Imperfect Knowledge, Adaptive Learning, and the Bias Against Activist Monetary Policies*

Alberto Locarno
Research Department, Banca d'Italia and London School of Economics

The paper studies the implications for the effectiveness of discretionary monetary policymaking of departing from the assumption of rational expectations. Society, whose welfare function is quadratic, can appoint a central banker whose preferences are either quadratic or lexicographic, to achieve the best mix of inflation and output stability. The focus on lexicographic preferences is justified on the grounds that they imply a strict ordering of policy objectives, which is typical of the mandate of several central banks. Both the private sector and the monetary policymaker have incomplete knowledge of the working of the economy and rely upon adaptive learning to form expectations and decide policy moves. The model economy is assumed to be subject to recurrent unobserved shifts, and the monetary authority, who has private information on the shocks hitting the economy, cannot credibly commit. The main finding of the paper is that when agents rely on an adaptive learning technology, a bias against activist policies arises. The paper also shows that when society has quadratic utility, a strategy based on a strict ordering of objectives is close to optimal for a wide range of values of the inflation aversion parameter.

JEL Codes: E52, E31, D84.

1. Introduction

Effective policymaking requires that the monetary authority commit to a systematic approach to policy. As long as price setting depends

*This paper benefited from conversation with Margaret Bray, Giuseppe Ferrero, Alessandro Secchi, and Marco Vega.

on expectations of the future, a credible central bank may be able to face an improved short-run trade-off between inflation and output and, accordingly, it may reduce inflation at lower costs. However, for policymakers to succeed in steering expectations and reaping the benefits of commitment, agents must be able to fully anticipate the future impact of monetary decisions on the economy, which is feasible only if the economic environment is stationary and expectations are rational. When, on the contrary, the structure of the economy is subject to recurrent shifts and knowledge is imperfect, agents must rely on alternative methods for anticipating future events, and this may dramatically alter the policy trade-offs.

A recent stream of literature, epitomized by Orphanides and Williams (2002), has shown that imperfect knowledge makes stabilization policies more difficult. This happens when central banks put too much weight on output stabilization, because overly activist policies are prone to generate episodes in which the public's expectations of inflation become uncoupled from the policy objective. An additional complication arises from the fact that if knowledge is incomplete, committing to a systematic policy becomes problematic (because the private sector cannot easily verify whether the central bank is delivering on its promises) and the gains from commitment are severely reduced (because expectations, being backward rather than forward looking, cannot be manipulated to increase policy effectiveness).

This paper focuses on how the relaxation of the rational expectations hypothesis changes the way monetary policy is set. It applies a principal-agent approach to the time-inconsistency problem: in order to improve the discretionary equilibrium, society (the principal) assigns a loss function to the central bank (the agent) that may differ from society's preferences. The assumption underlying this approach is that it is possible to commit the monetary authority to a particular loss function, whereas the minimization of the loss function occurs under discretion. Society has standard quadratic preferences on output and inflation, and the central banker is endowed with either a quadratic or a lexicographic preference ordering. The literature on monetary policy almost always relies on the assumption of quadratic preferences, but such a loss function does not reflect the task assigned to the monetary authority of most developed countries and does not account for the role of factors that are vital in actual policymaking.[1] Lexicographic preferences overcome a few of these shortcomings: they capture a hierarchical ordering of alternatives, which is typical of the mandate of virtually every inflation-targeting central bank,[2] and allow the policymaker to focus on different policy objectives under different circumstances.

The main findings of the paper are the following. First, it is confirmed that when agents do not possess complete knowledge about the structure of the economy and rely on an adaptive learning technology, a bias toward conservatism arises, suggesting that society is better off by appointing a policymaker whose degree of inflation aversion is higher than its own; even if there is no intrinsic dynamics in the economy, agents' and policymakers' attempts to learn adaptively introduce inertia in the system, which makes it costly for the central bank not to respond promptly and forcefully to shocks. Second, what matters for society's welfare is that the degree of inflation aversion of the monetary authority is high enough to prevent expectations from fluctuating too much: the specific form the loss function of the central bank takes is of second-order importance. Third, the bias against stabilization policies and toward conservatism, and the relative efficiency of alternative monetary strategies, do not depend on whether the memory of the learning process is finite or infinite.

The paper is related to the literature in several ways. It parallels, under more general conditions, the work of Terlizzese (1999) and Driffil and Rotondi (2003) in deriving the properties of a monetary strategy that has price stability as its primary objective and does not allow inflation to rise above an upper limit; it models two-sided learning in the vein of Evans and Honkapohja (2002); and it uses the study of Orphanides and Williams (2002) as a

[1] The following are just a few of these factors: (i) uncertainty (see al-Nowaihi and Stracca 2002); (ii) imperfect observability of the state variables (see Svensson and Woodford 2000); and (iii) path dependence and differential valuation of deviations from the inflation and output targets (see Orphanides and Wilcox 1996).

[2] Buiter (2006) presents a list of the central banks whose mandates reflect a lexicographic preference ordering, with price stability ranking first and all other desiderata coming after.

benchmark for the quantitative analysis. The most closely related of these contributions is the work of Orphanides and Williams: model simulations are designed so as to replicate their experiments, and the objective of this paper is to a large extent the same as theirs—namely, to understand how the economy responds to alternative monetary strategies when agents have bounded rationality and imperfect knowledge.

The original contribution of this work is to extend the findings by Orphanides and Williams (2002). It tries to do that in three different ways. First, it assumes that not only private agents but also the policymaker have imperfect knowledge; under this framework, policy effectiveness ends up depending on both inflation and output variability so that the bias toward conservatism, if confirmed, cannot be attributed to the limited role of output volatility in the reference model. Second, it tests whether society can increase welfare by appointing a policymaker whose preference ordering is lexicographic. Third, it tests the claim about the negative impact of imperfect knowledge on economic stabilization under a set of alternative learning mechanisms.

The paper is organized as follows. Section 2 outlines the model used in the paper and contrasts the implications of assuming quadratic or, alternatively, lexicographic preferences. Section 3 introduces econometric learning and studies how different policies affect the speed at which learning algorithms converge to the rational expectations equilibrium. Section 4 presents some evidence, obtained by means of simulation, on the distortions on monetary policymaking caused by assuming that agents have bounded rationality; the focus is on whether adaptive learning induces a bias toward conservatism and whether appointing a central banker whose preferences are lexicographic is welfare improving. Section 5 concludes.

2. The Model

The model presented in this section has two basic components: (i) the unobservability of the supply shock and (ii) an unknown and time-varying output-inflation trade-off. The model is first solved under rational expectations and then under adaptive learning.

2.1 The Structure of the Economy and the Delegation Problem

The economy is characterized by an expectations-augmented Phillips-curve relationship, linking inflation surprises $\pi - \pi^e$ to the output gap y:

$$y = \alpha(\pi - \pi^e) + \varepsilon. \tag{1}$$

Inflation is the policy instrument and is controlled without error by the monetary authority; the natural level of output is normalized at 0. Output also responds to a zero-mean supply shock ε, unobservable to the central bank and the private sector and uniformly distributed on the closed interval $[-\mu, \mu]$. A signal z, conveying noisy information on ε, is observed by the policymaker after expectations have been determined; it is assumed that $z = \varepsilon + \xi$, with ξ following a uniform distribution with the same support as ε, i.e., $\xi \sim U[-\mu, \mu]$.[3]

The final component of the model is the assumption that α, the output-inflation trade-off, is a random variable. Since α is time varying, the effects of monetary policy on output depend on the value of the trade-off. It is assumed that $\alpha = \bar{\alpha} + \widetilde{\alpha} \sim IID(\bar{\alpha}, \sigma_\alpha^2)$ and that it is independent of all the other shocks in the economy.[4] Notice that the model is entirely static, so that no issue of strategic interaction between the monetary authority and the private sector arises.

The timing of the model is shown in figure 1. The signal z materializes before the central bank chooses the inflation rate but after private agents have set their inflation expectations for the period. The information advantage of the central bank creates a role for

[3] The assumption that both ε and ξ follow a uniform distribution ensures that in the rational expectations equilibrium a closed-form solution exists. The additional hypothesis that both shocks share the same support helps to keep the distribution of z simple.

[4] The stochastic variable α can be interpreted as an index of monetary policy effectiveness. It can be either discrete or continuous. What is relevant is the i.i.d. assumption, which avoids introducing dynamic elements and strategic interactions into the optimization problem of the central bank. Ellison and Valla (2001) show that strategic interactions create a connection between the activism of the central bank and the volatility of inflation expectations: the latter reacts to the former because an activist policy produces more information, helping private agents learn. The value of experimentation in policymaking is also studied in Wieland (2003).

Figure 1. Timing of the Model

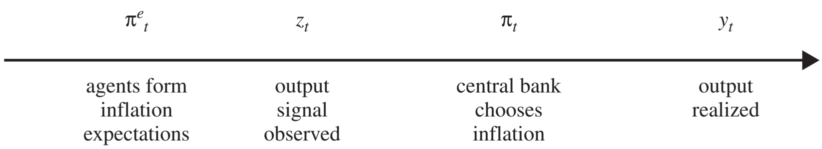

stabilization policies and takes into account the fact that policy decisions can be made more frequently than most wage and price decisions.

Society has quadratic utility and dislikes both inflation and output variability. Its welfare function is

$$W^S = -E[(y-k)^2 + \beta^S \pi^2], \tag{2}$$

where k is the target level of output; the expectation operator is due to the unobservability of the output-inflation trade-off α and output shock ε. The assumption that $k > 0$ is usually justified on the grounds that the presence of distortions in the labor and goods markets leads to an inefficiently low level of output in equilibrium; alternatively, $k > 0$ is interpreted as arising from political pressures on the central bank, as in Mishkin and Westelius (2006). In the principal-agent approach, society, whose preferences are quadratic, can appoint a central banker whose loss function may differ from its own. It is assumed that the choice is restricted to policymakers with either quadratic or lexicographic preferences; in the first case, the inflation aversion parameter β can be different from β^S.

2.2 Central Bank Loss Function

Though in general a lexicographic-preference ordering cannot be represented by a function, in the simplified case in which the monetary authority has been given only two objectives, such an ordering can be described by a loss function involving only the secondary objective, subject to a constraint involving the primary target. It is therefore assumed that the central bank aims at stabilizing output around a nonzero level, provided that inflation is kept below a known upper bound.[5] Though there is some dispute about the correctness of formulations depicting central bankers as affected by an inflation bias, the assumption is retained because otherwise the only rational expectation for inflation would be the zero target itself and the inflation constraint would never be binding.

