CHAPTER 19

Inflation Expectations, Adaptive Learning and Optimal Monetary Policy☆

Vitor Gaspar,* Frank Smets, and David Vestin†**

*Banco de Portugal
**European Central Bank, CEPR and University of Groningen
†Sveriges Riksbank and European Central Bank

Contents

☆ The views expressed are the authors' own and do not necessarily reflect those of the Banco de Portugal or the European Central Bank or the Sveriges Riksbank. We thank Klaus Adam, Krisztina Molnar, and Mike Woodford for very useful comments.

ISSN 0169-7218, DOI: 10.1016/S0169-7218(11)03025-5

Abstract

This chapter investigates the implications of adaptive learning in the private sector's formation of inflation expectations for the conduct of monetary policy. We first review the literature that studies the implications of adaptive learning processes for macroeconomic dynamics under various monetary policy rules. We then analyze optimal monetary policy in the standard New Keynesian model, when the central bank minimizes an explicit loss function and has full information about the structure of the economy, including the precise mechanism generating private sector's expectations. The focus on optimal policy allows us to investigate how and to what extent a change in the assumption of how agents form their inflation expectations affects the principles of optimal monetary policy. It also provides a benchmark to evaluate simple policy rules. We find that departures from rational expectations increase the potential for instability in the economy, strengthening the importance of anchoring inflation expectations. We also find that the simple commitment rule under rational expectations is robust when expectations are formed in line with adaptive learning.
JEL classification: E52.

Keywords

Adaptive Learning
Optimal Monetary Policy
Policy Rules
Rational Expectations

1. INTRODUCTION

The importance of anchoring the private sector's medium-term inflation expectations for the effective conduct of monetary policy in the pursuit of price stability is widely acknowledged both in theory and in practice. For example, Trichet's (2009) claim in a recent speech that "it is absolutely essential to ensure that inflation expectations remain firmly anchored in line with price stability over the medium term" can be found in many other central bank communications. In a 2007 speech on the determinants of inflation and inflation expectations, Chairman Bernanke stated that "The extent to which inflation expectations are anchored has first-order implications for the performance of inflation and the economy more generally." The fear is that when medium-term inflation expectations become unanchored, they get ingrained in actual inflation or deflation, making it very costly to re-establish price stability. This is reflected in a letter by Chairman Volcker to William Poole: "I have one lesson indelible in my brain: don't let inflation get ingrained. Once that happens, there is too much agony in stopping the momentum."[1]

Following the seminal article of Muth (1961), it has become standard to assume rational or model-consistent expectations in modern macroeconomics. For example, in the context of a microfounded New Keynesian model Woodford (2003) systematically explored the

[1] Quoted in Orphanides (2005).

implications of rational expectations for the optimal conduct of monetary policy. However, rational expectations assume economic agents who are extremely knowledgeable (Evans & Honkapohja, 2001), an assumption that is too strong given the pervasive model uncertainty agents have to face. A reasonable alternative is to assume adaptive learning. In this case, agents have limited knowledge of the precise working of the economy, but as time goes by, and available data change, they update their beliefs and the associated forecasting rule. Adaptive learning may be seen as a minimal departure from rational expectations in an environment of pervasive structural change. It also better reflects reality where economists formulate and estimate econometric models to make forecasts and re-estimate those as new data becomes available. Moreover, some authors (see Section 2) have found that adaptive learning models are able to reproduce important features of empirically observed inflation expectations

This chapter analyzes the implications of private sector adaptive learning for the conduct of monetary policy. Using the baseline New Keynesian model with rational expectations, Woodford (2003) argued that monetary policy is first and foremost about the management of expectations, in particular inflation expectations.[2] In this chapter, we investigate whether this principle still applies when agents use adaptive learning instead of rational expectations. In Section 3 we introduce the standard New Keynesian model and some basic results and notation that will be used in the remainder of the chapter.

