha) - \tilde{\pi}_t = -(\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n) + \epsilon_t. \tag{16}$$

Notice that the term $(1/\alpha)\ln(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n)$ is nonnegative but $-(\alpha/2)\sigma_\pi^2$ can be positive or negative depending on the sign of the preference parameter α.

Consider first the case $\alpha < 0$, where the central banker weights positive deviations from the target less severely than negative ones. Since $-(\alpha/2)\sigma_\pi^2$ is larger than zero, the mean inflation deviation from the target is unambiguously positive. Consider now the case $\alpha > 0$, meaning that the central banker weights positive deviations from the target more severely than negative ones. Then, $(\alpha/2)\sigma_\pi^2 > 0$ and for certain parameter values, it is possible that its magnitude be large enough to yield $\pi_t(\alpha) - \tilde{\pi}_t < 0$. Hence, inflation could be on average below the announced target. In summary, under a asymmetric preferences, inflation could be systematically above or below its target depending on the central banker's preference parameter α. In contrast, it is shown below that the quadratic model predicts that average inflation is no smaller than the announced target. In the special case examined by Green (1996) and Svensson (1997) where $0 < k < 1$, average inflation is always larger than the announced target.

The asymmetric model predicts that the conditional variance of inflation is helpful in predicting its mean. In a cross section, this means, for example, that if $\alpha > 0$, countries with more volatile inflation should have smaller average inflation deviations from their target. Some limited evidence to this effect is presented in section 3.4. In a time-series, note that the preference parameter is statistically identified as the coefficient of the conditional variance of inflation and from the nonlinear term on unemployment.

As a comparison, it is useful to derive the implications of the model with quadratic preferences. To that effect, take the limit of (16) as $\alpha \to 0$ to obtain

$$\pi_t(0) - \tilde{\pi}_t = \lambda\phi(1-k)E_{t-1}u_t^n + \epsilon_t.$$

Since the term $\lambda\phi(1-k)E_{t-1}u_t^n$ is nonnegative, the average realization of $\pi_t(0) - \tilde{\pi}_t$ is also nonnegative. Thus, this model predicts that when $0 < k < 1$, inflation is systematically above the announced target. Because this overshooting is rationally anticipated by the public, it follows that the targeting policy can only be imperfectly credible [see Svensson (1997)]. In addition, quadratic preferences imply that the relation between inflation deviations and unemployment is linear and the conditional variance has no explanatory power on the current deviation from the target. If one generalizes Svensson's model to allow the case where the central banker targets the natural rate of unemployment ($k = 1$), the model predicts that the average inflation deviation from the target is zero.

A more direct comparison of both models can be obtained by examining the variable

$$\begin{aligned}(\pi_t(\alpha) - \tilde{\pi}_t) - (\pi_t(0) - \tilde{\pi}_t) \quad &= -(\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln\left(1 + \alpha\lambda\phi(1-k)E_{t-1}u_t^n\right) \\ &\quad -\lambda\phi(1-k)E_{t-1}u_t^n,\end{aligned}$$

that explicitly compares the deviations predicted by the quadratic and asymmetric models. This relation is plotted in figure 3 as a function of the preference parameter α and for different values of $E_{t-1}u_t^n$, namely $2\%, 4\%$, and 6%. In constructing this graph, I have set $\lambda = 2$, $\sigma_\pi^2 = 1$, $\phi = 0.5$, and $k = 0.8$. Notice that when $\alpha \to 0$ both models predict exactly the same deviation from the target and, consequently, their difference is zero. For negative values of α, that is when the central banker weights more severely negative than positive deviations from the target, the model with asymmetric preferences predicts positive deviations that are larger than the (also positive deviations) predicted by the quadratic model. Notice that the nonlinear relation between unemployment and inflation can be quite relevant for certain large negative values of α. On the other hand, for positive values of α, the model with asymmetric preferences predicts negative deviations or positive deviations that are smaller than the one predicted by the specification with quadratic preferences, and, consequently, their difference is negative. Also note that for $\alpha > 0$, the relationship between the variables is approximately linear.

As previous game-theoretical models of monetary policy [see, among others, Barro and Gordon (1983), Rogoff (1985), and Svensson (1997)], the rates of inflation and unemployment are predicted to be positively related when $0 < k < 1$. However, in this case, the relationship is nonlinear. The preliminary analysis of the data in section 3.2 suggests that the nonlinear model indeed yields a more accurate fit of the inflation/unemployment observations, than the linear specification.

In the special case where the central banker targets the natural rate of unemployment ($k = 1$), the model predicts no systematic relationship between inflation and unemployment. However, an inflation bias can still arise in the case where $\alpha < 0$. This suggests that the result that the inflation bias is zero when $k = 1$ [see McCallum (1995, 1997) and Blinder (1998)] is not robust to the generalization of the policy maker's loss function.[13] In the more plausible case where $\alpha > 0$, a pure deflation bias arises and inflation is (on average) below its target for any positive α.

Finally, consider the effect on the inflation deviation from target of an innovation that increases the public's forecast of the natural rate. In the quadratic case:

$$\partial(\pi_t(0) - \tilde{\pi}_t)/\partial E_{t-1}u_t^n = \lambda\phi(1-k),$$

that is positive, constant and independent of the expected rate of unemployment when the innovation takes place. On the other hand, in the asymmetric case:

$$\partial(\pi_t(\alpha) - \tilde{\pi}_t)/\partial E_{t-1}u_t^n = \lambda\phi(1-k)/(1+\alpha\lambda\phi(1-k)E_{t-1}u_t^n),$$

that is also positive but depends on the expected unemployment rate. Note that, for example, for $\alpha > 0$, the innovation always has a smaller effect than in the quadratic model and decreases with the rate of unemployment. The empirical predictions of the models with quadratic and asymmetric preferences are summarized in table 2.

[13] On this see also Cukierman (2000) and Ruge-Murcia (2000).

3 Econometric Analysis

3.1 The Data

The empirical predictions of the model are examined using monthly observations of the rates of inflation and unemployment for Canada, Sweden, and the United Kingdom. These inflation-targeting countries feature time series that are sufficiently long to allow the meaningful estimation and testing of the model. In addition, their monetary policy objectives appear not to involve (at least explicitly) exchange rate considerations. Thus, despite very occasional interventions in foreign exchange markets, their exchange rate regime could be characterized as floating [see Bernanke *et al.* (1999)]. This is important because modeling the objectives of a central banker explicitly concerned about the nominal exchange rate, would require a more general model with the loss function appropriately augmented to include this institutional feature.

For example, Finland had an inflation target of 2% per year between February 1993 and December 1998, and was simultaneously a member of the Exchange Rate Mechanism (ERM).[14] Under the ERM, the nominal exchange rate was to be kept in a $\pm 15\%$ range around a target value and later fixed at an "irrevocable" rate prior to monetary union. Because of this, some authors argue that Finland's monetary policy might not correspond to one of pure inflation targeting [see Gerlach (1999, p. 1263) for a discussion].

Australia, New Zealand, and Israel also employ inflation targets in the conduct of their monetary policy. Unfortunately, their statistical offices publish data on inflation and/or unemployment only on a quarterly basis. The Australian Bureau of Statistics collects price data every quarter, pricing most articles in the first two months of each quarter. Only certain items (for example, milk, bread, fresh meat, fresh fruit and vegetables, holiday travel and accommodation) are surveyed each month. Similarly, Statistics New Zealand prices most items at the mid-point of each quarter. Food is the only commodity group of the Consumers Price Index (CPI) for which an index is prepared each month. Although monthly CPI data is available for Israel, the labor force survey of the Central Bureau of Statistics takes place only once per quarter. All this means that, since inflation targeting is a relatively recent policy, the number of observations available to estimate and test the model for these countries is too small to yield reliable results.[15]

For Canada, Sweden, and the United Kingdom, estimation is carried using the "targeted" measure of inflation. That is, the 12-month percentage change in a price index that, in the case of Canada and the United Kingdom, excludes certain components usually included in the CPI. However, in order to assess the robustness of the results, estimates using the broader CPI (RPI in the United Kingdom) are also presented in section 3.5.

The Canadian inflation targets were announced in February 1991 and originally envisaged a reduction in the annual rate of inflation to 3% by the end of 1992, 2.5% by mid-1994, and 2% by the end of 1995. In December 1993, the inflation target of 2% per year was extended

[14] In January 1999, Finland became part of the European Monetary Union and monetary policy was in effect transferred to the European Central Bank.

[15] At the time of collecting the sample (fall of 2000), there are only 23, 39, and 34 observations of both inflation and unemployment for Australia, New Zealand, and Israel, respectively.

to the end of 1998 and, in February 1998, it was further renewed to December 2001. I follow the literature [for example, Lafrance (1997)], in interpreting the Canadian targets during the period December 1992 to December 1995 as the linear interpolation of the announced values. The targets are regarded as mid-points in a band of plus or minus one percentage point and apply to a measure of "core" inflation that excludes the volatile food and energy components and the first-round effect of changes in indirect taxes. The annual percentage change of the CPI and Core CPI were obtained from the *Weekly Financial Statistics* published by the Bank of Canada (http://www.bankofcanada.ca).

In Sweden, an inflation target of 2% per year was announced by the Governing Board of the Riskbank in January 1993. It refers to the annual change in the CPI with no exclusions and, as in Canada, involves a tolerance range of $\pm 1\%$ around the target. The raw CPI series was taken from the Web Site of the Riksbank (http://www.riksbank.se).

In the United Kingdom, the Chancellor of the Exchequer announced an inflation target of between 1% and 4% per year in October 1992. Initially, targets were meant to apply "through the end of the present parliament" alone. In June 1995, the government reinterpreted the inflation target as a numerical value of 2.5% per year. However, since June 1997 there is a 1% threshold around the target rate that, if breached by the inflation rate, prompts and explanatory letter from the Monetary Policy Committee to the Chancellor of the Exchequer. Throughout, the targeted measure of inflation is the annual change in the Retail Price Index excluding mortgage interest payments (or RPIX). Both the RPI and the RPIX were taken from the *OECD Main Economic Indicators*. Plots of the rates of inflation and announced targets for Canada, Sweden, and the United Kingdom are presented in figure 4.