In formal terms, the problem solved by the central bank is[6]

$$\begin{array}{ll} \min_\pi & \frac{1}{2}E(y-k)^2 \\ s.t. & \begin{cases} \pi \leq \bar{\pi} \\ y = \alpha(\pi - \pi^e) + \varepsilon. \end{cases} \end{array} \tag{3}$$

Notice that k cannot exceed μ, the upper bound of the output shock: in that which follows, it will be assumed that k is not too high and, in particular, that $k = \frac{\mu}{6}$. Under the standard hypothesis of time separability of preferences, the problem is static and involves no trade-off between current and future utility, so that the optimal policy does not have to rely on the strategic interactions described by Bertocchi and Spagat (1993), Ellison and Valla (2001), and Wieland (2003).

To highlight the implications of endowing the monetary authority with lexicographic preferences, the policy problem is also analyzed under the standard assumption that the loss function is quadratic. In this case, the problem solved by the central bank can be formulated as

$$\begin{array}{ll} \min_\pi & \frac{1}{2}E[(y-k)^2 + \beta\pi^2] \\ s.t. & y = \alpha(\pi - \pi^e) + \varepsilon, \end{array} \tag{4}$$

[5]The existence of a lower bound on inflation is neglected in this paper. In a model where the output shock is observable and the trade-off between output and inflation is time invariant, Terlizzese (1999) shows that the main features of the monetary policy problem are largely unaffected by the inclusion of a lower bound on inflation. Intuitively, what explains this result is the asymmetric nature of the inflation bias that is assumed to characterize the monetary authority's preferences: if the central bank aims at pushing output above the natural level, it will tend to inflate, so that while the upper bound will often be binding, the lower one will not.

[6]Buiter (2006) provides an alternative, more restrictive representation of a lexicographic-preference ordering.

where β measures the weight attached to the inflation objective relative to output stabilization.

Regardless of the specific form of the loss function, be it (3) or (4), the model features an inflation bias, arising from the policymaker's incentive to create surprise inflation in order to keep output above the natural level. Many economists, however, think that such a feature makes the model irrelevant for monetary policy analysis.[7] Jensen (2003) argues that the Barro-Gordon-type models can be rescued by noting that their main implications—notably, the inflation bias result—can be maintained without having to resort to the presumption that the central bank engenders inflation surprises to fool the public.

2.3 Signal Extraction and the Rational Expectations Equilibrium

Given the structure of the problem, the issues of estimating the unobserved output shock and setting the optimal inflation rate can be kept separated and solved sequentially. Before deciding the optimal policy, the central bank has to solve a signal extraction problem. The first step is therefore to derive the probability distribution of $z = \varepsilon + \xi$ and the conditional mean $E(\varepsilon|z)$. In lemma 1 the density function of the signal z is derived, while in proposition 1 the first moment of the distribution of the output shock ε conditional on z is calculated.

LEMMA 1. *If $z = \varepsilon + \xi$ and ε and ξ are independent uniform random variables, with support on the interval $[-\mu, \mu]$, then the density function of z is equal to $f(z) = \frac{1}{2\mu} + \frac{1}{4\mu^2}[\min(z, 0) - \max(0, z)]$.*

Proof. See Locarno (2006).

PROPOSITION 1. *If ε and ξ are uniform random variables, defined on the same close interval $[-\mu, \mu]$, and $z = \varepsilon + \xi$, then the optimal estimate of ε conditional on z is $E(\varepsilon|z) = \frac{z}{2}$.*

Proof. See Locarno (2006).

[7] See, for instance, the quotations from Blinder, Vickers, and Issing listed in Jensen (2003).

Given the assumption regulating the flow of information, the central bank sets the inflation rate on the basis of the observed signal and the private-sector inflation expectations. Under lexicographic preferences, it will choose the inflation rate that solves the first-order condition $E[\alpha(\alpha(\pi - \pi^e) + \varepsilon - k) \mid z] = 0$—provided the inflation constraint is satisfied—and will choose $\pi = \bar{\pi}$ otherwise, i.e.,

$$\pi = \begin{cases} \pi^e - \dfrac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2}\left(\dfrac{z}{2} - k\right) = \pi^e - \dfrac{z}{\phi} + \dfrac{2k}{\phi} & \text{if } z \geq 2k + \phi(\pi^e - \bar{\pi}), \\ \bar{\pi} & \text{otherwise,} \end{cases} \quad (5)$$

where $\phi^{-1} \equiv \frac{\bar{\alpha}}{\bar{\alpha}^2+\sigma_\alpha^2}\frac{1}{2}$. The optimal policy depends, in a nonlinear way, on the value of z: when output shocks are strongly negative and the primary objective is at risk, the central bank acts as an inflation nutter; when the signal indicates more-favorable disturbances, it displays more activism, favoring output stabilization. Notice that the optimal policy depends on the parameters of the distribution of the output-inflation trade-off α. Two cases are considered, one corresponding to the rational expectations equilibrium (REE), and the other assuming bounded rationality and least-squares learning. In the first case, it is assumed that α is not observed but $\bar{\alpha}$ and σ_α^2 are known by both the central bank and the private sector;[8] in the second case, $\bar{\alpha}$ and σ_α^2 are unknown and must be estimated.

A few points illustrating the main properties of optimal policy are worth stressing. First, the optimal policy is not altered by the unobservability of the output shock, except that the policymaker responds to an efficient estimate of the state variable rather than to its actual value. It is a well-known fact that a linear model with a quadratic loss function and a partially observable state of the economy is characterized by certainty equivalence; since the assumed preference ordering is not quadratic, the result applies only when inflation is within the admissible range. Second, uncertainty about the multiplier of the policy instrument makes it optimal to react less than completely to the output shock. There is nothing to be gained

[8]This assumption could be justified if ε, though unobserved at the time when expectations and the inflation rate were set, was observed with a one-period lag.

by adopting more-activist policies in order to learn from experimentation, since the model is static and the loss in current welfare incurred in overreacting will not be compensated by future gains. The reduction in policy activism caused by parameter uncertainty, originally shown by Brainard (1967), reflects the direct impact of the monetary authority's instrument on the variability of the target variable.

For the equilibrium to be fully characterized, the solution for expected inflation must be provided. Under rational expectations, agents understand the incentives driving the actions of the central bank and have expectations that coincide on average with realizations. Accordingly, $\pi^e = \int \pi dF(z)$, where $F(z)$ is the distribution function of the signal z. Proposition 2 gives the full characterization of the rational expectations equilibrium under the simplifying assumption that $k = \frac{\mu}{6}$.[9]

PROPOSITION 2. *If the central bank has lexicographic preferences and output is determined as in (1), there is a unique REE, where* $\pi = \min\left[\pi^e - \frac{\bar{\alpha}}{\bar{\alpha}^2+\sigma_\alpha^2}\left(\frac{z}{2} - k\right), \bar{\pi}\right]$ *and* $\pi^e = \bar{\pi} - \frac{2k}{\phi} = \bar{\pi} - \frac{\mu}{3\phi}$.

Proof. See Locarno (2006).

Equilibrium is noncooperative Nash: the central bank and the private sector try to maximize their respective objective functions, taking as given the other player's actions. The assumption of rational expectations implicitly defines the loss function of the private sector as $E(\pi - \pi^e)^2$: given the public's understanding of the central bank's decision problem, its choice of π^e is the one minimizing disutility.

From the expression for π^e, it is apparent that the existence of an upper bound on inflation contributes to stabilizing expected inflation: for any value of k, the lower $\bar{\pi}$ is, the lower expected inflation is. Another feature of the policy is that the larger the support of the output shock, the closer to zero π^e. The reason for this result is straightforward: positive (and higher than k) output shocks trigger a

[9]Setting $k = \frac{\mu}{6}$ amounts to assuming that $2k + \phi(\pi^e - \bar{\pi}) = 0$ and implies that the central bank chooses its best static response when $z > 0$, while it has no discretion for negative values of the signal.

reaction from the central bank, which creates negative inflation surprises to stabilize output; large negative disturbances, on the other hand, cannot be neutralized, because too-high inflation rates are not admissible. Widening the support of ε has an asymmetric effect on the actions of the monetary authority: it increases the cases in which the central bank finds it optimal to deflate, but has no influence on its incentives to inflate.

Two features of the optimal monetary policy are worth stressing. First, for values of the signal in the nonempty interval $[-2\mu, 0)$, inflation is constant and equal to $\bar{\pi}$, which is higher than π^e: the central bank keeps price dynamics above inflation expectations and, in so doing, it sustains output, though it cannot cushion it against shocks. Second, even in the face of favorable output shocks, the policymaker is unable to fully stabilize output at the desired level. Two factors contribute to attenuate the policymaker's response: uncertainty about the output-inflation trade-off and unobservability of ε. The former reduces the response of the central bank by a factor of $\frac{2\alpha}{\phi}$, while the latter leaves part of the output shock, namely, $\varepsilon - \frac{z}{2}$, unchecked. Since $\frac{z}{2}$ is an unbiased estimate of ε, unobservability of the output shock increases volatility but does not affect the degree of activism of the policy response; on the contrary, unobservability of the output-inflation trade-off has a bearing on the policy strategy, since it favors more cautious policies. It turns out that output volatility is therefore smaller than $\frac{\mu^2}{3}$, the variance of the output shock, implying some degree of stabilization on the side of monetary policy.

To assess the distinguishing traits of the policy pursued by a central bank with lexicographic preferences, it is useful to contrast it with the optimal policy arising under the standard assumption of quadratic loss function. If (4) describes the central bank's problem, the policy instrument is set according to the rule

$$\pi = \frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta}\pi^e + \frac{\bar{\alpha}}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta}\left(k - \frac{z}{2}\right) = \rho\pi^e + \frac{\rho}{\phi}(2k - z)$$
$$= \pi^e - \frac{\rho}{\phi}z, \tag{6}$$

where $\rho \equiv \frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}^2 + \sigma_\alpha^2 + \beta}$, with $0 < \rho \leq 1$.

A few differences are apparent when comparing the two optimal policies. When output shocks are not too negative, rule (5) ensures more output stabilization, as the increase in the inflation rate that must be engendered to counteract the supply disturbance does not have a negative impact on welfare and hence does not bring on a trade-off between the output and the inflation objectives. The reverse is true when ε is large and negative, because in that case output stabilization is sacrificed to the primary objective of price stability. More activist policies are possible at the cost of larger inflation variability: in general, there exists a value $\bar{\beta}$ such that, for $\beta \in (\bar{\beta}, \infty)$, output variability under lexicographic preferences (henceforth, strategy 1) is lower than under quadratic utility (henceforth, strategy 2), and there exists a value $\underline{\beta}$ such that, for $\beta \in [0, \underline{\beta})$, $E(\pi - \pi^e)^2$ is smaller under strategy 1.[10] Since $\underline{\beta} < \bar{\beta}$, strategy 1 apparently cannot outperform strategy 2 in terms of both objectives, but in reality this is not necessarily the case, since what is important for social welfare is not $E(\pi - \pi^e)^2$ but $E\pi^2$. Under strategy 1 the central bank can use an additional instrument—the upper bound on inflation $\bar{\pi}$—to try to achieve both lower output volatility and lower mean-square inflation: the reason that strategy 1 is more appealing compared with Rogoff's (1985) solution is that the reduction of the inflation bias does not come at the cost of the output-stabilization objective;[11] the reason it is less valuable is that it tends to stabilize output too much when $\pi < \bar{\pi}$.