We provide a brief overview of the increasingly important line of research in monetary economics that studies the implications of adaptive learning processes for macroeconomic dynamics under various monetary policy rules (Section 4).[3] This literature typically investigates under which form of monetary policy rule the economy under learning converges to a rational expectations equilibrium. It was pioneered by Bullard and Mitra (2002), who applied the methodology of Evans and Honkapohja (2001) to monetary economics. This strand of the literature also discusses to what extent stability under learning can be used as a selection criterion for multiple rational expectations equilibria. Two key papers are Evans and Honkapohja (2003b, 2006), which analyze policy rules that are optimal under discretion or commitment in an environment with least-squares learning. They show that instabilities of "fundamental-based" optimal policy rules can be resolved by incorporating the observable expectations of the private agents in the policy rule.[4]

Furthermore, we analyze the optimal monetary policy response to shocks and the associated macroeconomic outcomes, when the central bank minimizes an explicit loss function and has full information about the structure of the economy (a standard assumption under rational expectations) including the precise mechanism generating private sector's

[2] A natural question to ask is whether the findings are robust in a model where the private sector is also learning about the output gap. The answer is yes. When the model incorporates a forward-looking IS equation (see Section 3), demand shocks are simply offset by interest rate changes. The same applies to fluctuations in output gap expectations. They do not create a trade-off between inflation and output gap stabilization and are, therefore, fully offset by changes in policy interest rates.

[3] See Evans and Honkapohja (2008a) for a recent summary of this literature.

[4] Most of this analysis is done within the framework of the standard two-equation New Keynesian model.

expectations.[5] This is in contrast to the literature summarized in Section 4, which only considers simple rules. The focus on optimal policy has two objectives. It allows investigating to what extent a relatively small change in the assumption of how agents form their inflation expectations affects the principles of optimal monetary policy. Second, it serves as a benchmark for the analysis of simple policy rules that would be optimal under rational expectations with and without central bank commitment, respectively. Here, the objective is to investigate how robust these policy rules are to changes in the way inflation expectations are formed.

As previously mentioned, the framework used in this chapter is the standard New Keynesian model. As shown by Clarida, Gali, and Gertler (1999) and Woodford (2003), in this model optimal policy under commitment leads to history dependence. In the model with rational expectations, credibility is a binary variable: the central bank either has the ability to commit to future policy actions and to influence expectations or not. With adaptive learning the private sector forms its expectations based on the past behavior of inflation. As a result, its outlook for inflation depends on the past actions of the central bank. Realizing this, following a cost-push shock the central bank will face an intertemporal trade-off between stabilizing output and anchoring future inflation expectations, in addition to the standard intratemporal trade-off between stabilizing current output versus current inflation.

Overall, in line with Orphanides and Williams (2005b) and Woodford (2010), we show that lessons for the conduct of monetary policy under model-consistent expectations are strengthened, when policy takes modest departures from rational expectations into account. The main intuition is that departures from rational expectations increase the potential for instability in the economy, strengthening the importance of managing (anchoring) inflation expectations. We also find that the simple commitment rule under rational expectations is robust when expectations are formed in line with adaptive learning. As a matter of fact, for the baseline calibration, macroeconomic outcomes, under the simple commitment rule, are surprisingly close to those under full optimal policy.

The rest of this chapter is organized as follows. Section 2 briefly reviews the evolution of measures of private sector inflation expectations in a number of industrial countries since the early 1990s. A number of papers have documented that with the establishment of monetary policy regimes focused on maintaining price stability, private sector medium-term inflation expectations have become much more anchored and do not respond very much to short-term inflation news. Section 3 then presents the basic New Keynesian model of inflation dynamics that will be used throughout most of the chapter and provides a characterization of the equilibrium under rational expectations as a benchmark for the analysis under adaptive learning. Section 4 gives an overview of the literature that studies the implications of adaptive learning processes for macroeconomic dynamics

[5] In doing so, we build on Svensson's (2003) distinction between "instrument rules" and "targeting rules." An instrument rule expresses the central bank's policy-controlled instrument, typically a short-term interest rate, as a function of observable variables in the central bank's information set. A targeting rule, in contrast, expresses it implicitly as the solution to a minimization problem of a loss function. Svensson stresses the importance of looking at optimal policy and targeting rules to understand modern central banking.