Note that by design, inflation targets are intended to apply to annual, rather than to monthly, inflation. However, because (*i*) annual inflation is the sum of the twelve most recent observations of monthly inflation[16], (*ii*) the policy holds continuously, and (*iii*) past rates of monthly inflation are predetermined at time t, it is possible estimate the model using data observed at a higher frequency than the time horizon for which the policy is defined. See Ruge-Murcia (forthcoming) for a discussion.

Unemployment is measured by the survey-based, seasonally-adjusted rate of unemployment published by the *OECD Main Economic Indicators*.[17] The sample periods are as follows: Canada from 1992:12 to 2000:6, Sweden from 1995:1 to 2000:6, and the United Kingdom from 1992:10 to 2000:6. The different sample periods start at the time when the inflation targeting policy was first launched in each of country and end with the latest available observation of the variables at the time the data was collected.

[16]This result relies on the definition of the rate of inflation as the change in the log of the price level. While, strictly speaking, this is an approximation to its percentage change, the numerical difference between the two measures is negligible for small rates of inflation.

[17]In preliminary work, I also used the seasonally unadjusted unemployment series. Results are similar to the ones reported and are available from the author upon request.

3.2 A Rough Look at the Data

Prior to estimation, this section uses simple summary statistics and plots to examine to what extend the model predictions are (or are not) borne by the data. Consider first the predictions of the quadratic and asymmetric models regarding the average inflation deviation from the target. Under quadratic preferences, inflation should be on average above or at the target. Under asymmetric preferences, inflation can be on average above or below the target depending on the sign and magnitude of the preference parameter α. A negative average inflation deviation from the announced target is suggestive of a positive value of α.

Table 3 contains the sample mean of the variable $\pi_t - \tilde{\pi}_t$ and the percentage of inflation observations above and below the target. For Canada and Sweden, the average deviation from the target is negative and quantitatively important (-0.61 and -1.23, respectively), and the proportion of observations below target is high (91.2% and 81.2%, respectively). In the case of Sweden, the average deviation is actually below the lower limit of its inflation target zone. On the other hand, for the United Kingdom, the average inflation deviation from the target is small and positive (0.15) and 69.9% of the observations are above $\tilde{\pi}_t$.[18]

Although these statistics are provocative and appear to challenge the usual assumption of quadratic preferences, they must be interpreted with caution for two reasons. *First*, since the series correspond to monthly observations of annual inflation, the data points are serially correlated and do not constitute independent evidence against quadratic preferences. *Second*, these statistics make no use of the restrictions embodied in the economic model constructed in section 2. A more formal analysis of the data is performed below in sections 3.4 and 3.5.

A second prediction common to game-theoretical models of monetary policy is that the rates of inflation and unemployment are positively related.[19] Analytical results in section 2 show that this prediction is robust to generalizing the functional form of the central banker's loss function. However, under asymmetric preferences the relation between inflation and unemployment is nonlinear or, more precisely concave. In the special case where $k = 1$, both models predict no relationship between inflation and unemployment.

Because the inflation target is predetermined and publicly known, the prediction that π_t and u_t are positively related means that the inflation deviation from the target and unemployment should be positively related as well. Figure 5 presents scatter plots of $\pi_t - \tilde{\pi}_t$ with the rate of unemployment in the horizontal axis. As a very rough evaluation of the model prediction, the figures include the fitted line of a Least Squares projection of $\pi_t - \tilde{\pi}_t$ on a constant and the unemployment rate. Because u_t is endogenous, this projection was carried out by Two-Stage Least Squares (TSLS) using lagged unemployment as an instrument for current unemployment. The estimates of this regression are reported in table 4. Notice that for Canada, the coefficient on unemployment is negative and significantly different from

[18] These statistics are robust to using broader aggregate price measures. For Canada, the average CPI inflation deviation from the target is -0.69, and 76.9% of its observations are below target. For the United Kingdom the average RPI inflation deviation from target is 0.06, and 55.9% of its observations are above target.

[19] When the models are written in terms of output, rather than unemployment, the prediction is that output and inflation are negatively related. This result is independent of assuming a neoclassical expectations-augmented Phillips curve (as done here) or a New Keynesian Phillips curve. For example, the first-order condition in Clarida *et al.* (1999, p. 1672) predicts a negative relation between output gap and inflation.

zero. On other hand, for Sweden and the United Kingdom the coefficients are positive, but significantly different from zero only in the latter case. Hence, the prediction that the inflation deviation from target and the unemployment rate are positively related is supported by UK and (to some extent by) Swedish data, but not by Canadian data.

Since the prediction of the model with asymmetric preferences is that the relation is positive but nonlinear, table 4 also reports the results of a TSLS regression of $\pi_t - \tilde{\pi}_t$ on a constant, u_t, and u_t^2. Notice that the coefficients on u_t are positive and statistically different from zero at the 5% level in all cases. The coefficients on u_t^2 are negative and statistically different from zero at the 5% level in all cases. This result supports the notion of a nonlinear, concave relation between the inflation deviation from target and unemployment. Finally, note that the R^2's of the nonlinear model are considerably larger than the ones of the linear model. (This result is independent of whether one compares adjusted or unadjusted R^2's.)

3.3 Reduced-Form of the Model

The simple game-theoretical model developed in section 2 consists of three variables, namely the natural rate of unemployment and the rates of inflation and unemployment. However, although data on inflation and unemployment is readily available, the natural rate is not directly observable.[20] In order to allow the estimation of the model using observations on inflation and unemployment alone, a reduced-form version is now constructed. Taking conditional expectations in both sides of (3) and substituting into (12) yields

$$u_t = \psi + \delta u_{t-1}^n + \sum_{i=1}^{q-1} \theta_i \Delta u_{t-i}^n + \zeta_t - \lambda \epsilon_t + \eta_t.$$

Adding and substracting $\psi + \delta u_{t-1} + \sum_{i=1}^{q-1} \theta_i \Delta u_{t-i}$, noting that $u_t^n - u_t = \lambda \epsilon_t - \eta_t$, and substracting u_{t-1} in both sides:

$$\begin{aligned} \Delta u_t = \; & \psi - (1-\delta) u_{t-1} + \sum_{i=1}^{q-1} \theta_i \Delta u_{t-i} \\ & + \zeta_t - \lambda \epsilon_t + \eta_t + \delta(\lambda \epsilon_{t-1} - \eta_{t-1}) + \sum_{i=1}^{q-1} \theta_i (\lambda \Delta \epsilon_{t-i} - \Delta \eta_{t-i}). \end{aligned} \tag{17}$$

Equation (17) describes the process of the rate of unemployment as a function of its lagged values and a linear combination of current and past structural disturbances. An advantage of this representation is that it does not include the unobserved natural rate as one of the regressors. However, with only data on π_t and u_t, one cannot disentangle the residuals of each structural disturbance to construct empirical counterparts for the lagged ϵ_t and η_t that enter (17). Hence, the unemployment process cannot be estimated without additional statistical restrictions.

[20] A number of authors [for example, Staiger, Stock, and Watson (1997)] construct estimates of the natural rate using data on inflation and unemployment. However, in the context of this model, it is clear that such estimates provide no additional information beyond that already contained in π_t and u_t.

Consider a strategy that involves (*i*) assuming that ζ_t, ϵ_t and η_t are mutually uncorrelated with each other at all leads and lags,[21] and (*ii*) using time series results on the aggregation of $ARMA$ processes [see Harvey (1981, p. 43) and Hamilton (1994, ch. 4.7)]. Two of these results are relevant here. *First*, adding two moving average (MA) processes that are uncorrelated with each other at all leads and lags yields another MA process with order equal to the maximum order of its two components. *Second*, adding a white noise series to a moving average with which it is uncorrelated at all leads and lags produces a new MA process of the same order. Since the sequence $\zeta_t - \lambda\epsilon_t + \eta_t + \delta(\lambda\epsilon_{t-1} - \eta_{t-1}) + \sum_{i=1}^{q-1} \theta_i(\lambda\Delta\epsilon_{t-i} - \Delta\eta_{t-i})$ aggregates a white noise and two moving averages of order q, the two results above imply that there exists a $MA(q)$ process, say $w_t + \sum_{i=1}^{q} \tau_i w_{t-i}$, with *exactly* the same statistical properties as the original series. Then, the process for Δu_t can be written in reduced-form as an unrestricted $ARMA(q-1, q)$:

$$\Delta u_t = \psi - (1-\delta)u_{t-1} + \sum_{i=1}^{q-1} \theta_i \Delta u_{t-i} + w_t + \sum_{i=1}^{q} \tau_i w_{t-i}. \tag{18}$$

Equation (18) makes apparent a number of advantages of the strategy outlined above. *First*, it involves weaker assumptions than alternative identification schemes [for example, Ireland (1999)]. *Second*, although restrictions are imposed on the variance-covariance matrix of the structural disturbances, no restrictions arise on the variance-covariance matrix of the reduced-form disturbances (see below). *Third*, estimation is straightforward and can be carried out using standard procedures. For example, lagged values of w_t could be proxied empirically by lagged residuals of (18). *Finally*, since the $ARMA$ process is unrestricted, it could be well approximated by a finite-order autoregressive or moving average processes, if necessary.

Regarding the rate of inflation, note that taking E_{t-1} in both sides of (12), implies

$$E_{t-1}u_t^n = E_{t-1}u_t.$$

This result follows from the fact that unemployment differs from the natural rate only by a mean-zero and serially-uncorrelated random term. Hence, the forecast of u_t^n is numerically equivalent to the forecast of u_t, when both are based on the same information set, I_{t-1}. Furthermore, by (18), the latter can be constructed on the basis of past observations of unemployment alone. With this result, equation (11) can be written as

$$\pi_t(\alpha) = \tilde{\pi}_t - (\alpha/2)\sigma_\pi^2 + (1/\alpha)\ln(1 + \alpha\gamma E_{t-1}u_t) + \epsilon_t, \tag{19}$$

where $\gamma = \lambda\phi(1-k) \geqslant 0$ is a constant coefficient.

Finally, from the assumptions about the structural shocks, it follows that the reduced-form disturbance, w_t, and ϵ_t are serially uncorrelated and jointly normally distributed with zero mean:

$$\left[\begin{array}{c} w_t \\ \epsilon_t \end{array} \right] \Bigg| \, I_{t-1} \quad \sim N(\mathbf{0}, \mathbf{\Psi}),$$

[21] Note that this entails the restriction that the off-diagonal elements of Ω are zero.

where Ψ represents a 2×2 positive-definite variance-covariance matrix. Since w_t is an aggregation of structural shocks, w_t and ϵ_t are contemporaneously correlated and the off-diagonal element in Ψ is nonzero.