3. Adaptive Learning and Monetary Policy Regimes

In the real world, where shifts in policies and in the economic structure are by no means rare events, people often face the problem of understanding whether and how the environment has changed and which is the least costly way to adapt decision rules to suit the new framework. In such a context, a strict application of the rational expectations hypothesis (REH) would not be a convincing

[10] See Locarno (2006), proposition 4, for a formal proof.

[11] This is easily seen by noting that the choice of $\bar{\pi}$ has no effect on the set of values of z corresponding to inactive policies. This implies that it is always optimal for the policymaker to set $\bar{\pi}$ so that π^e is equal to 0.

theoretical solution. Alternatives have long been suggested. Herbert Simon (1957), for instance, supported some kind of bounded rationality and proposed to create a theory with behavioral foundations where agents learn in the same way as econometricians do. An increasingly important stream of literature, which builds on the pioneering work of Bray (1982) and Marcet and Sargent (1989), recently revived in particular by Evans and Honkapohja (2001), has introduced a specific form of bounded rationality, called adaptive learning, where agents adjust their forecast rule as new data becomes available over time. This approach provides an asymptotic justification for the REH and makes nonlearnable solutions in models with multiple equilibria irrelevant.

The central idea behind adaptive learning is that at each period t, private agents have a perceived law of motion (PLM) that they use to make forecasts. The PLM relates the variables of interest, whose future values are to be anticipated, to a set of state variables; the projection parameters are estimated using least squares. Forecasts generated in this way are used in decisions for period t, which yields the temporary equilibrium, also called the actual law of motion (ALM). The temporary equilibrium provides a new data point, and agents are then assumed to reestimate the projection parameters with data through period t and to use the updated forecast functions for period $t + 1$ decisions. The learning dynamics continue with the same steps in subsequent periods.

The updating of the projection coefficients may be represented in terms of a system of recursive equations having, under reasonable assumptions for the PLM, the rational expectations equilibrium as a fixed point. The recursive equations describe the mapping between the PLM and the ALM. Convergence of the adaptive learning process may be studied by means of the associated ordinary differential equation (ODE):[12] stability holds whenever the real parts of the eigenvalues of the Jacobian of the ODE are negative, i.e., whenever the system is E-stable.

[12]For large t, the stochastic recursive algorithm is well approximated by an ordinary differential equation, provided a few regularity conditions are met. These regularity conditions involve the stochastic process driving the state variables, the deterministic gain sequence, and the function governing the revision in the projection coefficients.

The connection between E-stability and the convergence of least-squares learning is a great advantage, since E-stability conditions are often easier to work out. However, focusing on asymptotic approximations puts aside any consideration of the learning speed. Benveniste, Metiver, and Priouret (1990) show that root-t convergence of the learning process holds when the real part of the largest eigenvalue of the Jacobian of the ODE is less than $-\frac{1}{2}$.[13] When this condition on the eigenvalues is not met, no analytic results on the asymptotic distribution are known, since the importance of initial conditions fails to die out quickly enough. Marcet and Sargent (1992) suggest a numerical procedure to obtain an estimate of the rate of convergence. The starting point is the assumption that there is a δ for which

$$t^{\delta}(\theta_t - \theta) \xrightarrow{D} F, \tag{7}$$

where θ_t is the vector of parameters of the PLM, θ is its asymptotic limit, and F is some nondegenerate, well-defined, mean-zero distribution. Marcet and Sargent show that, for large t, a good approximation of the rate of convergence δ is given by the expression

$$\delta = \frac{1}{\log l} \log \sqrt{\frac{E(\theta_t - \theta)^2}{E(\theta_{tl} - \theta)^2}}. \tag{8}$$

Given t and l, the expectations can be approximated by simulating a large number of independent realizations of length t and $l \times t$, and calculating the mean square across realizations.

In this section, adaptive learning is introduced to analyze the implications of imperfect knowledge on policy outcomes. The question to be answered is how the interaction between learning and central bank preferences affects aggregate welfare. For the sake of clarity, the case where only the private sector learns is considered first; then the model is expanded to incorporate central bank learning.

[13] Root-t convergence means convergence at a rate of the same order as the root square of the sample size.

3.1 Private-Sector Learning

Suppose that private agents have nonrational expectations, which they try to correct through adaptive learning. Assume also that the policymaker does not explicitly take agents' learning into account and continues to set policy according to either (5) or (6). The evolution of output and inflation is therefore described by the system

$$y = \alpha(\pi - \hat{E}^P \pi) + \varepsilon$$
$$\pi = \begin{cases} \min\left[\hat{E}^{CB}\pi - \frac{z}{\phi} + \frac{2k}{\phi}, \bar{\pi}\right] \\ \rho \hat{E}^{CB}\pi - \frac{\rho}{\phi} z + \frac{\rho}{\phi} 2k, \end{cases} \tag{9}$$

where the inflation rate depends on the monetary authority's preferences. $\hat{E}^P\pi$ represents the current estimate of the inflation rate of the private sector, while $\hat{E}^{CB}\pi$ is the value of inflation expectations used in the central bank's control rule. It is assumed that private agents run regressions to set $\hat{E}^P\pi$, while the monetary authority, which observes $\hat{E}^P\pi$ before moving, has rational expectations and therefore sets $\hat{E}^{CB}\pi = \hat{E}^P\pi$.

At each period t, private agents have a PLM for inflation, which they use to make forecasts. This PLM takes the form $\hat{E}^P\pi_t = a_{Pt}$,[14] where

$$a_{Pt} = a_{Pt-1} + \frac{1}{t}(\pi_{t-1} - a_{Pt-1}). \tag{10}$$

The estimate a_{Pt} is updated over time using least squares; $\frac{1}{t}$ represents the gain parameter, which is a decreasing function of the sample size. Equation (10), in line with the literature, is in recursive form, uses data up to period $t-1$, and requires a starting value at time $t = 0$. The PLM has the same form as the RE solution for expected inflation: private agents estimate the parameter of the reduced form and set $\hat{E}^P\pi_t = a_{P_t}$.

[14]A time index is used only when strictly necessary—for instance, when tracking the evolution over time of least-squares learning.

Consider first the case in which the central bank acts as if it had a lexicographic ordering of preferences. The ALM turns out to be

$$\pi_t = \begin{cases} a_{Pt} - \dfrac{z_{t-1}}{\phi} + \dfrac{2k}{\phi} & \text{if } z_{t-1} \geq 0, \\ \bar{\pi} & \text{otherwise.} \end{cases} \tag{11}$$

The mapping between the PLM and the ALM generates the stochastic recursive algorithm

$$a_{Pt} = \begin{cases} a_{Pt-1} + \dfrac{1}{t}\left(-\dfrac{z_{t-1}}{\phi} + \dfrac{2k}{\phi}\right) & \text{if } z_{t-1} \geq 0, \\ a_{Pt-1} + \dfrac{1}{t}\left(\bar{\pi} - a_{Pt-1}\right) & \text{otherwise,} \end{cases} \tag{12}$$

which is approximated by the following ODE,

$$\frac{d}{d\tau} a_P = h(a_P) = \lim_{t\to\infty} E(\pi_{t-1} - a_P), \tag{13}$$

where

$$\begin{aligned} \lim_{t\to\infty} E(\pi_{t-1} - a_P) &= (\bar{\pi} - a_P) \int_{-2\mu}^{0} \left(\frac{1}{2\mu} + \frac{z}{4\mu^2}\right) dz \\ &\quad + \int_{0}^{2\mu} \left(-\frac{z}{\phi} + \frac{2k}{\phi}\right)\left(\frac{1}{2\mu} - \frac{z}{4\mu^2}\right) dz \\ &= \frac{1}{2}(\bar{\pi} - a_P) + \frac{k - \mu/3}{\phi} \\ &= \frac{1}{2}(\bar{\pi} - a_P) - \frac{\mu}{6\phi}. \end{aligned}$$

Notice that the fixed point of the ODE, namely, $a_P = h^{-1}(0) = \bar{\pi} - \frac{2k}{\phi} = \bar{\pi} - \frac{\mu}{3\phi}$, coincides with the unique REE for expected inflation. The theorems on the convergence of stochastic recursive algorithms can be applied so that convergence is governed by the stability of the associated ODE.[15] Since $\frac{d}{da_P} h(a_P) = -\frac{1}{2} < 0$, the ODE is

[15] Chapter 6 of Evans and Honkapohja (2001) studies the conditions under which convergence of stochastic recursive algorithms is ensured. In the case of

(globally) stable and hence adaptive learning asymptotically converges to the REE. Notice that the size of the eigenvalue of $h(a_P)$ depends on the share of the support of z corresponding to an active policy: in general, root-t convergence does not hold.

Consider now the situation in which the central bank has quadratic preferences. The optimal policy for the monetary authority is to set $\pi_t = \rho a_{Pt} - \frac{\rho}{\phi} z_t + \frac{\rho}{\phi} 2k$. Compared with the RE case, the central bank does not completely offset inflation expectations and the parameter k enters explicitly the control rule: both features disappear asymptotically, provided $a_P \rightarrow \pi^e = \frac{\bar{\alpha}}{\beta} k$.