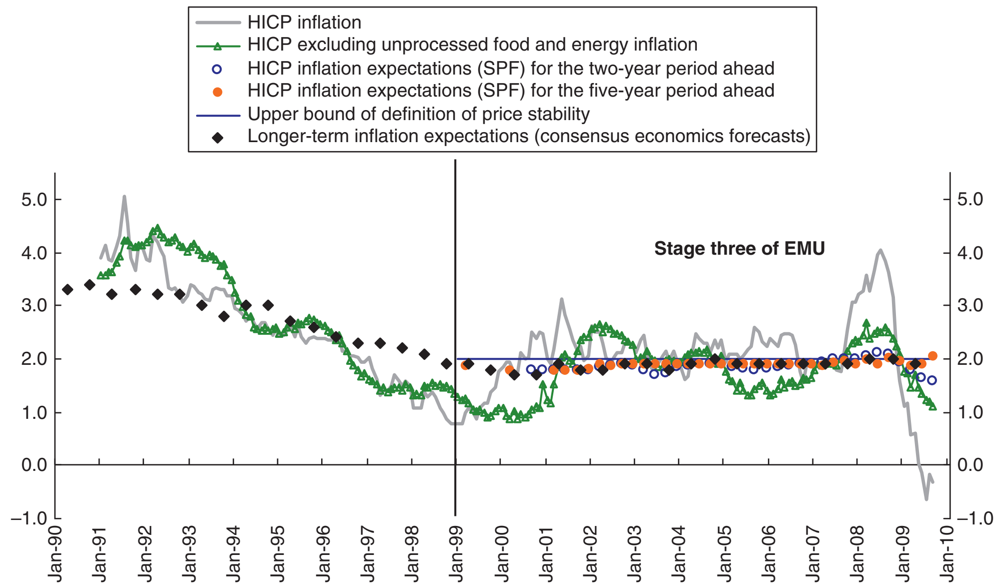

Figure 1 Inflation and inflation expectations in the Euro Area.

under various monetary policy rules. Section 5 instead discusses the implications of adaptive learning for optimal monetary policy in the baseline model. Finally, Section 6 contains a number of additional reflections related to alternative forms of expectations formation and Section 7 concludes the chapter.

2. RECENT DEVELOPMENTS IN PRIVATE-SECTOR INFLATION EXPECTATIONS

One of the striking developments in macroeconomic performance over the past two decades has been the firm anchoring of inflation expectations in many countries all over the world.[6] Figure 1 illustrates this for the Euro Area. Following the convergence of inflation and inflation expectations in the run-up to the establishment of Economic and Monetary Union (EMU) in January 1999, inflation expectations have been closely tied to the European Central Bank's (ECB) objective of keeping headline inflation close to, but below 2%. Moreover, in spite of the short-term volatility of headline inflation around this objective, both medium- and long-term inflation expectations have been very stable.

This is more generally true for many industrial and emerging countries. Figures 2 and 3 plot longer-term (5 to 10 years) and one-year ahead Consensus inflation forecasts for a

[6] This is also one of the conclusions of the chapter by Boivin, Kiley, and Mishkin (2010) in Volume 3A of the *Handbook of Monetary Economics* on changes in the monetary transmission mechanism. One of the changes highlighted is that inflation expectations respond less to changes in monetary policy.