3.4 Estimation

Recall that unemployment can be written in reduced-form as an unrestricted $ARMA(q-1, q)$ process. However, since (*i*) the estimation of $ARMA$ processes is frequently complicated by common factors and (*ii*) any stationary $ARMA$ process has an autoregressive representation, the unemployment process is estimated here in autoregressive form. (Results using a low-order $ARMA$ process yielded virtually the same results as reported below and are available from the author upon request.)

Under asymmetric preferences, inflation depends nonlinearly on the public's unemployment forecast. In order to examine the robustness of the results to the use of different forecasting models of unemployment, this paper considers two processes for u_t. The first one is a stationary specification that corresponds to the case where $0 \leq \delta < 1$ in (3). The second one is a unit root model that is obtained when $\delta = 1$ in (3). For both models, the lag length of the AR representation was determined using Akaike's Information Criterion (AIC). After estimating processes with lag length 1 to 9, results indicated that the most appropriate stationary specifications for the rates of unemployment in Canada, Sweden, and the United Kingdom are $AR(3)$, $AR(4)$, and $AR(4)$, respectively. When modeling unemployment as nonstationary, the most appropriate specifications are $AR(2)$, $AR(3)$, and $AR(3)$, respectively.

Specification tests for both forecasting models of unemployment are reported in table 5. Durbin-Watson and Breusch-Godfrey tests for serial correlation of the residuals are reported in rows 1 and 2. The Breusch-Godfrey test statistics were calculated as the product of the number of observations and the uncentered R^2 of the Ordinary Least Square (OLS) regression of the unemployment residuals on a constant, lagged unemployment rates, and two lagged residuals. Under the null hypothesis of no serial correlation, the statistic is distributed chi-square with as many degrees of freedom as lagged residuals are included in the regression. For all countries and models the null of hypothesis of no serial correlation cannot be rejected at the 5% level.

Lagrange Multiplier (LM) tests for neglected Autoregressive Conditional Heteroskedasticity (ARCH) are reported in row 3. The statistics were calculated as the product of the number of observations and the uncentered R^2 of the OLS regression of the squared unemployment residual on a constant and two of its lags. Under the null hypothesis of no conditional heteroskedasticity, the statistic is distributed chi-square with as many degrees of freedom as squared residuals are included in the regression. The null of hypothesis of no conditional heteroskedasticity cannot be rejected by the data at the 5% level in any country.[22]

The bivariate process of inflation and unemployment was estimated by the numerical maximization of their joint log likelihood function. Since this function imposes the cross-equation restrictions that arise from the dependence of inflation on $E_{t-1}u_t$, its maximization

[22]Note, however, that for Sweden the null hypothesis would be rejected at the 10% significance level.

yields consistent and efficient Full Information Maximum Likelihood (FIML) estimates of the parameters. Asymptotic standard errors were computed using as estimate of the variance-covariance matrix the inverse of the Hessian of the log likelihood function at the maximum. In order to assess the robustness of the results to deviations from the assumption of normality, Quasi-Maximum Likelihood standard errors [White (1982)] were also computed and used in statistical inference. A well-known feature of nonlinear models is that their log likelihood function might present numerous local maxima. In order to address this issue, the robustness of the global maximum was examined using the method of simulated annealing. This approach is very efficient computationally and by linking different points in the domain through a Markov chain insures that ultimately any point on the surface will be visited. For more details the reader is referred to Judd (1998, ch. 8.3).

Parameter estimates for Canada, Sweden and the United Kingdom are reported in tables 6, 7 and 8, respectively. Several results are apparent from these tables. *First*, estimates appear robust to assuming that the public forecasts unemployment using a stationary or a nonstationary specification. This result is not surprising because the one-step-ahead predictions of a persistent variable computed using a nonstationary or persistent stationary process are likely to be quite similar. The fact that estimates of the θ-coefficients are different simply reflects the fact that their interpretation is different in each model. For example, in the stationary specification, θ_1 is the coefficient of lagged unemployment, while in the nonstationary specification, it is the coefficient of the lagged first-difference of unemployment.

Second, estimates of γ vary substantially across countries. As expected from the empirical analysis in section 3.2 and figure 5, the estimate of γ is negative and statistically different from zero for Canada but positive for Sweden and the United Kingdom. Consider first the results for Canada. Note that the conclusion that γ is negative and statistically different from zero does not depend on the estimate of the standard error employed to construct the t-statistic. This result indicates that unemployment is helpful in forecasting the rate of inflation (as expected) but in an opposite direction as predicted by the model. Thus, it would appear that this simple game-theoretical model of monetary policy might not completely capture the statistical relation between inflation and unemployment in the case of Canada.

Consider now the results for Sweden. In this case, $\hat{\gamma}$ is positive but whether it is statistically different from zero or not depends on the estimate of the standard error employed to construct the t-statistic. A Likelihood Ratio test of the restriction $\gamma = 0$ yields a statistic of 0.62 that is smaller than the 5% critical value of chi-square variable with one degree of freedom.[23] Hence, the restriction cannot be rejected at the 5% significance level. The finding that unemployment is not helpful in forecasting inflation is still consistent with a version of the model where the central banker targets the natural rate of unemployment. However, since the sample size for Sweden is the smallest among the countries considered, it is also possible that this result simply reflects the larger uncertainty regarding the model parameters that is associated with finite samples.

Finally, consider the results for the United Kingdom. As in the case of Sweden, the statistical significance of the estimate depends on the estimate of the standard error employed to construct the t-statistic. However, a Likelihood Ratio test of the restriction $\gamma = 0$ yields

[23]This statistic (and the one reported for the United Kingdom below) were obtained using the stationary model for the rate of unemployment. Conclusions were the same when using the unit-root model.

a statistic of 15.26 that is larger than the 5% critical value of chi-square variable with one degree of freedom. Hence, the restriction is rejected at the 5% significance level. The result that $\gamma\ (=\lambda\phi(1-k)) > 0$ is not meant to suggest that the Bank of England targets a rate of unemployment below the natural rate, but simply that unemployment is helpful in forecasting the inflation rate in a manner consistent with the simple game-theoretical model. LM tests of the overidentifying restrictions are reported in section 3.5.

Third, estimates of the preference parameter α are positive and quantitatively large: 4.54 for Canada, 3.42 for Sweden, and 2.64 for the United Kingdom. This means that for the countries under consideration, positive inflation deviations from the target appear to be weighted more severely than negative ones in the central banker's loss function, even if they are of the same magnitude. In the case of Canada, the null hypothesis that the true preference parameter is zero is rejected at the 1% significance level against the two-sided alternative that it is different from zero.[24] In the case of the United Kingdom, the hypothesis is rejected at the 1% level when one computes the t-statistic using the robust standard error, and at the 10% when one employs the asymptotic standard error. In the case of Sweden, the hypothesis is rejected at the 1% level when the t-statistic is computed using the robust standard error but cannot be rejected when one uses the asymptotic standard error. When one considers the test of the same hypothesis ($\alpha = 0$) against the one-sided alternative that $\alpha > 0$, the null hypothesis is rejected in all cases at the 10% level (1% level in most cases). However, given the limited sample sizes currently available to estimate inflation target models, these results should be interpreted with caution.

The result that $\alpha > 0$ for Canada and Sweden was to some extent anticipated by the sample statistics reported in table 3. However, recall that for the United Kingdom, these statistics revealed a small, but positive, average inflation deviation from target (0.15) and a substantial proportion of observations above target (69.9%). Then, the result that α is positive and significantly different from zero for the UK, means that the average deviation from target and the percentage of observations above target are simply too low to be consistent with a model with quadratic preferences.

Figure 6 plots the central banker's loss functions implied by the estimates of the preference parameter α and compares them with the usual quadratic loss function (dotted line). Note that although there are some differences in the numerical estimates of α for the three countries, their loss function in the range of interest are similar. In all cases, negative deviations from the target are penalized much less than under a quadratic loss function. Small positive deviations from the target (say between 0 and 0.5) are penalized roughly in the same manner in all loss functions, including the quadratic. Finally, positive deviations from the target, are penalized much more severely in the estimated asymmetric loss functions than in the quadratic.

Under asymmetric preferences the marginal cost of departing from the target is convex when $\alpha > 0$. Hence, uncertainty raises the expected marginal cost of deviating from $\tilde{\pi}_t$ and induces a prudent behavior on the part of the central banker. The analytical counterpart of this prudence motive is given by the term $(\alpha/2)\sigma_\pi^2$ in (11). Note that it increases with the

[24]Note that even though $\alpha = 0$ forms part of the domain and the limit of the log-likelihood function as $\alpha \to 0$ exists, strictly speaking the function is not continuous at the point $\alpha = 0$. To circumvent this problem, I have used Wald-type t-tests to assess the significance of $\hat{\alpha}$.

conditional variance of inflation and the asymmetry in preferences. Prudence moderates the central banker's inflation bias (if there is any) and for large enough values of either α or σ_π^2, can override it. Then, a prudent central banker would pursue a more conservative monetary policy than her quadratic counterpart, in the sense that expected inflation is lower. Depending on the preference parameters, it might be possible to observe realizations of inflation that are below target for seemingly long periods of time.

¿From the above discussion it follows that, for a given preference parameter $\alpha > 0$, countries with a larger conditional variance of inflation, should have a lower (that is, more negative) average inflation deviation from the target. Although the number of countries in the sample is too small to allow a fully-fledged cross-section analysis, one can still verify graphically to what extent this prediction is supported by the data. Figure 7 plots the relation between the conditional variance of inflation and the average deviation from target. Notice that Sweden has the largest conditional variance of inflation (1.14) and the smallest average deviation from the target (-1.23). On the other hand, Canada and the United Kingdom have much smaller conditional variances of inflation (0.129 and 0.130, respectively) and larger (that is, less negative) average inflation deviations from the target (-0.63 and 0.15, respectively). The larger conditional variance of inflation in Sweden and the prudence motive that arises under asymmetric preferences, could explain why average $\pi_t - \tilde{\pi}_t$ is much lower in Sweden than in Canada and the United Kingdom, even though the estimate of the preference parameter is roughly the same.