If agents use recursive least squares, then expectations evolve according to the equations

$$\begin{aligned} a_{Pt} &= a_{Pt-1} + \frac{1}{t}(\pi_{t-1} - a_{Pt-1}) \\ &= a_{Pt-1} + \frac{1}{t}\left[(\rho - 1)a_{Pt-1} - \frac{\rho}{\phi} z_{t-1} + \frac{\rho}{\phi} 2k\right] \end{aligned} \tag{14}$$

and

$$h(a_P) = \lim_{t\rightarrow\infty} E\left[(\rho - 1)a_{Pt} - \frac{\rho}{\phi} z_{t-1} + \frac{\rho}{\phi} 2k\right] = (\rho - 1)a_P + \frac{\rho}{\phi} 2k. \tag{15}$$

Also in this case, the fixed point of the ODE, namely, $a_P = h^{-1}(0) = \frac{\rho}{(1-\rho)\phi} 2k = \frac{\bar{\alpha}}{\beta} k$, coincides with the unique REE and the system is (globally) stable. In fact, $\frac{d}{da_P} h(a_P) = \rho - 1 < 0$, since ρ is positive and smaller than 1. Whether root-t convergence holds depends on the size of β: the higher the weight the central bank attaches to the inflation objective, the faster agents learn. The explanation of this result is quite intuitive: the attempt to offset output shocks requires generating inflation surprises—i.e., moving the inflation rate away from expectations—so that every period agents will have to revise their estimate with values of π that may substantially differ from the unconditional mean. A similar result applies

interest, they amount to show that the process π, which is a linear function of z, is bounded and stationary and that the function driving the updating of the projection parameter, namely, $\pi_{t-1} - a_{Pt-1}$, is bounded and is twice countinuously differentiable (with respect to both π_{t-1} and a_{Pt-1}), with bounded second derivatives. Whether stability holds locally or globally depends on whether the regularity conditions hold on an open set around the equilibrium or for all admissible values of a_P. These regularity conditions, in the present case, are clearly met.

to the case of lexicographic preferences: by reducing the support of the signal corresponding to an active policy (i.e., to a policy that aims at avoiding excessive output fluctuations), expectations adjust more quickly to the long-term equilibrium. Notice that for reasonable parameterization of the model, the value of β must be large for $\frac{d}{da_P}h(a_P)$ to be less than $-\frac{1}{2}$, meaning that convergence holds only in the case of a highly inflation-averse central bank.

The previous result is interestingly similar to the one in Orphanides and Williams (2002), who use a dynamic model based on aggregate supply and demand equations. They find that, with imperfect knowledge, the ability of private agents to forecast inflation depends on the monetary policy in place, with forecast errors on average smaller when the central bank responds more aggressively to inflationary pressures. Significantly improved economic performances can be achieved by placing greater emphasis on controlling inflation: indeed, more-aggressive policies reduce the persistence of inflation and facilitate the formation of expectations, which in turn enhances economic stability and mitigates the influence of imperfect knowledge on the economy. The conclusion of the paper turns out to be quite similar to Rogoff's (1985) solution to the central bank's credibility problem under discretion: to improve welfare, the responsibility of the conduct of monetary policy must be delegated to a policymaker who is more inflation averse than society.

3.2 Private-Sector and Central Bank Learning

Consider now the case in which $\bar{\alpha}$ and σ^2_α (or, alternatively, ϕ) are not known to the policymaker. The central bank needs to estimate them since both parameters affect the policy rule through the degree of responsiveness to the signal z. As usual, it is assumed that they are gauged by means of least squares and that the estimate is updated every time new realizations of y and π are available. This form of bounded rationality corresponds to the case in which ε is never observed, so that $\bar{\alpha}$ and σ^2_α cannot be directly estimated on the basis of past realizations of the output shock.

To account for central bank learning, the previous model must be augmented with a new set of recursive equations, which are the same irrespective of the monetary authority's preferences, as learning involves parameters rather than variables so that the values to be estimated are not related to agents' behavior.

The system of recursive least-squares equations is now the following:

$$\begin{aligned}
a_{Pt} &= a_{Pt-1} + \frac{1}{t}(\pi_{t-1} - a_{Pt-1}) \\
\hat{\bar{\alpha}}_t &= \hat{\bar{\alpha}}_{t-1} + \frac{1}{t}S_{\pi,t}^{-1}(\pi_{t-1} - a_{Pt-1}) \\
&\quad \times \left[\left(y_{t-1} - \frac{z_{t-1}}{2}\right) - \hat{\bar{\alpha}}_{t-1}(\pi_{t-1} - a_{Pt-1})\right] \\
S_{y,t} &= S_{y,t-1} + \frac{1}{t}\left[\left(y_{t-1} - \frac{z_{t-1}}{2}\right)^2 - S_{y,t-1}\right] \\
S_{\pi,t} &= S_{\pi,t-1} + \frac{1}{t}[(\pi_{t-1} - a_{Pt-1})^2 - S_{\pi,t-1}]
\end{aligned} \tag{16}$$

or, more compactly,

$$\theta_t = \theta_{t-1} + \frac{1}{t}Q(t, \theta_{t-1}, X_t),$$

where $\theta_t = (a_{Pt}, \hat{\bar{\alpha}}_t, R_{y,t}, R_{\pi,t})'$, $R_{y,t-1} = S_{y,t}, R_{\pi,t-1} = S_{\pi,t}$, and $X_t = (1, \alpha_t, z_t, \varepsilon_t)'$. The first equation is the same as in the previous section and captures private-sector learning, while the others refer to the central bank's inference problem: $\hat{\bar{\alpha}}_t$ is an estimate of the mean value of the output-inflation trade-off; $R_{y,t}$ measures the sample variance of $y - \frac{z}{2}$, the policy-driven component of the output gap;[16] and $R_{\pi,t}$ is the second moment of the inflation surprise. As shown below, the central bank calculates the statistics $R_{y,t}$ and $R_{\pi,t}$ as an intermediate step in estimating the optimal response coefficient to the signal z in the policy rule.

While the recursion for $R_{\pi,t}$ is obvious, as it is simply the estimate of the variance of the inflation surprise, the other two equations require some explanations. To understand the recursion for $\hat{\bar{\alpha}}_t$, notice that the output equation can be rearranged as

$$y - \frac{z}{2} = \bar{\alpha}(\pi - a_P) + \left[\varepsilon - \frac{z}{2} + \widetilde{\alpha}(\pi - a_P)\right]. \tag{17}$$

[16] Since $E(\varepsilon|z) = \frac{z}{2}$, the difference $y - \frac{z}{2}$ represents (an unbiased estimate of) the share of the output gap that depends on the inflation surprise only.

The central bank observes the signal z and can efficiently estimate $\bar{\alpha}$ by regressing $y - \frac{z}{2}$ on the inflation surprise $(\pi - a_P)$. Using $y - \frac{z}{2}$ as the regressand allows for the elimination of the simultaneity bias: the inflation surprise, being a linear function of the signal z only, is uncorrelated with $\varepsilon - \frac{z}{2}$ (the residual of the regression of ε on z) and with $\widetilde{\alpha}$, by assumption orthogonal to all other shocks in the model.

The justification for the recursion for $R_{y,t}$ is somewhat more complicated. A biased estimator of $E(\alpha^2)$ can be obtained from the sample average of the squared (policy-driven component of the) output gap, scaled by the second moment of the inflation surprise,

$$\frac{E\left(y - \frac{z}{2}\right)^2}{E(\pi - a_P)^2} = \frac{E(\alpha^2)E(\pi - a_P)^2 + E\left(\varepsilon - \frac{z}{2}\right)^2}{E(\pi - a_P)^2}$$
$$= \bar{\alpha}^2 + \sigma_\alpha^2 + \frac{2\frac{\mu^2}{3}}{E(\pi - a_P)^2}.$$

The bias can be easily calculated, since it depends on $E(\pi - a_P)^2$ and on known parameters. The sample estimate of $\bar{\alpha}^2 + \sigma_\alpha^2$ is therefore obtained by using the expression $\psi_t \equiv \frac{R_{y,t} - 2\frac{\mu^2}{3}}{R_{\pi,t}}$.[17]

Whether the stochastic recursive algorithm converges depends on the associated ODE, i.e., on the Jacobian of the matrix $h(\theta) = \lim_{t\to\infty} EQ(t, \theta, X_t)$. In the case of lexicographic preferences, it can be shown that the ODE is

$$\begin{bmatrix} \frac{d}{d\tau} a_P \\ \frac{d}{d\tau} \hat{\bar{\alpha}} \\ \frac{d}{d\tau} R_y \\ \frac{d}{d\tau} R_\pi \end{bmatrix} = h(\theta) = \begin{bmatrix} \frac{1}{2}(\bar{\pi} - a_P) - \frac{\hat{\bar{\alpha}}}{2\frac{R_y - 2\mu^2/3}{R_\pi}} \frac{\mu}{6} \\ R_\pi^{-1} E(\pi - a_P)^2 (\bar{\alpha} - \hat{\bar{\alpha}}) \\ E\left(y - \frac{z}{2}\right)^2 - R_y \\ E(\pi - a_P)^2 - R_\pi \end{bmatrix}, \tag{18}$$

[17] One alternative could have been to regress $(y - \frac{z}{2})^2$ on $(\pi - a_P)^2$. The drawback of this approach is that the bias is a more convoluted function of the model parameters than it is in the case considered in the paper.

while for the standard quadratic case it is equal to

$$\begin{bmatrix} \frac{d}{d\tau}a_P \\ \frac{d}{d\tau}\hat{\hat{\alpha}} \\ \frac{d}{d\tau}R_y \\ \frac{d}{d\tau}R_\pi \end{bmatrix} = h(\theta) = \begin{bmatrix} -\frac{\beta}{\frac{R_y - 2\mu^2/3}{R_\pi} + \beta}a_P + \frac{\hat{\hat{\alpha}}}{\frac{R_y - 2\mu^2/3}{R_\pi} + \beta}\frac{\mu}{6} \\ R_\pi^{-1}E(\pi - a_P)^2(\bar{\alpha} - \hat{\hat{\alpha}}) \\ E\left(y - \frac{z}{2}\right)^2 - R_y \\ E(\pi - a_P)^2 - R_\pi \end{bmatrix}. \tag{19}$$

It is apparent that while the specific form of the loss function does not affect the inference problem of the central bank, it has a bearing on private-sector learning.

Both systems are recursive. $R_\pi \to E(\pi - a_P)^2$ from any starting point, which implies that $R_\pi^{-1}E(\pi - a_P)^2 \to I$, provided R_π is invertible along the path. The same happens for R_y. Hence, the stability of the differential equation for $\hat{\hat{\alpha}}$ may be assessed regardless of the remaining part of the system. Conditional on $\hat{\hat{\alpha}}$, R_y, and R_π approaching the true parameter values, convergence to the REE of private-sector expectations is determined on the basis of the eigenvalues of the ODE for a_P. It is noticeable that the probability limit of the latter does not depend on the information set of the central bank and is the same whether or not the monetary authority knows the full structure of the economy. Conditions for learnability of the REE under both lexicographic and quadratic preferences are stated in the next proposition.

PROPOSITION 3. *Assume that the economy is endowed with agents that rely on adaptive learning to form expectations; moreover, assume that the central bank has only incomplete information about the structure of the economy and uses recursive least squares (RLS) to estimate the unknown parameters. Then, the asymptotic behavior of the system is described by (18) and (19) and, regardless of whether the policymaker has quadratic or lexicographic preferences, the discretionary REE is unique and E-stable: the estimates* $(\hat{\hat{\alpha}}_t, \psi_t)$ *converge locally to* $(\bar{\alpha}, \bar{\alpha}^2 + \sigma_\alpha^2)$ *and the expectations of private agents tend in the limit to the RE values.*

Proof. See the appendix.