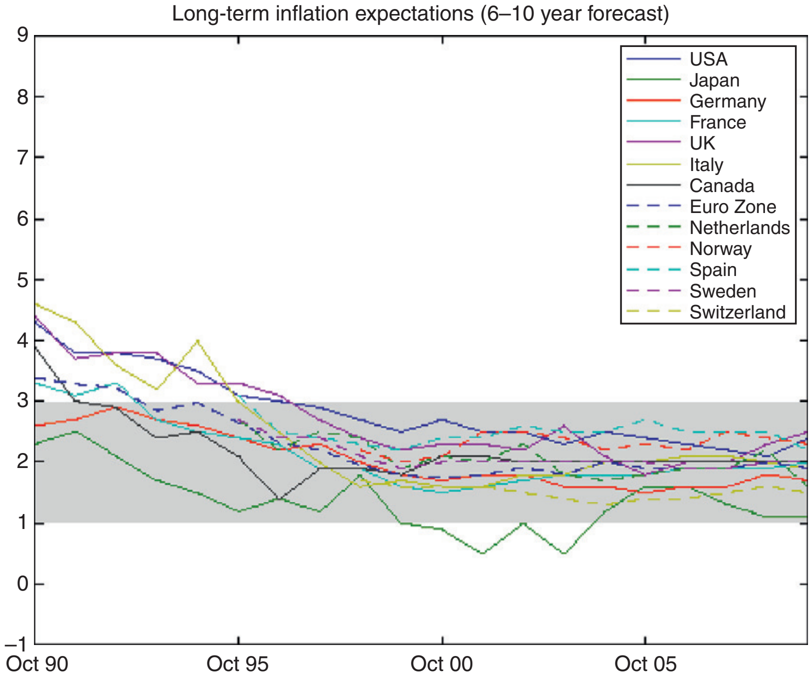

Figure 2 Long-term inflation expectations in selected OECD countries. (*Source:* Consensus Economics ©.)

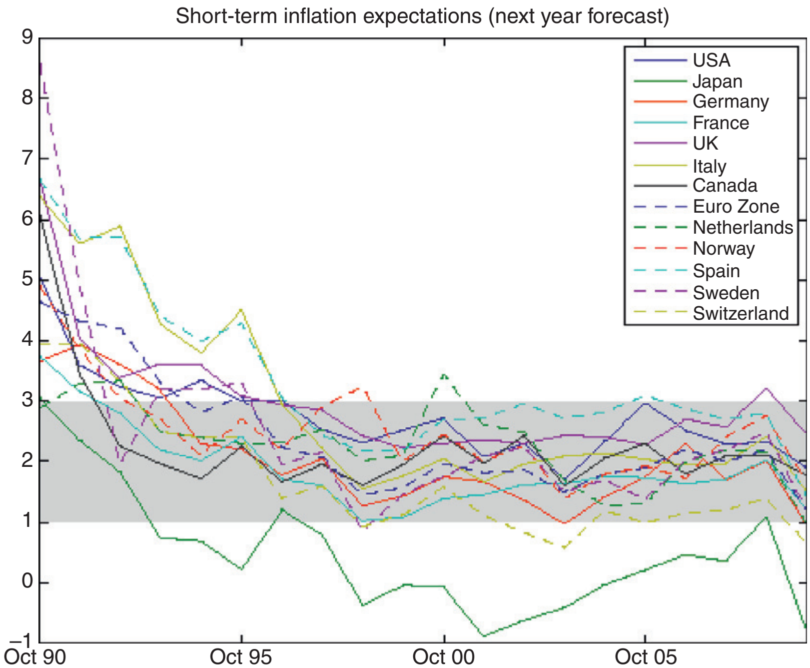

Figure 3 One-year ahead inflation expectations in selected OECD countries. (*Source:* Consensus Economics ©.)

number of Organization for Economic Cooperation and Development (OECD) countries. Figure 2 shows that, with the exception of Japan, longer term inflation expectations are very stable and are falling within a narrow band around 2%. One-year ahead forecasts are somewhat more variable but remain more or less within the 1 to 3% interval, even at the end of the sample following the most severe recession since the World War II in many countries.

A number of empirical studies have confirmed the visual impression that inflation expectations have become much more anchored as central banks have increasingly focused on achieving and maintaining price stability. Walsh (2009) and Blinder et al. (2008) summarized the evidence and concluded that inflation expectations have become well anchored in both inflation targeting (IT) and many non-IT countries. For example, Castelnuovo, Nicoletti-Altimari, and Palenzuela (2003) used survey data on long-term inflation expectations in 15 industrial countries since the early 1990s to find that in all countries, except Japan, long-term inflation expectations are well-anchored and generally increasingly so over the past two decades. Another interesting study (Beechey, Johannsen, & Levin, 2007) compares the recent evolution of long-run inflation expectations in the Euro Area and the United States, using evidence from financial markets and surveys of professional forecasters and shows that inflation expectations are well anchored in both economies, although surprises in macro-economic data releases appear to have a more significant effect on forward inflation compensation in the United States than in the Euro Area.