3.5 Specification Tests and Robustness Analysis

This section reports specification tests of the model, examines the robustness of the results to using broader measures of inflation, and discusses other theories that could also account for the empirical results.

Because inflation depends on the public's unemployment forecast, the model imposes overidentifying restrictions on the joint process of inflation and unemployment. These restrictions were tested by means of a Lagrange Multiplier test. The alternative specification was constructed to nest the model as a special case and differed from the null in that the unemployment coefficients in (19) were unrestricted. Results for both unemployment models and all countries are reported in the first row of table 9. Under the null hypothesis, the LM test statistics is asymptotically distributed chi-square with as many degrees of freedom as restrictions are tested. The number of restrictions tested for Canada, Sweden, and the United Kingdom are 3, 4, and 4, respectively, for the stationary unemployment model, and 2, 3, and 3, respectively, for the unit-root model. Comparing the statistics with the 5% critical value of the appropriate distribution, indicates that the restrictions cannot be rejected at the 5% significance level. Although the p-values for the United Kingdom are 0.20 and 0.12 for stationary and unit-root models, respectively, they are still above the standard 10% critical value. Note, however, that since the small-sample distribution of the statistic might differ from the asymptotic one, the above results should be best regarded as indicative only.[25]

[25] For the case of Sweden, results might also reflect low power because the coefficient of unemployment is very imprecisely estimated.

The inflation equation (19) allows for the intercept term $-\alpha\sigma_\pi^2/2$. However, since α is also a coefficient of the unemployment forecast, and σ_π^2 appears in the variance-covariance matrix of the residuals, the intercept is restricted. The restriction is tested by means of a LM test where the alternative allows for a free intercept. Results are reported in the second row of table 9. Under the null hypothesis, the LM statistic is asymptotically distributed chi-square with one degree of freedom. Comparing the statistics with the 5% critical value leads to the conclusion that the restrictions imposed by the model on the intercept cannot be rejected at the 5% significance level.

Less favorable to the model are the tests for serial correlation of the inflation residuals reported in rows 3 and 4 of table 9. Comparing the Durbin-Watson test statistics with the upper and lower bounds of 5% critical value of its tabulated distribution leads to the rejection of the null hypothesis of no serial correlation in favor of the alternative of positive serial correlation. Similarly, Pormanteau tests for the first order autocorrelation of the residuals yield statistics that are well above their 5% critical value. (Under the null hypothesis of no autocorrelation, the test statistic is distributed chi-square with as many degrees of freedom as autocorrelations are tested for.) However, note that since (19) does not include lagged inflation rates in the right-hand side, serial correlation reduces the efficiency, but does not affect the consistence of the FIML estimates reported above.

Consider the situation where the conditional variance of inflation is assumed to be a function of lagged squared residuals, as in the ARCH model proposed by Engle (1982). In this case, lagged squared rates of inflation are implicitly included among the explanatory variables. Since residuals are serially correlated, estimates are likely to be biased and inconsistent.[26] In preliminary work, I considered parameterizing the conditional variance of inflation as function of lagged squared changes in the oil price, that could be plausibly assumed to be exogenous [see Hamilton (1983)]. However, its coefficient is not statistically different from zero and results are similar to the ones obtained under the assumption of conditional homoskedasticity.[27]

One could also consider a more general model with asymmetric preferences over both inflation and unemployment. The model solution differs from (19) in two ways. *First*, although the relationship between π_t and u_t is still positive and nonlinear, the functional form is slightly different. *Second*, along with $E_{t-1}u_t$, the conditional variance of unemployment becomes one of the arguments of the nonlinear function. Since the tests for neglected ARCH reported in table 5, indicate that unemployment is conditionally homoskedastic, this amounts to including a second constant term in the nonlinear function. The restricted intercept $-(\alpha/2)\sigma_\pi^2$ remain unchanged. Hence, this generalization of the model, does not appear to fundamentally alter the model predictions. In related research, Ruge-Murcia (2000) estimates a model with asymmetric unemployment preferences for the G7 countries. However, estimates of the reduced-form parameters do not support the hypothesis as asymmetric unemployment preferences for Canada and the United Kingdom.

The estimates reported above are based on the "targeted" inflation rates. That is, the

[26] Monte-Carlo experiments (not reported) indicate that when the residuals are positively autocorrelated, the estimate of the preference parameter α is biased downwards. This explains the results in the previous version of this paper, where only mild evidence of preference asymmetry was reported.

[27] This results are not reported to save space, but are available from the author upon request.

measures of inflation to which the official targets are meant to apply. In the cases of Canada and the United Kingdom, they are (respectively) Core and RPIX inflation, that exclude a number of components usually included in the CPI. In order to (*i*) investigate to what extent results are robust to the measure of inflation and (*ii*) examine if the central banker's preferences are different when cast in terms of broader inflation measures, the models were reestimated for Canada and the United Kingdom using CPI and RPI inflation, respectively. Results are reported in tables 10 and 11. Note that in contrast to previous results, in this case it is not possible to reject the null hypothesis of quadratic preferences. Hence, it appears that although the central banker treats asymmetrically the "targeted" inflation deviation from its goal , the loss associated with a "headline" inflation deviation depends primarily on its magnitude and not on its sign.

Finally, this section examines whether two alternative explanations could account for the empirical results. *First*, note that this paper relaxes the usual linear-quadratic framework in a particular dimension. That is, it relaxes the assumption of a quadratic objective function but preserves the linear constraint (*i.e.,* the expectations-augmented Phillips curve). Alternatively, one could consider a model where the objective function is quadratic but the supply function is nonlinear. This is the strategy followed by Nobay and Peel (2000). However, these authors show analytically that the nonlinearity of the supply schedule yields ambiguous implications for average inflation. Numerical simulations indicate that a convex supply function produces upward, rather than downward, inflation deviations from the target. Only a less-plausible, concave supply function yields deviations consistent with observed inflation.

Second, it could be argued that the positive estimates of α might be the result of misspecification in the statistical distribution of the model disturbances. The assumption of normality is useful in deriving close-form analytical results, but there is no theoretical reason to expect that $\boldsymbol{\xi}_t$ is normally distributed in practice. However, t-tests based on Quasi-Maximum Likelihood standard errors that are robust to deviations from normality, tend to reject the null hypothesis of quadratic preferences more strongly than those based on asymptotic standard errors. Whether misspecification in the distribution of the shocks could affect the point estimate of α was examined by means of Monte-Carlo experiments. The data is assumed to be generated by the following (quadratic) specification:

$$u_t = 0.5 + 0.9u_{t-1} + w_t,$$

and

$$\pi_t - \tilde{\pi}_t = 0.8E_{t-1}u_t + \epsilon_t,$$

where $w_t \sim N(0, 0.2^2)$. The disturbance ϵ_t is drawn from an asymmetric chi-square distribution with one degree of freedom and is independent of w_t.[28] Since the mean of a chi-square distribution equals the number of degrees of freedom, the distribution of ϵ_t is centered around zero by substracting -1 from each draw. Recall that the variance of a chi-square distribution is twice the number of degrees of freedom. The values of the parameters where chosen to be roughly on line with the estimates reported above. For example, the unconditional mean of unemployment is $0.5/(1-0.9) = 5\%$, but its process is highly persistent as suggested by the

[28] Recall that for all estimated models, the correlation coefficient between ϵ_t and w_t, namely $\rho_{\epsilon w}$, was found to be insignificantly different from zero.

autocorrelations presented in table 1. The R^2 associated with the unemployment process is calculated to be $R_u^2 = 0.9^2 = 0.81$. Experiments are based on 100 replications using sample sizes of 2000 and 200 observations. For the unemployment process, 100 extra observations were generated in every replication. Then, for the estimation of the model, the initial 100 observations were discarded in order to limit the effect of starting values used to generate the observations of u_t.

Tables 12 and 13 present the results of estimating the quadratic model with asymmetric shocks using an asymmetric-preference model that assumes normally distributed disturbances. In both the large (table 12) and small sample (table 13) experiments, all parameter estimates are close to their true value. The type I error computed for testing the null hypothesis that the parameters take their true values are well within their 95% confidence interval. The only exception is the estimate of σ_π $(= \sigma_\epsilon)$ where there is considerable degree of overrejection regardless of the sample size. Focusing more precisely on the preference parameter, note that if the true model were quadratic ($\alpha = 0$), estimates using a misspecified normal distribution when the true distribution is asymmetric would be unlikely to lead to the conclusion that the central banker's preferences are asymmetric.

4 Conclusions

This paper makes two contributions to the game-theoretical literature on optimal monetary policy. *First*, it constructs a tractable model that generalizes the usual quadratic loss function to allow asymmetries in the preferences of the central banker. Rather that being concerned only with the size of the inflation deviations from the target, the policy maker can associate different losses to positive and negative deviations, even if they are of the same magnitude. This more general framework of analysis is important because it addresses recent empirical literature that suggests that asymmetries might be an important feature of the data. It is shown that some of the conclusions obtained under the assumption of symmetry are not robust to the functional form of the central banker's loss function (though some of them are).

Second, and in contrast to most of the related literature, this paper derives and tests the econometric implications of the model on the joint probability distribution of inflation and unemployment. While the quadratic model predicts a linear relation between the inflation deviations from target and the rate of unemployment, the asymmetric model predicts that inflation depends on its conditional mean and is nonlinearly related to unemployment. The quadratic model predicts that the inflation deviations from target should be positive on average. Since this overshooting is rationally anticipated by the public, the policy can only be imperfectly credible. Under asymmetric preferences, inflation can be on average above of below the target depending on the parameters of the loss function. Data from three inflation targeting countries, namely Canada, Sweden, and the United Kingdom, provides limited but encouraging support for notion of asymmetric preferences. This suggest that departures from the linear-quadratic framework could relevant in actual policy making and that inflation targeting might be a credible framework for the conduct of monetary policy.