As in the case when only the private sector learns, the effect of preferences on the speed of convergence is not clear. For small values of β, a central bank setting policy so as to minimize a quadratic loss function seems to be less effective in driving the economy toward the REE, while the opposite is true when β is high. However, central banks' imperfect knowledge introduces an additional layer of interaction between monetary policy and economic outcomes, and model dynamics cannot be properly analyzed by focusing only on asymptotic distributions. In particular, when the learning process is disturbed by several sources of shocks, the ODE becomes an acceptable approximation to the stochastic recursive algorithm only for large values of t, and the asymptotic distribution is not of much help in understanding the properties of the system. The problem is even more serious in models where there are multiple equilibria, since in such cases, in early time periods, when estimates are based on very few degrees of freedom, large shocks can displace θ_t outside the domain of attraction of the ODE, and the system can therefore converge to any of the equilibrium points.[18] It follows then that when the agents' information set is severely constrained, both the asymptotic and the finite sample behavior of the system are relevant. Theoretical results are therefore no longer sufficient and it becomes necessary to rely on simulation experiments and numerical results.

4. Imperfect Knowledge and Policy Effectiveness

Model simulations are used to illustrate how learning affects the dynamic properties of inflation, inflation expectations, and output. First the performance of the forecasting rules is assessed; then the issue of the relative speed of convergence is considered; finally, the output-inflation variability trade-off under alternative monetary regimes is assessed. To account for the finding by Orphanides and

[18]When there is a unique equilibrium and the ODE is stable, it can be shown that $\theta_t \to \theta^*$ with probability 1 from any starting point. When there are multiple equilibria, however, such a strong result does not apply, unless one artificially constrains θ_t to an appropriate neighborhood of the locally stable equilibrium θ^*. In the earlier literature, local convergence was obtained by making an additional assumption about the algorithm, known as the *projection facility*. As a reference, see Evans and Honkapohja (2001, section 6.4).

Table 1. Baseline Calibrated Parameters

Parameter	Value
$\bar{\alpha}$	1.75
σ_α	0.5
μ	0.0175
k	0.0029

Williams (2002)—that policies that put too much weight on output stabilization can generate episodes in which the public's expectations of inflation become uncoupled from the policy objective—additional simulations are run mimicking the impact of a string of negative supply shocks on the economy. Finally, as a further check on how much the results depend on the chosen learning mechanism, the assumption of infinite memory is dropped and the case of perpetual learning is considered.[19]

Each experiment is based on 500 replications, and all simulations cover an interval of 2,000 periods. Subsamples of 500 observations are also considered in order to estimate the convergence speed. Initial conditions for the lagged variables in the RLS algorithm are randomly drawn from the distribution corresponding to the REE. Results reported in the tables are calculated excluding the first 150 periods so as to minimize the impact of initial observations, which could be too far away from the equilibrium solution. The model is calibrated according to the estimates in Ellison and Valla (2001); the selected parameter values are reported in table 1. Concerning β, the relative weight in the loss function of the inflation objective, three values are considered—namely, $\beta = \{.176, 1, 5.667\}$. Under lexicographic preferences, it is assumed that $\bar{\pi}$ is chosen so as to drive inflation expectations to 0.

Tables 2 and 3 report the simulation results: the first table describes the "plain" RLS learning rule (unconstrained estimation, or UE), while the second shows results for the case of constrained

[19]Perpetual learning is sometimes used as a synonym of constant-gain learning. A constant-gain algorithm is preferable when the agents believe that the economic environment is subject to frequent structural changes. In such cases, observations from the distant past are no longer informative and can instead become a source of distortion.

Table 2. Least-Squares Learning and the Volatility of Output and Inflation (Unconstrained Estimator—Decreasing-Gain Sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, the performance index, and the rate at which estimates of α_p, $\hat{\hat{\alpha}}$, and ψ converge to the REE. The performance index is the ratio between the social loss function under the delegated policymaker and that which would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to have infinite memory, implying a decreasing-gain sequence. Recursive least-squares estimates are unconstrained (UE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic Preferences		**Quadratic Preferences**	
	RE	**T = 2,000**	**RE**	**T = 2,000**
			$\beta = 0.176$	
Mean α_p	0.0000	−0.0005	0.0289	0.0283
Policy Rule Coefficient	0.2642	0.1807	0.2508	0.2308
Output Variability	0.0084	0.0088	0.0074	0.0076
Inflation Variability	0.0022	0.0023	0.0291	0.0298
Performance Index: $\beta^s = 0.176$	0.3494	0.3628	1.0000	1.0000
$\beta^s = 1.0$	0.8203	0.8453	9.8428	9.7605
$\beta^s = 5.667$	1.0925	1.1462	54.5380	55.0000
Convergence Speed	α_p: $\delta = 0.3351$		α_p: $\delta = 0.3570$	
	α: $\delta = 0.4951$		α: $\delta = 0.2435$	
	ψ: $\delta = 0.4609$		ψ: $\delta = 0.4884$	
			$\beta = 1.0$	
Mean α_p	0.0000	−0.0005	0.0051	0.0048
Policy Rule Coefficient	0.2642	0.1807	0.2029	0.1844
Output Variability	0.0084	0.0088	0.0076	0.0078
Inflation Variability	0.0022	0.0023	0.0059	0.0059
Performance Index: $\beta^s = 0.176$	0.3494	0.3628	0.3105	0.3168
$\beta^s = 1.0$	0.8203	0.8453	1.0000	1.0000
$\beta^s = 5.667$	1.0925	1.1462	2.8307	2.8287
Convergence Speed	α_p: $\delta = 0.3351$		α_p: $\delta = 0.2185$	
	α: $\delta = 0.4951$		α: $\delta = 0.2728$	
	ψ: $\delta = 0.4609$		ψ: $\delta = 0.4892$	
			$\beta = 5.667$	
Mean α_p	0.0000	−0.0005	0.0009	0.0008
Policy Rule Coefficient	0.2642	0.1807	0.0975	0.0775
Output Variability	0.0084	0.0088	0.0086	0.0088
Inflation Variability	0.0022	0.0023	0.0017	0.0016
Performance Index: $\beta^s = 0.176$	0.3494	0.3628	0.3627	0.3680
$\beta^s = 1.0$	0.8203	0.8453	0.8319	0.8340
$\beta^s = 5.667$	1.0925	1.1462	1.0000	1.0000
Convergence Speed	α_p: $\delta = 0.3351$		α_p: $\delta = 0.4080$	
	α: $\delta = 0.4951$		α: $\delta = 0.3642$	
	ψ: $\delta = 0.4609$		ψ: $\delta = 1.0000$	

Table 3. Least-Squares Learning and the Volatility of Output and Inflation (Constrained Estimator—Decreasing-Gain Sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, the performance index, and the rate at which estimates of α_p, $\hat{\hat{\alpha}}$, and ψ converge to the REE. The performance index is the ratio between the social loss function under the delegated policymaker and that which would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to have infinite memory, implying a decreasing-gain sequence. Recursive least-squares estimates are constrained to belong to a subset of the parameter space (CE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic Preferences		**Quadratic Preferences**	
	RE	**T = 2,000**	**RE**	**T = 2,000**
			$\beta = 0.176$	
Mean α_p	0.0000	−0.0003	0.0289	0.0267
Policy Rule Coefficient	0.2642	0.2365	0.2508	0.2411
Output Variability	0.0084	0.0087	0.0074	0.0075
Inflation Variability	0.0022	0.0025	0.0291	0.0271
Performance Index: $\beta^s = 0.176$	0.3494	0.4153	1.0000	1.0000
$\beta^s = 1.0$	0.8203	0.9092	9.8428	8.7535
$\beta^s = 5.667$	1.0925	0.7491	54.5380	28.4679
Convergence Speed	α_p: $\delta = 0.4193$		α_p: $\delta = 0.2972$	
	α: $\delta = 0.4670$		α: $\delta = 0.4557$	
	ψ: $\delta = 0.4765$		ψ: $\delta = 0.4892$	
			$\beta = 1.0$	
Mean α_p	0.0000	−0.0003	0.0051	0.0047
Policy Rule Coefficient	0.2642	0.2365	0.2029	0.1913
Output Variability	0.0084	0.0087	0.0076	0.0077
Inflation Variability	0.0022	0.0025	0.0059	0.0056
Performance Index: $\beta^s = 0.176$	0.3494	0.4153	0.3105	0.3475
$\beta^s = 1.0$	0.8203	0.9092	1.0000	1.0000
$\beta^s = 5.667$	1.0925	0.7491	2.8307	1.5986
Convergence Speed	α_p: $\delta = 0.4193$		α_p: $\delta = 0.1616$	
	α: $\delta = 0.4670$		α: $\delta = 0.4525$	
	ψ: $\delta = 0.4765$		ψ: $\delta = 0.4918$	
			$\beta = 5.667$	
Mean α_p	0.0000	−0.0003	0.0009	0.0014
Policy Rule Coefficient	0.2642	0.2365	0.0975	0.0805
Output Variability	0.0084	0.0087	0.0086	0.0105
Inflation Variability	0.0022	0.0025	0.0017	0.0026
Performance Index: $\beta^s = 0.176$	0.3494	0.4153	0.3627	0.6033
$\beta^s = 1.0$	0.8203	0.9092	0.8319	1.2990
$\beta^s = 5.667$	1.0925	0.7491	1.0000	1.0000
Convergence Speed	α_p: $\delta = 0.4193$		α_p: $\delta = 0.3331$	
	α: $\delta = 0.4670$		α: $\delta = 0.4544$	
	ψ: $\delta = 0.4765$		ψ: $\delta = 0.4870$	

estimation (CE). The latter amounts to impose some minimal restriction on the admissible regions of the estimates of $\bar{\alpha}$ and $\bar{\alpha}^2+\sigma_{\alpha}^2$: in the first case, it is assumed that only positive values of $\hat{\bar{\alpha}}_t$ are admissible, since surprises cannot have a negative impact on output; in the second, it is assumed that only values of ψ_t greater than $\hat{\bar{\alpha}}_t^2$ are sensible, since variances cannot be negative. The constraints act as substitutes for a *projection facility*:[20] though they cannot guarantee almost-sure convergence of the learning algorithm, they can in principle contribute to reducing the number of nonconvergent replications.