One way to explain the improved stability of inflation expectations over the past two decades is that private agents have adjusted their forecasting models to reflect the lower volatility and persistence in inflation. A number of studies have used least-squares learning models to explain expectations data (e.g. Branch, 2004; Branch & Evans, 2006; Orphanides & Williams, 2005a; Basdevant, 2005; and Pfajfar & Santoro, 2009). Moreover, Milani (2006, 2007, 2009) has incorporated least-squares learning into otherwise standard New Keynesian models for the United States and the Euro Area as a way to explain the changing persistence in the macroeconomic data (Murray, 2007; Slobodyan & Wouters, 2009).

3. A SIMPLE NEW KEYNESIAN MODEL OF INFLATION DYNAMICS UNDER RATIONAL EXPECTATIONS

Throughout most of the chapter, we use the following standard New Keynesian model of inflation dynamics, which under rational expectations can be derived from a consistent set of microeconomic assumptions, as extensively discussed in Woodford (2003):

$$\pi_t - \gamma\pi_{t-1} = \beta\left(E_t^*\pi_{t+1} - \gamma\pi_t\right) + \kappa x_t + u_t, \tag{1}$$

where E_t^* is an expectations operator, π_t is inflation, x_t is the output gap, and u_t is a cost-push shock (assumed i.i.d.). Furthermore, β is the discount rate; κ is a function of the underlying structural parameters including the degree of Calvo price stickiness, α; and γ captures the degree of intrinsic inflation persistence due to partial indexation of prices to past inflation.

In addition, we assume as a benchmark that the central bank uses the following loss function to guide its policy decisions:

$$L_t = (\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2. \tag{2}$$

Woodford (2003) showed that, under rational expectations and the assumed microeconomic assumptions, such a loss function can be derived as a quadratic approximation of the (negative of the) period social welfare function, where $\lambda = \kappa/\theta$ measures the relative weight on output gap stabilization and θ is the elasticity of substitution between the differentiated goods. We implicitly assume that the inflation target is zero. To keep the model simple, we first abstract from any explicit representation of the transmission mechanism of monetary policy and assume that the central bank controls the output gap directly.

As discussed in the introduction, we consider two assumptions regarding the formation of inflation expectations in Eq. (1): rational expectations and adaptive learning. Moreover, we assume that with the exception of the expectations operator, Eqs. (1) and (2) are invariant to these assumptions.[7] In this section, we first solve for optimal policy under rational expectations with and without commitment by the central bank. This will serve as a benchmark for the analysis of adaptive learning in the remainder of the chapter.

Defining $z_t = \pi_t - \gamma\pi_{t-1}$, Eqs. (1) and (2) can be rewritten as:

$$z_t = \beta E_t z_{t+1} + \kappa x_t + u_t \tag{1'}$$

$$L_t = z_t^2 + \lambda x_t^2. \tag{2'}$$

3.1 Optimal policy under discretion

If the central bank cannot commit to its future policy actions, it will be unable to influence expectations of future inflation. In this case, there are no endogenous state variables and, since the shocks are independent and identically distributed, the rational expectations solution (which coincides with the standard forward-looking model) must have the property $E_t z_{t+1} = 0$. Thus:

$$z_t = \kappa x_t + u_t \tag{1''}$$

Hence, the problem reduces to a static optimization problem. Substituting Eq. (1″) into Eq. (2′)) and minimizing the result with respect to the output gap implies the following policy rule:

$$x_t = -\frac{\kappa}{\kappa^2 + \lambda} u_t. \tag{3}$$

[7] It is clear that in general both the inflation equation (1) and the welfare function (2) may be different when adaptive learning rather than rational expectations are introduced at the micro level (Preston, 2005). In this paper, we follow the convention in the adaptive learning literature and assume that the structural relations (besides the expectations operator) remain identical when moving from rational expectations to adaptive learning.