Table 1. Autocorrelations of the Rate of Unemployment

Autocorrelation	Country		
	Canada	Sweden	UK
1	0.966	0.935	0.987
2	0.930	0.875	0.971
3	0.892	0.823	0.954
4	0.849	0.773	0.937
5	0.803	0.724	0.917
6	0.751	0.685	0.896
7	0.702	0.639	0.875
8	0.657	0.598	0.851
9	0.614	0.556	0.827
10	0.577	0.512	0.802

Notes: The statistics were calculated using quarterly, seasonally-adjusted data for the periods 1990:01 to 2000:06 (Canada), 1993:01 to 2000:06 (Sweden), and 1991:01 to 2000:06 (United Kingdom).

Table 2. Summary of Predictions

Prediction	Model	
	Quadratic	Asymmetric
Positive relation between $\pi_t - \tilde{\pi}_t$ and u_t	Yes	Yes
Nonlinear relation between $\pi_t - \tilde{\pi}_t$ and u_t	No	Yes
Average deviation from target is nonnegative	Yes	Not necessarily
Conditional variance helps forecast $\pi_t - \tilde{\pi}_t$	No	Yes

Table 3. Percentage of Observations Above and Below Target and Average Inflation Deviation from Target

	Country		
	Canada	Sweden	UK
Average Deviation	−0.61	−1.23	0.15
Above	8.79	18.18	69.89
Below	91.21	81.82	30.11
Observations	91	66	93

Notes: The statistics were calculated using monthly observations for the periods 1992:12 to 2000:06 (Canada), 1995:01 to 2000:06 (Sweden), and 1992:10 to 2000:06 (United Kingdom).

Table 4. Results of Two-Stage Least Squares Regressions

	Country and Model					
	Canada		Sweden		UK	
	Linear	Nonlinear	Linear	Nonlinear	Linear	Nonlinear
Intercept	0.23	-8.20^{**}	-1.99^{*}	-12.87^{*}	-0.40^{**}	-1.31^{**}
	(0.32)	(1.76)	(0.81)	(5.46)	(0.11)	(0.45)
u_t	-0.09^{**}	1.79^{**}	0.11	3.42^{*}	0.08^{**}	0.36^{**}
	(0.03)	(0.39)	(0.11)	(1.65)	(0.01)	(0.14)
u_t^2		-0.10^{**}		-0.24^{*}		-0.02^{*}
		(0.02)		(0.12)		(0.01)
R^2	0.07	0.27	0.02	0.07	0.23	0.25

Notes: The figures in parenthesis are standard errors. The superscripts ** and * denote the rejection of the null hypothesis that the coefficient is zero at the 1% and 5% significance levels, respectively.

Table 5. Tests of Unemployment Residuals

Test	Country and $ARIMA$ Model					
	Canada		Sweden		UK	
	$(3,0,0)$	$(2,1,0)$	$(4,0,0)$	$(3,1,0)$	$(4,0,0)$	$(3,1,0)$
DW	1.91	1.96	2.01	2.00	2.00	1.99
BG	1.43	1.25	2.91	2.84	2.47	1.93
ARCH	0.48	0.08	5.60	5.47	2.44	2.81

Notes: DW and BG stand for Durbin-Watson and Breusch-Godfrey, respectively. Under the null hypothesis of no serial correlation up to order 2, the Breusch-Godfrey test statistic is distributed χ^2 with 2 degrees of freedom. Under the null hypothesis of no conditional heteroskedasticity, the LM statistic for neglected ARCH is distributed χ^2 with 2 degrees of freedom. The superscripts ** and * denote the rejection of the null hypothesis at the 1% and 5% significance levels, respectively.

Table 6. Maximum Likelihood Estimates
Canada
Core Inflation

Parameter	Unemployment Model					
	$ARIMA\ (3,0,0)$			$ARIMA\ (2,1,0)$		
	Estimate	S.E.	Robust S.E.	Estimate	S.E.	Robust S.E.
α	4.54	(0.87)	[0.32]	4.53	(0.87)	[0.32]
γ	−0.02	(0.003)	[0.001]	−0.02	(0.003)	[0.001]
σ_π	0.36	(0.02)	[0.02]	0.36	(0.02)	[0.02]
ψ	−0.005	(0.13)	[0.10]	−0.08	(0.02)	[0.02]
θ_1	0.82	(0.10)	[0.10]	−0.18	(0.10)	[0.10]
θ_2	−0.11	(0.13)	[0.14]	−0.28	(0.10)	[0.10]
θ_3	0.28	(0.10)	[0.10]			
σ_w	0.21	(0.02)	[0.02]	0.21	(0.02)	[0.02]
$\rho_{\epsilon w}$	−0.04	(0.10)	[0.10]	−0.04	(0.10)	[0.11]
L	−133.39			−133.28		

Notes: S.E. and L denote, respectively, the standard error and the value of the joint log-likelihood function at the maximum.

Table 7. Maximum Likelihood Estimates
Sweden
CPI Inflation

Parameter	Unemployment Model					
	$ARIMA\ (4,0,0)$			$ARIMA\ (3,1,0)$		
	Estimate	S.E.	Robust S.E.	Estimate	S.E.	Robust S.E.
α	3.42	(2.58)	[0.47]	3.42	(2.59)	[0.48]
γ	0.45	(2.97)	[0.08]	0.45	(2.98)	[0.09]
σ_π	1.07	(0.09)	[0.07]	1.07	(0.09)	[0.07]
ψ	−0.06	(0.22)	[0.21]	−0.04	(0.03)	[0.03]
θ_1	0.80	(0.12)	[0.15]	−0.20	(0.12)	[0.15]
θ_2	0.16	(0.15)	[0.18]	−0.04	(0.12)	[0.12]
θ_3	0.32	(0.15)	[0.17]	0.27	(0.12)	[0.12]
θ_4	−0.27	(0.13)	[0.13]			
σ_w	0.26	(0.02)	[0.02]	0.26	(0.02)	[0.02]
$\rho_{\epsilon w}$	0.05	(0.12)	[0.13]	0.05	(0.12)	[0.13]
L	−17.75			−17.75		

Notes: S.E. and L denote, respectively, the standard error and the value of the joint log-likelihood function at the maximum.

Table 8. Maximum Likelihood Estimates
UK
RPIX Inflation

Parameter	Unemployment Model					
	$ARIMA\ (4,0,0)$			$ARIMA\ (3,1,0)$		
	Estimate	S.E.	Robust S.E.	Estimate	S.E.	Robust S.E.
α	2.64	(1.52)	[0.57]	2.62	(1.51)	[0.56]
γ	0.08	(0.06)	[0.01]	0.08	(0.06)	[0.01]
σ_π	0.36	(0.03)	[0.03]	0.36	(0.03)	[0.03]
ψ	0.005	(0.03)	[0.03]	−0.03	(0.01)	[0.01]
θ_1	1.05	(0.10)	[0.11]	0.06	(0.10)	[0.11]
θ_2	0.35	(0.15)	[0.14]	0.40	(0.09)	[0.10]
θ_3	−0.23	(0.15)	[0.16]	0.17	(0.10)	[0.12]
θ_4	−0.18	(0.10)	[0.11]			
σ_w	0.08	(0.006)	[0.007]	0.07	(0.006)	[0.006]
$\rho_{\epsilon w}$	0.14	(0.11)	[0.15]	0.10	(0.11)	[0.15]
L	−244.42			−243.76		

Notes: S.E. and L denote, respectively, the standard error and the value of the joint log-likelihood function at the maximum.

Table 9. Tests of Inflation Residuals and Overidentifying Restrictions

Test	Country and $ARIMA$ Model					
	Canada		Sweden		UK	
	$(3,0,0)$	$(2,1,0)$	$(4,0,0)$	$(3,1,0)$	$(4,0,0)$	$(3,1,0)$
LM (Overidentifying)	1.00	0.73	0.06	0.05	6.13	5.80
LM (Intercept)	1.83	1.67	0.0002	0.0001	0.20	0.22
DW	0.33^{**}	0.33^{**}	0.09^{**}	0.09^{**}	0.31^{**}	0.31^{**}
Portmanteau	61.37^{**}	62.34^{**}	57.04^{**}	57.04^{**}	58.36^{**}	58.36^{**}

Notes: DW stands for Durbin-Watson. The superscripts ** and * denote the rejection of the null hypothesis at the 1% and 5% significance levels, respectively.

Table 10. Maximum Likelihood Estimates
Canada
CPI Inflation

Parameter	Unemployment Model					
	$ARIMA\ (3,0,0)$			$ARIMA\ (2,1,0)$		
	Estimate	S.E.	Robust S.E.	Estimate	S.E.	Robust S.E.
α	-0.25	(0.75)	[0.19]	-0.25	(0.75)	[0.20]
γ	-0.10	(0.06)	[0.01]	-0.10	(0.06)	[0.01]
σ_π	0.81	(0.06)	[0.06]	0.81	(0.06)	[0.06]
ψ	-0.11	(0.17)	[0.17]	-0.08	(0.02)	[0.02]
θ_1	0.81	(0.10)	[0.11]	-0.19	(0.10)	[0.11]
θ_2	-0.11	(0.12)	[0.13]	-0.30	(0.10)	[0.10]
θ_3	0.30	(0.10)	[0.10]			
σ_w	0.21	(0.02)	[0.02]	0.21	(0.02)	[0.02]
$\rho_{\epsilon w}$	0.17	(0.11)	[0.10]	0.16	(0.10)	[0.10]
L	-71.40			-71.39		

Notes: S.E. and L denote, respectively, the standard error and the value of the joint log-likelihood function at the maximum.

Table 11. Maximum Likelihood Estimates
UK
RPI Inflation

Parameter	Unemployment Model					
	$ARIMA\ (4,0,0)$			$ARIMA\ (3,1,0)$		
	Estimate	S.E.	Robust S.E.	Estimate	S.E.	Robust S.E.
α	-1.46	(1.05)	[0.74]	-1.38	(1.01)	[0.77]
γ	-0.07	(0.09)	[0.05]	-0.06	(0.08)	[0.05]
σ_π	0.76	(0.06)	[0.04]	0.76	(0.06)	[0.04]
ψ	-0.003	(0.03)	[0.03]	-0.02	(0.01)	[0.01]
θ_1	1.06	(0.10)	[0.10]	0.06	(0.10)	[0.10]
θ_2	0.37	(0.15)	[0.14]	0.43	(0.09)	[0.11]
θ_3	-0.22	(0.15)	[0.15]	0.21	(0.10)	[0.12]
θ_4	-0.22	(0.10)	[0.11]			
σ_w	0.08	(0.006)	[0.007]	0.08	(0.006)	[0.006]
$\rho_{\epsilon w}$	0.16	(0.10)	[0.09]	0.16	(0.10)	[0.09]
L	-174.15			-173.80		

Notes: S.E. and L denote, respectively, the standard error and the value of the joint log-likelihood function at the maximum.