Expected inflation, the estimate of the policy rule coefficient, the standard deviation of output and inflation, and the performance index are reported in the tables, as is the convergence speed for a_P, $\hat{\bar{\alpha}}_t$, and ψ_t. The performance index is defined as

$$\frac{Ey^2_{LEX}+\beta^S E\pi^2_{LEX}}{Ey^2_{QUA(\beta^S)}+\beta^S E\pi^2_{QUA(\beta^S)}} \text{ or } \frac{Ey^2_{QUA(\beta)}+\beta^S E\pi^2_{QUA(\beta)}}{Ey^2_{QUA(\beta^S)}+\beta^S E\pi^2_{QUA(\beta^S)}},$$

depending on which strategy is evaluated. Ey^2_{LEX} is the second moment of output under strategy 1, and the other terms have a similar meaning; β is in general different from β^S, though the case $\beta=\beta^S$ is also considered. The index is equal to the ratio between the social loss function under the delegated policymaker and the loss that would be obtained by appointing a central banker with the same preferences as society's: when $PI<1$, it is efficient to delegate to a policymaker with a utility function that is different from society's. For ease of comparison, the values of the performance index are also shown for the REE.

Regardless of the central bank type, the estimates of expected inflation are precise but biased downward: the higher the values of

[20]Convergence of the learning process to the REE holds almost surely when there is a unique solution and the ODE is globally stable; in general, convergence with probability 1 is not guaranteed, since the ODE is not a reliable approximation of the stochastic recursive algorithm for small values of t. Almost-sure convergence holds only when the algorithm is supplemented with a *projection facility*, i.e., when θ_t is artificially constrained to remain in an appropriate neighborhood of θ (see section 6.4 and corollary 6.8 in Evans and Honkapohja 2001). The hypothesis of a projection facility, however, is often criticized because it cannot be easily justified on economic grounds and it is clearly inappropriate for decentralized markets.

β, the higher the accuracy of the estimate of a_p. Imposing constraints on the support of $\bar{\alpha}$ and $\bar{\alpha}^2 + \sigma_\alpha^2$ reduces the size of the bias in the case of lexicographic preferences but increases it if the policymaker has quadratic utility. One of the most striking findings of the simulation experiment concerns the minor effect on output and inflation variability of assuming imperfect central bank knowledge: the increase in output and inflation volatility, compared with the REE case, is surprisingly small—in most cases limited to a few percentage points. This is remarkable, since the estimation problem faced by the monetary authority is quite convoluted, requiring dealing with nonlinearities and computing higher-order moments. In the vast majority of cases, the increase in volatility of output and inflation remains well below 10 percent regardless of the preferences of the central bank, suggesting that the cost for the policymaker of having partial knowledge of the working of the economy is not disproportionately large. It is worth remembering, however, that the model lacks intrinsic dynamics, which explains why deviations from the REE tend to be short lived.

The analysis of the performance index provides a few interesting insights to the relative efficiency of the two strategies. The first finding is that, with decreasing gain, learning strategy 1 performs well regardless of the exact value of β^S: in the UE case, it is at least as good as strategy 2 unless the degree of inflation aversion is extremely high; in the CE case, it is uniformly better. The second finding is that the shorter the memory of the learning process, the poorer the performance of strategy 1, possibly because of the nonlinearity of the policy rule: when the policymaker has no discretion, the inflation surprises are not informative about the output-inflation trade-off and the recursive estimates become less accurate—in particular, when the size of the sample is small. The third finding is that the relative ranking of the two strategies is the same regardless of the way expectations are formed, which suggests that the learning process converges quite quickly to the REE. A closer inspection of the simulation evidence shows that for low values of β, a central bank endowed with lexicographic preferences is more effective in keeping inflation expectations under control, and it is also successful in stabilizing output fluctuations, even when bounds are not imposed on the RLS algorithm. The situation is reversed when β is high. A downward bias is evident in the estimate of the parameter measuring the response to the signal z, but it mostly disappears in the CE case; the imprecision in guessing the value of $\frac{1}{\phi}$ is responsible for some undesired fluctuations in output, while the excessive volatility of inflation is not attributable to the surprise component but rather to movements in private-sector inflation expectations, which under adaptive learning are not constant as in the REE. Since the bounds imposed on the RLS algorithm mimic the working of a projection facility, the rejection rate in the CE case turns out to be substantially lower.[21] Indeed, because of the complexity of the filtering problem facing the monetary authority, in a large number of replications, shocks displace the recursive algorithm outside of the domain of attraction of the ODE, and the estimate of the optimal response coefficient in the policy rule remains far off the true value. If the estimate of ϕ is very large at time t, the monetary authority has no incentive to respond aggressively to the signal z, and changes in y mostly reflect output shocks ε: the data become uninformative about the output-inflation trade-off and the estimate of ϕ gets larger and larger. Expectations become self-fulfilling and the economy gets stuck indefinitely on a suboptimal path, characterized by too passive a monetary policy, as if the policymaker's degree of inflation aversion were enormously higher than society's.

Except for high βs, strategy 2 underperforms strategy 1. It is, however, more effective in enhancing agents' learning process, as witnessed by the more precise estimate of the coefficient of the policy rule, which guarantees that the equilibrium under imperfect

[21]The rejection rate is calculated on the basis of the RLS estimate of the second moment of the output-inflation trade-off. Replications are considered diverging if the estimate of $\bar{\alpha}^2 + \sigma_\alpha^2$ is at least three times as large as the true value. The initial 150 observations are not used in the calculation. In the case of lexicographic preferences, absent constraints, the rejection rate turns out to be quite high (some 20 percent); it falls by a factor of 4 if the RLS algorithm is augmented with lower bounds. In the case of quadratic preferences, the number of diverging replications is on average much smaller and so is the gain obtained by imposing constraints on the learning process; for high βs, however, the rejection rate gets larger and becomes similar to the one observed under lexicographic preferences. The estimated number of diverging replications decreases sizably if less-restrictive criteria are used. Notice that divergence pertains to central bank learning and is defined in terms of the estimates of the policy parameters $\frac{1}{\phi}$ and $\frac{\rho}{\phi}$, which become very close to 0: neither the output gap nor the inflation rate are actually deviating boundlessly from equilibrium.

knowledge matches very closely the REE. There is no clear evidence that too-low or too-high inflation aversion can have a negative impact on the accuracy of the estimate of $\frac{\rho}{\phi}$. Concerning the speed of convergence, the two rules are more or less equivalent, though under CE, strategy 1 seems to be preferable; δ is very close to one half, so that convergence to a Gaussian distribution of both sequences $\{\theta_t\}^{LEX}$ and $\{\theta_t\}^{QUA}$ cannot be ruled out. As expected, strategy 2 reduces output variability more than strategy 1 when β is low, while the opposite result holds when inflation aversion is very high. The simulation results confirm that when learning substitutes rational expectations, a bias against activist policies arises. Hawkish policies are welfare enhancing because they offset the negative impact of two distortions: (i) the policymaker's desire to push output above the natural level and (ii) the uncoupling of expectations from policy objectives. While the latter distortion is present in the paper by Orphanides and Williams (2002), the former is not.

Tables 4 and 5 present evidence for the case of perpetual learning; the statistic for the speed of convergence is of course not shown since, under constant-gain learning, θ_t may at most converge to a probability distribution but not to a nonstochastic point. No meaningful differences are apparent compared with the previous case. Given the structure of the model, there is no benefit in discarding observations, so it is no surprise that in most cases RLS estimates are less accurate and policies are less successful in stabilizing both output and inflation; the deterioration in policy effectiveness seems to be relatively stronger for strategy 1.

The previous evidence suggests that a benevolent government may be better off appointing a central banker whose preferences are lexicographic if the degree of inflation aversion is not known with certainty or if it changes over time, since strategy 1 can get very close to maximize welfare for a large set of values of β^S. Strategy 1 can be implemented by giving the central bank a mandate that specifies an upper (and possibly a lower) bound on inflation and does not require that the government finds the perfect policymaker with the right preferences.[22]

[22]This feature also belongs to inflation-zone targeting; see Mishkin and Westelius (2006). It is easily seen that when the cost of overshooting (undershooting) the upper (lower) bound gets extremely large ($C \longrightarrow \infty$ in

Table 4. Least-Squares Learning and the Volatility of Output and Inflation (Unconstrained Estimator—Constant-Gain Sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, and the performance index. The performance index is the ratio between the social loss function under the delegated policymaker and that which would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant-gain sequence. Recursive least-squares estimates are unconstrained (UE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic Preferences		**Quadratic Preferences**	
	RE	**T = 2,000**	**RE**	**T = 2,000**
			$\beta = 0.176$	
Mean α_p	0.0000	−0.0005	0.0289	0.0288
Policy Rule Coefficient	0.2642	0.2084	0.2508	0.2001
Output Variability	0.0084	0.0090	0.0074	0.0076
Inflation Variability	0.0022	0.0026	0.0291	0.0291
Performance Index: $\beta^s = 0.176$	0.3494	0.3974	1.0000	1.0000
$\beta^s = 1.0$	0.8203	0.9326	9.8428	1.2276
$\beta^s = 5.667$	1.0925	1.2008	54.5380	48.8786
			$\beta = 1.0$	
Mean α_p	0.0000	−0.0005	0.0051	0.0051
Policy Rule Coefficient	0.2642	0.2084	0.2029	0.4747
Output Variability	0.0084	0.0090	0.0076	0.0077
Inflation Variability	0.0022	0.0026	0.0059	0.0059
Performance Index: $\beta^s = 0.176$	0.3494	0.3974	0.3105	0.3163
$\beta^s = 1.0$	0.8203	0.9326	1.0000	1.0000
$\beta^s = 5.667$	1.0925	1.2008	2.8307	2.5821
			$\beta = 5.667$	
Mean α_p	0.0000	−0.0005	0.0009	0.0009
Policy Rule Coefficient	0.2642	0.2084	0.0975	0.0503
Output Variability	0.0084	0.0090	0.0086	0.0090
Inflation Variability	0.0022	0.0026	0.0017	0.0018
Performance Index: $\beta^s = 0.176$	0.3494	0.3974	0.3627	0.3944
$\beta^s = 1.0$	0.8203	0.9326	0.8319	0.8952
$\beta^s = 5.667$	1.0925	1.2008	1.0000	1.0000

Mishkin's and Westelius's notation) and the policymaker does not care much about inflation variability inside the range ($\omega_\pi \longrightarrow 0$), strategy 1 and inflation-zone targeting become more and more alike.