Under the optimal discretionary policy, the output gap only responds to the current cost-push shock. In particular, following a positive cost-push shock to inflation, monetary policy is tightened and the output gap falls. The strength of the response depends on the slope of the New Keynesian Phillips curve, κ, and the weight on output gap stabilization in the loss function, λ.[8]

Using Eq. (3) to substitute for x_t in (1″):

$$z_t = \frac{\lambda}{\kappa^2 + \lambda} u_t. \tag{4}$$

Or, expressing the semi-difference of inflation directly as a function of the output gap:

$$z_t = -\frac{\lambda}{\kappa} x_t \tag{5}$$

This equation expresses the usual trade-off between inflation and output gap stability in the presence of cost-push shocks. In the standard forward-looking model (corresponding to $\gamma = 0$), there should be an appropriate balance between inflation and the output gap. The higher the λ, the higher inflation is in proportion to (the negative of) the output gap, because it is more costly to move the output gap. When κ increases, inflation falls relative to the output gap. When $\gamma > 0$, it is the balance between the quasi-difference of inflation and the output gap that matters. If last period inflation was high, current inflation will likely be high as well.

3.2 Optimal monetary policy under commitment

As shown earlier, under discretion optimal monetary policy only responds to the exogenous shock and there is no inertia in policy behavior. In contrast, as discussed extensively in Woodford (2003), if the central bank is able to credibly commit to future policy actions, optimal policy will feature a persistent "history dependent" response. In particular, Woodford (2003) showed that optimal policy will now be characterized by the following equation:

$$z_t = -\frac{\lambda}{\kappa}(x_t - x_{t-1}). \tag{6}$$

In this case, the expressions for the output gap and inflation can be written as:

$$x_t = \partial x_{t-1} - \frac{\kappa}{\lambda} u_t, \tag{7}$$

and

$$z_t = \frac{\lambda(1-\partial)}{\kappa} x_{t-1} + \partial u_t, \tag{8}$$

[8] The reaction function in Eq. (3)) contrasts with the one derived in Clarida et al. (1999). They assumed that the loss function is quadratic in inflation (instead of the quasi-difference of inflation, zt) and the output gap. They found that, in this case, lagged inflation appears in the expression for the reaction function, corresponding to optimal policy under discretion.

where $\partial = \left(\tau - \sqrt{\tau^2 - 4\beta}\right)/2\beta$ and $\tau = 1 + \beta + k^2/\lambda$ (see Clarida et al., 1999).

Comparing Eqs. (3) and (7), it is clear that under commitment optimal monetary policy is characterized by history dependence in spite of the fact that the shock is temporary. The intuitive reason for this is that under commitment perceptions of future policy actions help stabilize current inflation through their effect on expectations. By ensuring that, under rational expectations, a decline in inflation expectations is associated with a positive cost-push shock, optimal policy manages to reduce the impact of the shock and spread it over time.

3.3 Optimal instrument rules

Conditions (5) and (6) describe the optimal policy under discretion and commitment, respectively, in terms of the target variables in the central bank's loss function. To implement those policies, it is also useful to provide a reaction function for the policy-controlled interest rate. Consider the following "IS curve," which links the output gap to the short-term nominal interest rate:

$$x_t = -\varphi(i_t - E_t\pi_{t+1}) + E_t x_{t+1} + g_t \tag{9}$$

where i_t denotes the nominal short-term interest rate and g_t is a random demand shock. Such an equation can be derived from the household's consumption Euler equation, where φ is a function of the intertemporal elasticity of substitution.

Combining the IS curve (9), the price-setting equation (1) and the first-order optimality condition (6), respectively, treating the private expectations as given, the optimal expectations-based rule under commitment is given by[9]

$$i_t = \delta_L x_{t-1} + \delta_\pi E_t^* \pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t \tag{10}$$

where the reaction coefficients are functions of the underlying parameters.[10] The optimal rule under discretionary policy is identical except for the fact that $\delta_L = 0$. As before, in the commitment case, the dependence of the interest rate on lagged output reflects the advantage of the effects on expectations of commitment to a rule.