Table 12. Monte Carlo Results
n=2000

Parameter	True Value	Estimate	S.E.	Type I Error	S.E.
α	0	0.0005	(0.059)	0.03	(0.017)
γ	0.8	0.811	(0.113)	0.04	(0.020)
σ_π	1.41	1.414	(0.022)	0.40	(0.049)
ψ	0.5	0.507	(0.049)	0.03	(0.017)
θ_1	0.9	0.899	(0.010)	0.03	(0.017)
σ_w	0.2	0.199	(0.003)	0.05	(0.022)

Notes: The experiment was based on 100 replications. The nominal size of the test was taken to be 0.05.

Table 13. Monte Carlo Results
n=200

Parameter	True Value	Estimate	S.E.	Type I Error	S.E.
α	0	0.011	(0.209)	0.03	(0.017)
γ	0.8	0.953	(0.679)	0.08	(0.027)
σ_π	1.41	1.388	(0.070)	0.45	(0.050)
ψ	0.5	0.586	(0.166)	0.02	(0.014)
θ_1	0.9	0.883	(0.033)	0.02	(0.014)
σ_w	0.20	0.199	(0.010)	0.06	(0.024)

Notes: The experiment was based on 100 replications. The nominal size of the test was taken to be 0.05.

References

[1] Bai, J. and Ng. S. (1998), "A Test for Conditional Symmetry in Time Series Models," Boston College, *Mimeo.*

[2] Barro, R. and Gordon, D. (1983), "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, **91:** 589-610.

[3] Beetsma, R. and Jensen, H. (1998), "Inflation Targets and Contracts with Uncertain Central Bank Preferences," *Journal of Money, Credit and Banking*, **30**: 384:403.

[4] Bernanke, B. S. and Mishkin, F. S. (1997), "Inflation Targeting: A New Framework for Monetary Policy?," *Journal of Economic Perspectives*, **11:** 97-116.

[5] Bernanke, B. S., Laubach, T., Mishkin, F. S., and Posen, A. S. (1999), *Inflation Targeting: Lessons from the International Experience.* Princeton University Press: Princeton.

[6] Blinder, A. S. (1998), *Central Banking in Theory and Practice*, The MIT Press: Cambridge.

[7] Clarida, R. and Gertler, M. (1997), "How the Bundesbank Conducts Monetary Policy," in *Reducing Inflation,* edited by D. Romer. University of Chicago Press: Chicago.

[8] Clarida, R., Galí, J., and Gertler, M. (1999), "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, **37**: 1661-1707.

[9] Campbell, J. Y. and Perron, P. (1991), "Pitfalls and Opportunities: What Macroeconomists Should Know About Unit Roots," *NBER Macroeconomics Annual,* **6**: 141-200.

[10] Christoffersen, P. F. and Diebold, F. X. (1997), "Optimal Prediction Under Asymmetric Loss," *Econometric Theory*, **13**: 808-817.

[11] Cukierman, A. (2000), "The Inflation Bias Result Revisited ," Tel-Aviv University, *Mimeo.*

[12] Eijffinger, S. C. W., Hoeberichts, M., and Schaling, E. (2000), "Why Money Talks and Wealth Whispers: Uncertainty and Mystique," *Journal of Money, Credit and Banking* **32**: 218-235.

[13] Engle, R. F. (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica,* **50**: 987-1007.

[14] Epstein, L. and Zin, S. (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica* **57**: 937-969.

[15] Fischer, S. (1994), "Modern Central Banking," in *The Nature of Central Banking,* edited by F. Capie, C. Goodhart, S. Fischer, and N. Schnadt. Cambridge University Press: Cambridge.

[16] Friedman, M. (1968), "The Role of Monetary Policy," *American Economic Review*, **58:** 1-17.

[17] Gerlach, S. (1999), "Who Targets Inflation Explicitly?," *European Economie Review* **43**: 1257-1277.

[18] Goldfajn, I. and Valdés, R. O. (1999), "The Aftermath of Depreciations," *Quarterly Journal of Economics,* **114**: 229-262.

[19] Granger, C. W. J. and Pesaran, M. H. (1996), "A Decision Theoretic Approach to Forecast Evaluation," University of Cambridge, *DAE Working Paper No. 9618.*

[20] Green, J. H. (1996), "Inflation Targeting: Theory and Implications," *IMF Staff Papers*, **43**: 779-795.

[21] Gul, F. (1991), "A Theory of Disappointment Aversion," *Econometrica* **59**: 667-686.

[22] Hamilton, J. D. (1983), "Oil and the Macroeconomy since World War II," *Journal of Political Economy*, **91:** 228-248.

[23] Hamilton, J. D. (1994), *Time Series Analysis,* Princeton: Princeton University Press.

[24] Harvey, A. C. (1981), *Time Series Models,* Philip Allan Publishers Limited: Oxford.

[25] Ireland, P. N. (1999), "Does the Time-Consistency Problem Explain the Behavior of Inflation in the United States?", *Journal of Monetary Economics,* **44:** 279-292.

[26] Judd, K. L. (1998), *Numerical Methods in Economics,* MIT Press: Cambridge.

[27] King, M. (1996), "How Should Central Banks Reduce Inflation: Conceptual Issues" in *Achieving Price Stability,* Federal Reserve Bank of Kansas City.

[28] Kydland, F. and Prescott, E. (1977), "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, **85**: 473-490.

[29] Lafrance, R. (1997), "An Overview of the Monetary Frameworks of Four Inflation-Targeting Countries," Bank of Canada, *Mimeo.*

[30] Lucas, R. E. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory*, **4:** 103-124

[31] Lucas, R. E. (1973), "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review*, **63:** 326-334.

[32] Lockwood, B. and Philippopoulos, A. (1994), "Insider Power, Unemployment Dynamics and Multiple Inflation Equilibria," *Economica*, **61**: 59-77.

[33] McCallum, B. T. (1995), "Two Fallacies Concerning Central Bank Independence," *American Economic Review Papers and Proceedings*, **85:** 201-211.

[34] McCallum, B. T. (1997), "Crucial Issues Concerning Central Bank Independence," *Journal of Monetary Economics*, **39:** 99-112.

[35] Muscatelli, V. A. (1999), "Inflation Contracts and Inflation Targets under Uncertainty: Why We Might Need Conservative Central Bankers," *Economica*, **66**: 241-254.

[36] Nobay, R. A. and Peel, D. A. (1998), "Optimal Monetary Policy in a Model of Asymmetric Central Bank Preferences," London School of Economics, *Mimeo.*

[37] Nobay, R. A. and Peel, D. A. (2000), "Optimal Monetary Policy with a Nonlinear Phillips Curve," *Economics Letters*, **67**: 159-164.

[38] Orphanides, A. and Wilcox, D. W. (1996), "The Opportunistic Approach to Disinflation," Federal Reserve Board, *Mimeo.*

[39] Persson, T. and Tabellini, G. (2000), "Political Economics and Macroeconomic Policy," in *Handbook of Macroeconomics,* edited by J. Taylor and M. Woodford. North-Holland: Amsterdam.

[40] Pesaran, M. H. and Ruge-Murcia, F. J. (1999), "Analysis of Exchange Rate Target Zones Using a Limited-Dependent Rational-Expectations Model with Jumps," *Journal of Business and Economics Statistics*, **17:** 50-66.

[41] Rogoff, K. (1985), "The Optimal Degree of Commitment to a Monetary Target," *Quarterly Journal of Economics,* **100**: 1169-1190.

[42] Romer, D. (1996), *Advanced Macroeconomics*, McGraw-Hill: New York.

[43] Ruge-Murcia, F. J. (forthcoming), "Uncovering Financial Markets' Beliefs about Inflation Targets," *Journal of Applied Econometrics.*

[44] Ruge-Murcia, F. J. (2000), "The Inflation Bias when the Central Banker Targets the Natural Rate of Unemployment," University of Montreal, *Mimeo.*

[45] Shimer, R. (1998), "Why is the U.S. UnemploymentRate so Much Lower?," *NBER Macroeconomics Annual:* 33-49.

[46] Staiger, D., Stock, J. H., and Watson, M. W. (1997), "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives,* **11**

[47] Svensson, L. E. O. (1997), "Optimal Inflation Targets, Conservative Central Banks, and Linear Inflation Contracts," *American Economic Review,* **87**: 98-114.

[48] Svensson, L. E. O. (1999), "Inflation Targeting as a Monetary Policy Rule," *Journal of Monetary Economics,* **43**: 607-654.

[49] Tetlow, R. J. (1999), "Inflation Targeting and Target Instability," Board of Governors of the Federal Reserve System, *Mimeo.*

[50] Tootell, G. M. B. (1994), "Restructuring, the NAIRU, and the Phillips Curve," *New England Economic Review of the Federal Reserve Bank of Boston*, September/October: 31-44..

[51] Varian, H. (1974), "A Bayesian Approach to Real Estate Assessment," in *Studies in Bayesian Economics in Honour of L. J. Savage,* edited by S. E. Feinberg and A Zellner. Noth-Holland: Amsterdam.

[52] Walsh, C. (1995), "Optimal Contracts for Independent Central Bankers," *American Economic Review,* **85**: 150-167.

[53] Walsh, C. (1998), *Monetary Theory and Policy,* The MIT Press: Cambridge.

[54] Weiner, S. E. (1993), "The Natural Rate and Inflationary Preasures," *Economic Review of the Federal Reserve Bank of Kansas City,* **79**: 5-9.

[55] White, H. (1982), "Maximum Likelihood Estimation of Misspecified Models," *Econometrica,* **50**: 1-25.

[56] Zellner, A. (1992), "Bayesian Estimation and Prediction Using Asymmetric Loss Functions," *Journal of the American Statistical Association,* **81**: 446-451.