Table 5. Least-Squares Learning and the Volatility of Output and Inflation (Constrained Estimator—Constant-Gain Sequence)

The table reports the estimated value (in 500 replications) of expected inflation, the coefficient of the optimal policy rule, the standard deviation of output and inflation, and the performance index. The performance index is the ratio between the social loss function under the delegated policymaker and that which would have been obtained by appointing a central banker with the same preferences as society's. Agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant-gain sequence. Recursive least-squares estimates are constrained to belong to a subset of the parameter space (CE case). In columns 1 and 3, the RE values for lexicographic and, respectively, quadratic preferences are shown; in columns 2 and 4, the same statistics are presented for the case when agents learn.

	Lexicographic Preferences		**Quadratic Preferences**	
	RE	**T = 2,000**	**RE**	**T = 2,000**
			$\beta = 0.176$	
Mean α_p	0.0000	−0.0003	0.0289	0.0288
Policy Rule Coefficient	0.2642	0.1917	0.2508	0.2070
Output Variability	0.0084	0.0101	0.0074	0.0075
Inflation Variability	0.0022	0.0023	0.0291	0.0290
Performance Index: $\beta^s = 0.176$	0.3494	0.5040	1.0000	1.0000
$\beta^s = 1.0$	0.8203	1.1546	9.8428	9.6551
$\beta^s = 5.667$	1.0925	1.4084	54.5380	51.4546
			$\beta = 1.0$	
Mean α_p	0.0000	−0.0003	0.0051	0.0051
Policy Rule Coefficient	0.2642	0.1917	0.2029	0.1603
Output Variability	0.0084	0.0101	0.0076	0.0077
Inflation Variability	0.0022	0.0023	0.0059	0.0058
Performance Index: $\beta^s = 0.176$	0.3494	0.5040	0.3105	0.3192
$\beta^s = 1.0$	0.8203	1.1546	1.0000	1.0000
$\beta^s = 5.667$	1.0925	1.4084	2.8307	2.6668
			$\beta = 5.667$	
Mean α_p	0.0000	−0.0003	0.0009	0.0009
Policy Rule Coefficient	0.2642	0.1917	0.0975	0.0588
Output Variability	0.0084	0.0101	0.0086	0.0089
Inflation Variability	0.0022	0.0023	0.0017	0.0016
Performance Index: $\beta^s = 0.176$	0.3494	0.5040	0.3627	0.3900
$\beta^s = 1.0$	0.8203	1.1546	0.8319	0.8799
$\beta^s = 5.667$	1.0925	1.4084	1.0000	1.0000

An additional set of simulations have been run to analyze the dynamic response of output and inflation to a sequence of unanticipated shocks. The experiment is designed supposing that the economy is perturbed by a string of negative output shocks, declining gradually in magnitude and vanishing after twelve periods. With rational expectations, the impact of the shocks is short lived and causes only a temporary fall in output and a rise in inflation, while under imperfect knowledge the response of the economy is prolonged and amplified by agents' learning. The objective of the experiment is to test whether the evidence reported by Orphanides and Williams (2002)—namely, that activist policies end up causing the perceived process for inflation to become uncoupled from the policymaker's objectives—is a general one and extends also to the theoretical framework adopted in this paper.

Tables 6 and 7 report the outcome of the experiment. Regardless of the central bank preferences, activist policies do not seem to pay off: the lower β is, the more volatile inflation and output are, especially the latter. Simulation evidence supports Rogoff's (1985) claim that it is welfare improving to appoint a central banker who places greater relative importance on the inflation objective than society does. No policy seems to be able to offset effectively the impact on activity of a sequence of negative output shocks: it still pays to be hawkish, but the attempts to reduce inflation volatility translate into output fluctuations that are much wider than under rational expectations. What worsens the performance of monetary policy is the uncoupling between expectations and policy targets: since expectations depend upon past values of inflation, they cannot be easily anchored unless the policymaker behaves as an inflation nutter. According to the performance index, strategy 1 is highly successful in promoting social welfare: in the UE case, except for very high values of β, it is more effective than strategy 2; in the CE case, it is uniformly better. A central banker endowed with lexicographic preferences ensures both robustness and effectiveness of monetary policy.

5. Conclusions

This paper focuses on the implications for the effectiveness of monetary policymaking of departing from the benchmark of rational expectations and applies a principal-agent approach to deal with the time-inconsistency problem that arises when the central bank cannot commit. It is assumed that society can delegate monetary policy to a central banker endowed with either quadratic or lexicographic preferences. Special attention is paid to the latter, which seems to describe more accurately the objectives of inflation-targeting central

Table 6. Dynamic Response to Contractionary Shocks (Decreasing-Gain Sequence)

The table reports a few statistics measuring how the equilibrium outcome under learning differs from the perfect knowledge—i.e., rational expectations—benchmark. Results are presented for both the "plain" RLS algorithm (UE) and the constrained version (CE); agents are assumed to have infinite memory, implying a decreasing-gain sequence. The first two columns refer to lexicographic preferences, while the next two refer to quadratic (dis)utility. To describe the dynamic response of output and inflation and to compare the outcomes under adaptive learning and rational expectations, three measures are computed: (1) the ratio of the volatility of the target variables under adaptive learning and under rational expectations, (2) the trough, and (3) the peak of the responses of output and inflation. In the last three lines of each section of the table, the value of the performance index is presented. All statistics are computed on the first 50 observations.

	Lexicographic Preferences		**Quadratic Preferences**	
	UE	**CE**	**UE**	**CE**
			$\beta = 0.176$	
$\sigma_y^{AL}/\sigma_y^{RE}$	6.6936	6.0532	1.9096	2.1572
Min y	−0.0177	−0.0149	−0.0132	−0.0196
Max y	0.0025	0.0025	0.0067	0.0067
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	2.8320	1.0090	1.0370	0.9868
Min π	−0.0028	−0.0012	0.0286	0.0261
Max π	0.0013	0.0013	0.0330	0.0330
Performance Index: $\beta^s = 0.176$	0.1634	0.1409	1.0000	1.0000
$\beta^s = 1.0$	0.5891	0.4687	17.9880	16.7520
$\beta^s = 5.667$	1.4124	0.8177	190.9600	155.5900
			$\beta = 1.0$	
$\sigma_y^{AL}/\sigma_y^{RE}$	6.6936	6.0532	2.6290	2.6582
Min y	−0.0177	−0.0149	−0.0138	−0.0144
Max y	0.0025	0.0025	0.0054	0.0054
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	2.8320	1.0090	1.0060	0.9812
Min π	−0.0028	−0.0012	0.0048	0.0048
Max π	0.0013	0.0013	0.0084	0.0084
Performance Index: $\beta^s = 0.176$	0.1634	0.1409	0.1460	0.1564
$\beta^s = 1.0$	0.5891	0.4687	1.0000	1.0000
$\beta^s = 5.667$	1.4124	0.8177	7.2675	6.2611
			$\beta = 5.667$	
$\sigma_y^{AL}/\sigma_y^{RE}$	6.6936	6.0532	5.7385	5.5524
Min y	−0.0177	−0.0149	−0.0162	−0.0149
Max y	0.0025	0.0025	0.0028	0.0028
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	2.8320	1.0090	0.6207	0.9806
Min π	−0.0028	−0.0012	−0.0003	0.0008
Max π	0.0013	0.0013	0.0025	0.0025
Performance Index: $\beta^s = 0.176$	0.1634	0.1409	0.1361	0.1365
$\beta^s = 1.0$	0.5891	0.4687	0.4769	0.4726
$\beta^s = 5.667$	1.4124	0.8177	1.0000	1.0000

Table 7. Dynamic Response to Contractionary Shocks (Constant-Gain Sequence)

The table reports a few statistics measuring how the equilibrium outcome under learning differs from the perfect knowledge—i.e., rational expectations—benchmark. Results are presented for both the "plain" RLS algorithm (UE) and the constrained version (CE); agents are assumed to use a finite number of observations in computing RLS estimates, implying a constant-gain sequence. The first two columns refer to lexicographic preferences, while the next two refer to quadratic (dis)utility. To describe the dynamic response of output and inflation and to compare the outcomes under adaptive learning and rational expectations, three measures are computed: (1) the ratio of the volatility of the target variables under adaptive learning and under rational expectations, (2) the trough, and (3) the peak of the responses of output and inflation. In the last three lines of each section of the table, the value of the performance index is presented. All statistics are computed on the first 50 observations.

	Lexicographic Preferences		**Quadratic Preferences**	
	UE	**CE**	**UE**	**CE**
			$\beta = 0.176$	
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	2.0835	1.6911
Min y	−0.0170	−0.0138	−0.0156	−0.0111
Max y	0.0025	0.0025	0.0067	0.0067
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.9863	0.9938
Min π	−0.0017	−0.0002	0.0283	0.0286
Max π	0.0025	0.0013	0.0330	0.0330
Performance Index: $\beta^s = 0.176$	0.1938	0.1217	1.0000	1.0000
$\beta^s = 1.0$	0.6838	0.4498	17.8460	19.4540
$\beta^s = 5.667$	1.2303	0.7748	168.8800	182.7800
			$\beta = 1.0$	
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	2.6372	2.3021
Min y	−0.0170	−0.0138	−0.0148	−0.0122
Max y	0.0025	0.0025	0.0054	0.0054
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.9321	0.9540
Min π	−0.0017	−0.0002	0.0046	0.0048
Max π	0.0025	0.0013	0.0084	0.0084
Performance Index: $\beta^s = 0.176$	0.1938	0.1217	0.1528	0.1281
$\beta^s = 1.0$	0.6838	0.4498	1.0000	1.0000
$\beta^s = 5.667$	1.2303	0.7748	6.2019	6.7111
			$\beta = 5.667$	
$\sigma_y^{AL}/\sigma_y^{RE}$	7.0633	5.5241	5.5729	5.3428
Min y	−0.0170	−0.0138	−0.0149	−0.0143
Max y	0.0025	0.0025	0.0028	0.0028
$\sigma_\pi^{AL}/\sigma_\pi^{RE}$	1.3607	0.7174	0.7987	0.8049
Min π	−0.0017	−0.0002	0.0005	0.0007
Max π	0.0025	0.0013	0.0025	0.0025
Performance Index: $\beta^s = 0.176$	0.1938	0.1217	0.1383	0.1308
$\beta^s = 1.0$	0.6838	0.4498	0.4987	0.4982
$\beta^s = 5.667$	1.2303	0.7748	1.0000	1.0000

banks. The main focus of the paper is on validating the claim that policies that are designed to be efficient under rational expectations can perform very poorly when knowledge is incomplete and agents learn adaptively.

The evidence shown in the paper confirms that, when agents do not possess complete knowledge of the structure of the economy and rely instead on an adaptive learning technology, a bias toward conservatism arises, suggesting that society is better off by appointing a policymaker whose degree of inflation aversion is higher than its own. The rationale for this finding is that agents' and policymakers' attempts to learn adaptively introduce inertia into the system and induce prolonged deviations of output and inflation from target, thereby raising the costs for the central bank of not responding promptly and forcefully to shocks. The paper also shows that the strategy that implements a lexicographic-preference ordering performs, on average, very well: it comes close to maximizing social welfare for a wide range of values of β^S and outperforms the strategies implementing quadratic preferences unless society is extremely inflation averse.

The findings of the paper closely resemble those of Orphanides and Williams (2002), which is surprising given the differences in the theoretical framework. First, the model adopted in this paper has no intrinsic dynamics, and the only source of persistence comes from the assumption that agents learn adaptively: the uncoupling between actual and perceived inflation is much harder to achieve with such a simple dynamic structure, though presumably, the lack of dynamics in the economy is compensated by the inertia induced by the attempts of the central bank to estimate the mean and variance of the output-inflation trade-off. Second, though only inflation expectations have a direct impact on the equilibrium outcome, output gap uncertainty affects the central bank's estimates of the moments of α and hence the policy setting: it is by no means obvious that a strategy that penalizes output variability might be conducive to higher welfare. The justification for the existence of a bias in favor of hawkish policies is to be found in the role of central bank learning: too-activist policies reduce the information content of the output gap and make estimates of the coefficients of the policy rule too volatile and unreliable.

Appendix

Proof of Proposition 3. Regardless of the preferences of the monetary authority, the recursive system representing the learning process is of the form $\theta_t = \theta_{t-1} + \frac{1}{t}Q(t, \theta_{t-1}, X_t)$, where $\theta_t = (a_{Pt}, \hat{\hat{\alpha}}_t, R_{y,t}, R_{\pi,t})'$ and $X_t = (1, \alpha_t, z_t, \varepsilon_t)$. To show the asymptotic stability of the REE under learning, one has to proceed as follows: first, it must be verified that there exists a nontrivial open domain containing the equilibrium point where the learning algorithm satisfies a few regularity conditions concerning the updating function $Q(t, \theta_{t-1}, X_t)$ and the stochastic process driving the state variables X_t; second, the local (or global) stability of the ODE associated with the stochastic recursive system must be established.

Consider first the case of lexicographic preferences. The system (16) has a unique equilibrium point θ^*, at which $a_P = \bar{\pi} - \frac{\mu/3}{\phi}$, $\hat{\hat{\alpha}} = \bar{\alpha}$, $R_\pi = \frac{2}{3\phi^2}(\frac{\mu^2}{3})$, and $R_y = (\bar{\alpha}^2 + \sigma_\alpha^2)[\frac{2}{3\phi^2}(\frac{\mu^2}{3})] + 2\frac{\mu^2}{3}$. It can easily be seen that θ^* is the REE. The stochastic process X_t is white noise, with finite absolute moments, so that regularity conditions (B.1) and (B.2) in Evans and Honkapohja (2001) are satisfied.[23] In addition, the gain sequence approaches zero asymptotically and is not summable. Finally, provided that R_π and R_y are nonzero along the learning path, $Q(t, \theta_{t-1}, X_t)$ satisfies a Lipschitz condition on a compact set containing the equilibrium point θ^*,[24] which ensures that regularity conditions (A.1)–(A.3) in Evans and Honkapohja (2001) are also met. Convergence of the learning process to the REE hinges therefore on the stability of the associated ODE (18). Notice that the system is recursive and the asymptotic behavior of the subsystem describing central bank learning can be assessed independently of the expectations formation mechanism of the private agents. Indeed, provided R_π and R_y are invertible along the convergence path, $R_y \to E(y - \frac{z}{2})^2$ and $R_\pi \to E(\pi - a_P)^2$ from any starting point; since $R_\pi^{-1}E(\pi - a_P)^2 \to I$, it is easily seen that $\hat{\hat{\alpha}}_t \to \bar{\alpha}$, since

[23] Chapter 6 in Evans and Honkapohja (2001) lists the regularity conditions required for the analysis of the asymptotic behavior of the stochastic recursive algorithm. Local stability is treated in section 6.2, while global convergence is analyzed in section 6.7.

[24] $Q(\theta_{t-1}, X_t)$ satisfies a Lipschitz condition if it is bounded and twice continuously differentiable, with bounded second derivatives.

the eigenvalue of the Jacobian of the corresponding differential equation has a negative real part. Conditional on $\hat{\hat{\alpha}}_t \to \bar{\alpha}$, convergence of private-sector inflation forecasts follows, since the associated ODE is stable.

A more formal proof of the convergence of the learning process to the REE requires proving that the Jacobian of the ODE, evaluated at the REE θ^*, has eigenvalues whose real parts are negative. In order to show that this is indeed the case, first notice that $R_\pi = E(\pi - a_P)^2$ at θ^*, which implies that R_π does not appear in the first three equations of the ODE evaluated at θ^*. A similar result holds for R_y. In the last two equations, the derivatives of R_π and R_y (and, accordingly, of $E(\pi - a_P)^2$ and $E(y - \frac{z}{2})^2$), though different from zero, cancel out, so that the Jacobian has the following upper-triangular, block-recursive structure:

$$Dh(\theta^*) = \begin{bmatrix} -\frac{1}{2} & -\frac{1}{\bar{\alpha}^2 + \sigma_\alpha^2}\frac{\mu}{12} & \frac{3}{2\mu}\frac{1}{\bar{\alpha}} & -\frac{3}{2\mu}\frac{\bar{\alpha}^2 + \sigma_\alpha^2}{\bar{\alpha}} \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}.$$

It is easily checked that its eigenvalues are $(-\frac{1}{2}, -1, -1, -1)$. They are all negative, and the system is therefore E-stable.

Consider now the case of quadratic preferences. The unique equilibrium point θ^* of the system (16) is now $a_P = \frac{\bar{\alpha}}{\beta}k$, $\hat{\hat{\alpha}} = \bar{\alpha}$, $R_\pi = E(\pi - a_P)^2 = 2(\frac{\rho}{\phi})^2(\frac{\mu^2}{3})$, and $R_y = E(y - \frac{z}{2})^2 = (\bar{\alpha}^2 + \sigma_\alpha^2)[2(\frac{\rho}{\phi})^2(\frac{\mu^2}{3})] + 2\frac{\mu^2}{3}$. The stochastic process X_t is the same as in the previous case, so that regularity conditions (B.1) and (B.2) in Evans and Honkapohja (2001) are satisfied. The same holds for the assumptions (A.1)–(A.3) on the gain sequence and the updating function $Q(t, \theta_{t-1}, X_t)$. The stability of the associated ODE (19) can be proved in the same way as for the system (18): provided that R_π and R_y are invertible along the convergence path, $R_\pi \to E(\pi - a_P)^2$ from any starting point; central bank estimates converge to the true parameter values $\bar{\alpha}$, since the eigenvalue of the Jacobian of the corresponding differential equation has a negative real part, and $a_P \to \frac{\bar{\alpha}}{\beta}k$, since the associated ODE is stable.

As in the previous case, the structure of the Jacobian justifies the sequential solution of the system. At θ^*, the derivative matrix of the ODE (19) is equal to

$$Dh(\theta^*) = \begin{bmatrix} -\frac{\beta}{\bar{\alpha}^2+\sigma_\alpha^2+\beta} & \frac{1}{\bar{\alpha}^2+\sigma_\alpha^2+\beta}\frac{\mu}{6} & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}.$$

The lower block for R_π and R_y can be solved first; then, triangularity of the upper block ensures that convergence for $\hat{\hat{\alpha}}_t$ does not depend on the asymptotic behavior of a_{Pt}. The eigenvalues of the Jacobian are $(-\frac{\beta}{\bar{\alpha}^2+\sigma_\alpha^2+\beta}, -1, -1, -1)$, and the system is therefore E-stable.

References

Benveniste, A., M. Metivier, and P. Priouret. 1990. *Adaptive Algorithms and Stochastic Approximations.* Berlin and Heidelberg: Springer-Verlag.

Bertocchi, G., and M. Spagat. 1993. "Learning, Experimentation and Monetary Policy." *Journal of Monetary Economics* 32 (1): 169–83.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Bray, M. M. 1982. "Learning, Estimation and Stability of Rational Expectations." *Journal of Economic Theory* 26 (2): 318–39.

Buiter, W. H. 2006. "How Robust Is the New Conventional Wisdom? The Surprising Fragility of the Theoretical Foundations of Inflation Targeting and Central Bank Independence." CEPR Discussion Paper No. 5772.

Driffill, J., and Z. Rotondi. 2003. "Monetary Policy and Lexicographic Preference Ordering." CEPR Discussion Paper No. 4247.

Ellison, M., and N. Valla. 2001. "Learning, Uncertainty and Central Bank Activism in an Economy with Strategic Interactions." *Journal of Monetary Economics* 48 (1): 153–71.

Evans, G. W., and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics.* Princeton, NJ: Princeton University Press.

———. 2002. "Expectations and the Stability Problem for Optimal Monetary Policy." Mimeo.

Jensen, H. 2003. "Explaining an Inflation Bias without Using the Word 'Surprise.'" Mimeo.

Locarno, A. 2006. "Imperfect Knowledge, Adaptive Learning and the Bias Against Monetary Policy." Temi di Discussione No. 590, Banca d'Italia.

Marcet, A., and T. J. Sargent. 1989. "Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models." *Journal of Economic Theory* 48 (2): 337–68.

———. 1992. "Speed of Convergence of Recursive Least Squares Learning with ARMA Perceptions." UPF Economics Working Paper No. 15.

Mishkin, F. S., and N. J. Westelius. 2006. "Inflation Band Targeting and Optimal Inflation Contracts." NBER Working Paper No. 12384.

al-Nowaihi, A., and L. Stracca. 2002. "Non-standard Central Bank Loss Functions, Skewed Risks and Certainty Equivalence." ECB Working Paper No. 129.

Orphanides, A., and D. W. Wilcox. 1996. "The Opportunistic Approach to Disinflation." FEDS Discussion Paper No. 24, Board of Governors of the Federal Reserve System.

Orphanides, A., and J. C. Williams. 2002. "Imperfect Knowledge, Inflation Expectations and Monetary Policy." Working Paper No. 27, Board of Governors of the Federal Reserve System.

Rogoff, K. 1985. "The Optimal Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

Simon, H. A. 1957. *Models of Man: Social and Rational.* New York: Wiley.

Svensson, L. E. O., and M. Woodford. 2000. "Indicator Variables for Optimal Policy." NBER Working Paper No. 7953.

Terlizzese, D. 1999. "A Note on Lexicographic Ordering and Monetary Policy." Mimeo, Bank of Italy.

Wieland, V. 2003. "Monetary Policy and Uncertainty about the Natural Unemployment Rate." CEPR Discussion Paper No. 3811.