Figure 1: Preferences

(a) Asymmetric Preferences

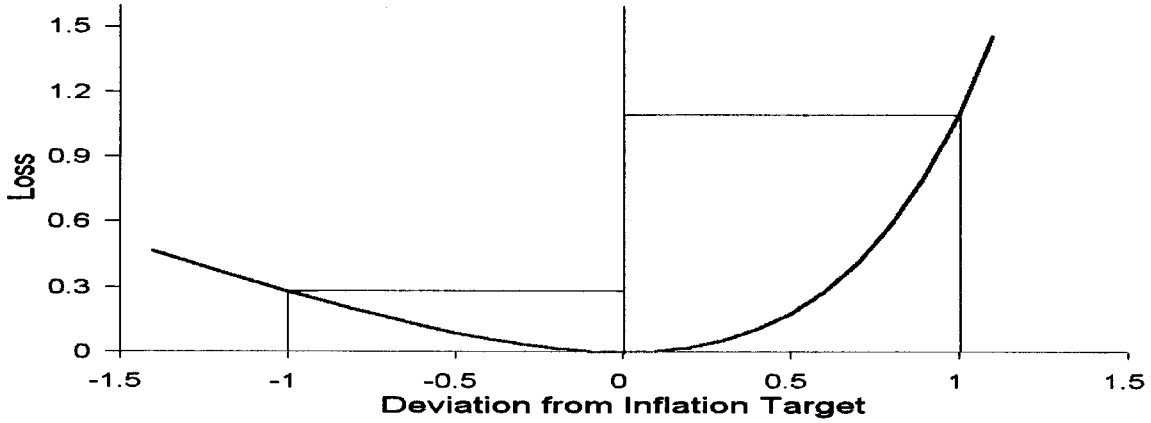

(b) Quadratic Preferences

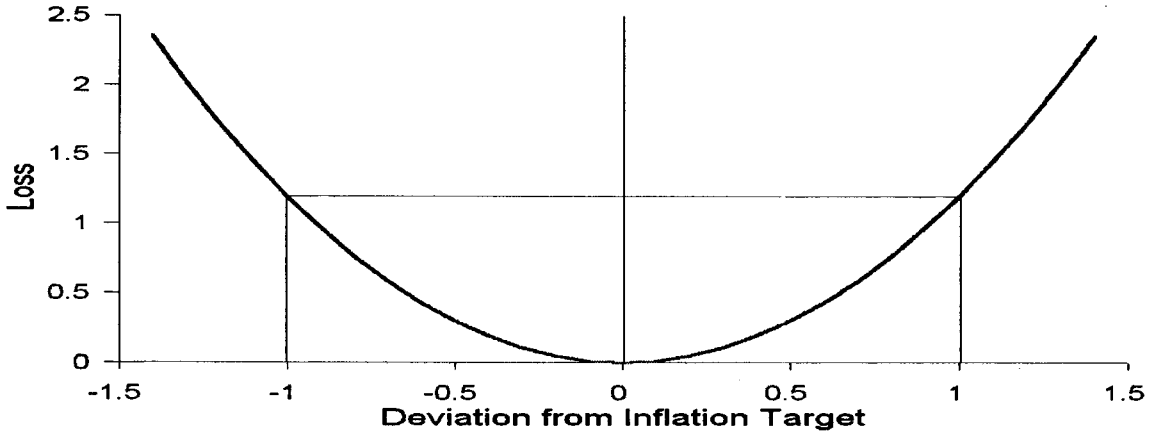

Figure 2: Reaction Functions

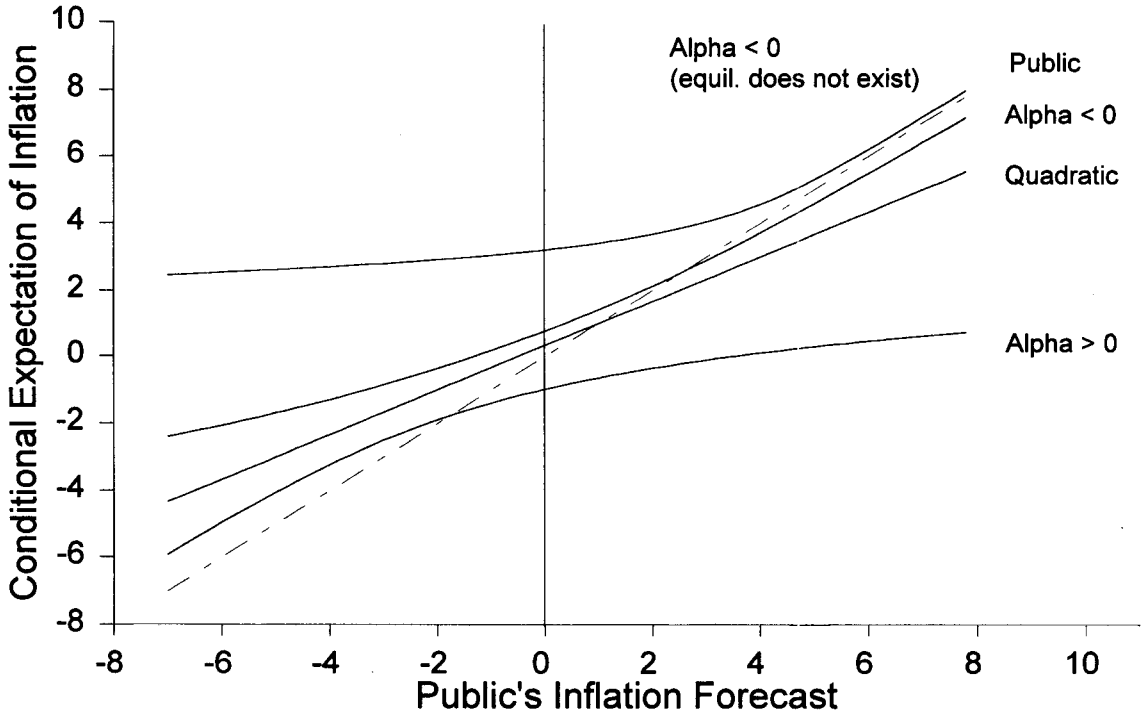

Figure 3: Difference in Predicted Inflation Deviation from Target

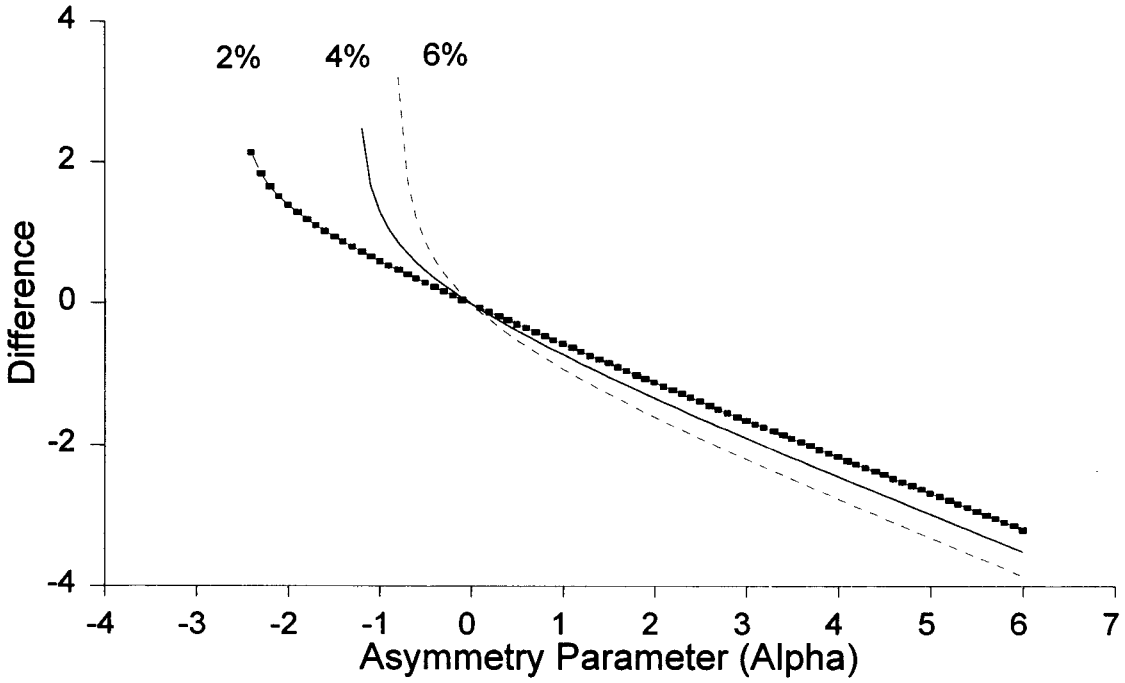

Figure 4: Inflation Rates

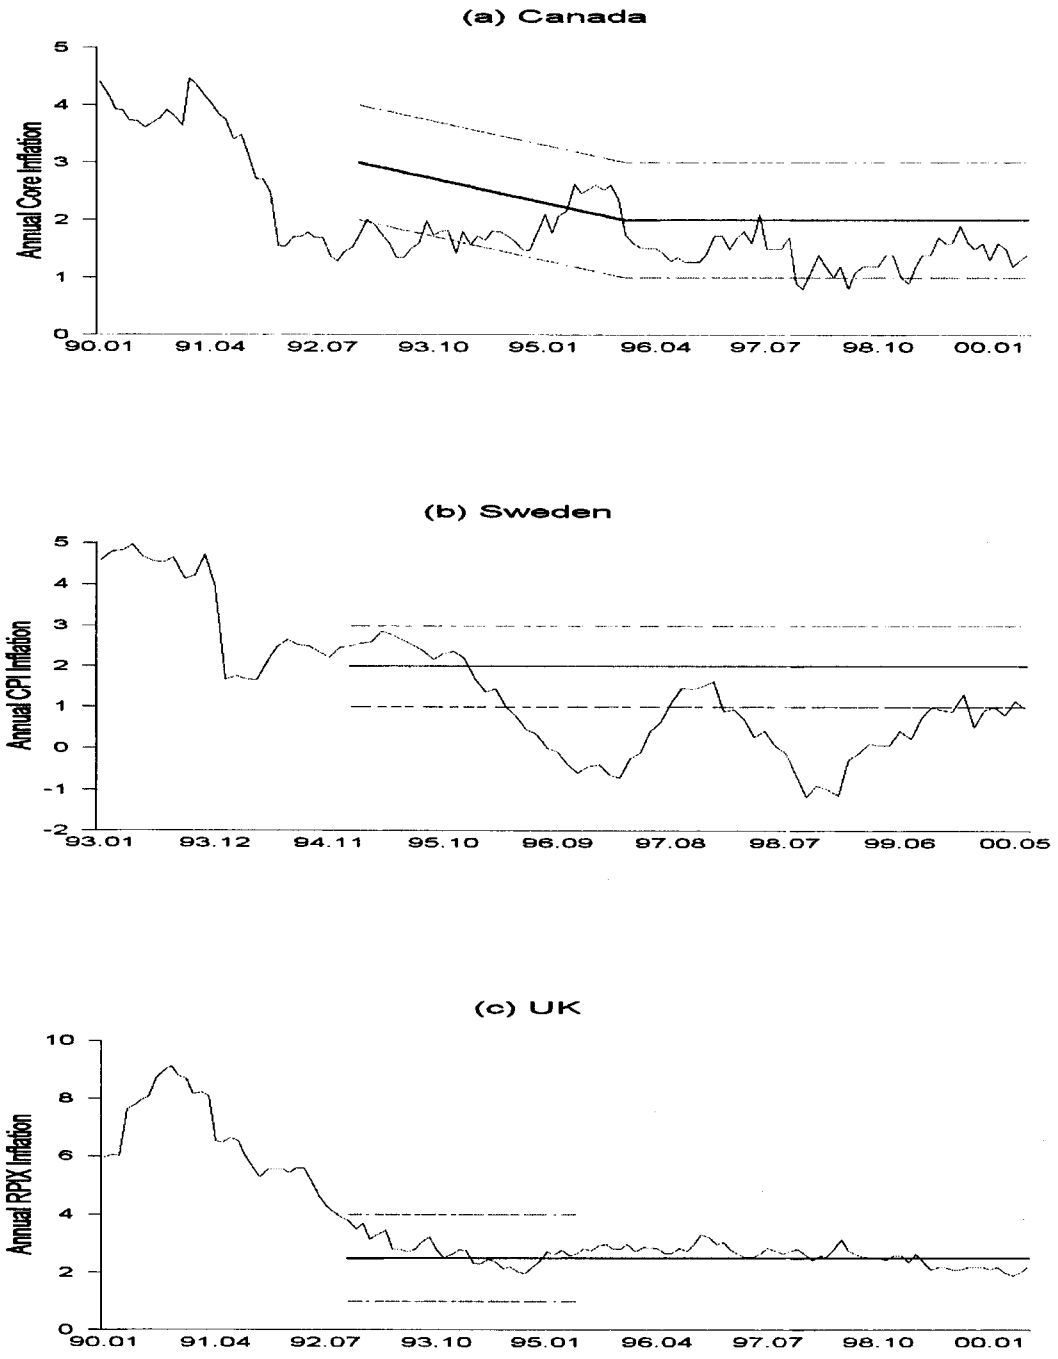

Figure 5: Unemployment and the Inflation Deviation from Target

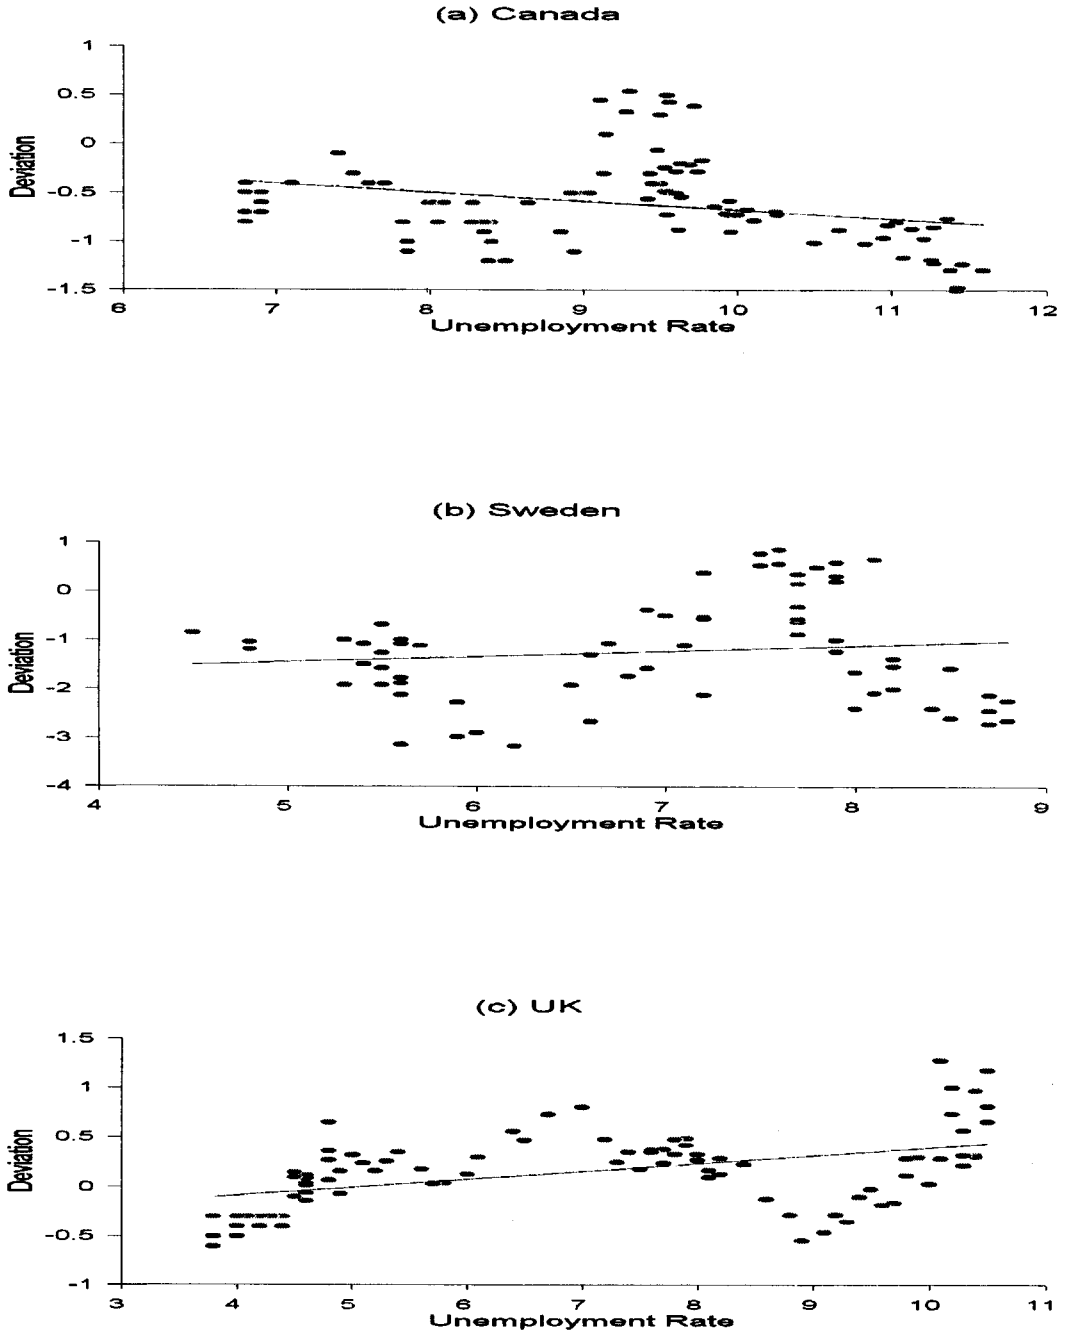

Figure 6: Estimated Preferences

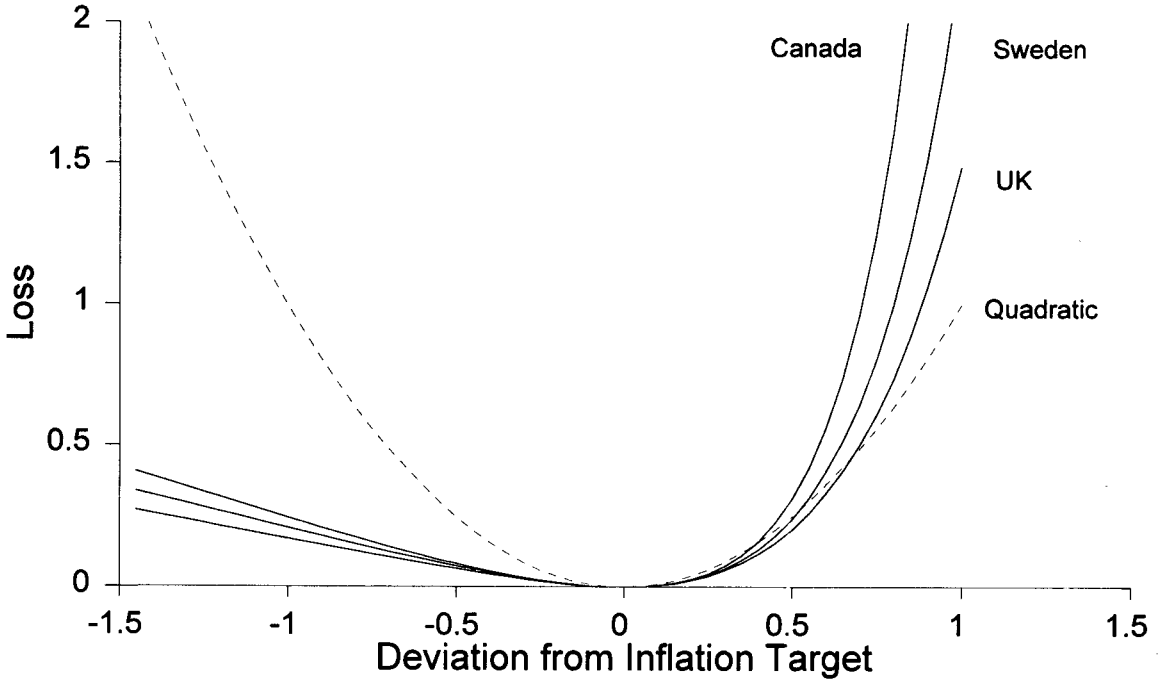

Figure 7: Conditional Variance and the Inflation Deviation from Target