Alternatively, the optimal policy can also be characterized by a fundamentals-based rule that only depends on the exogenous shocks and the lagged output gap,

$$i_t = \psi_L x_{t-1} + \psi_g g_t + \psi_u u_t, \tag{11}$$

where the ψ parameters are again determined by the structural parameters and the objective function.

[9] In this derivation we have for simplicity assumed that there is no indexation.
[10] See Evans and Honkapohja (2008a).

Finally, it is also useful to consider a well-known alternative instrument rule, the so-called Taylor rule, which prescribes a response of the interest rate to current inflation and the output gap as follows:

$$i_t = \chi_\pi \pi_t + \chi_x x_t \tag{12}$$

This rule is not fully optimal in the New Keynesian model presented earlier, but has been shown to be quite robust in a variety of models.[11]

4. MONETARY POLICY RULES AND STABILITY UNDER ADAPTIVE LEARNING

In Section 3, agents in the economy were assumed to have rational (or model-consistent) expectations. As argued in the introduction, such an assumption is extreme given pervasive model uncertainty. Moreover, certain policy rules may be associated with indeterminacy of the rational expectations equilibrium, and therefore might be viewed as undesirable (Bernanke & Woodford, 1997). If the monetary authorities actually followed such a rule, the system might be unexpectedly volatile as agents are unable to coordinate on a particular equilibrium among the many that exist. In contrast, when equilibrium is determinate, it is normally assumed that the agents can coordinate on that equilibrium.

To address whether such coordination would arise, one needs to show the potential for agents to learn the equilibrium of the model being analyzed.[12] Following the seminal papers by Bullard and Mitra (2002) and Evans and Honkapohja (2003b, 2006), a growing literature has taken on this task by assuming that the agents in the model do not initially have rational expectations, and that they instead form forecasts by using recursive learning algorithms — such as recursive least-squares — based on the data produced by the economy.[13] This literature uses the methodology of Evans and Honkapohja (1998, 2001) to ask whether the agents in such a world can learn the fundamental or minimum state variable (MSV) equilibrium of the system under a range of monetary policy feedback rules. It uses the criterion of expectational stability (E-stability) to calculate whether or not rational expectations equilibria are stable under real-time recursive learning dynamics. Stability under learning is suggested as an equilibrium selection criterion and as a criterion for a "desirable" monetary policy. In this section, we review this literature on the performance of various monetary policy rules (like the ones presented in Section 3) when agents in the economy behave as econometricians; that is, as new data comes in agents re-estimate a reduced-form equation to form their expectations of inflation and the output gap. Most of the field has been

[11] See, for example, the Chapter 15 by Taylor and Williams (2010) in this volume.

[12] See Marcet and Sargent (1989) for an early analysis of convergence to rational expectations equilibria in models with learning.

[13] Evans and Honkapohja (2008a) labeled this assumption as the "principle of cognitive consistency."

covered recently, and with authority, by Evans and Honkapohja (2008a). We will follow their presentation to a large extent.

4.1 E-Stability In The New Keynesian Model

In this section, we briefly illustrate the concepts of determinacy and E-stability using the simple system given by the New Keynesian Phillips curve (1) without indexation, the forward-looking IS curve (9), and the simple Taylor rule (12). Defining the vectors

$$y_t = \begin{bmatrix} x_t \\ \pi_t \end{bmatrix} \quad \text{and} \quad v_t = \begin{bmatrix} g_t \\ u_t \end{bmatrix},$$

the reduced form of this system can be written as:

$$y_t = ME_t^* y_{t+1} + Pv_t \tag{13}$$

with

$$M = \frac{1}{1+\varphi(\chi_x + \kappa\chi_\pi)} \begin{bmatrix} 1 & \varphi(1-\beta\chi_\pi) \\ \kappa & \kappa\varphi + \beta(1+\varphi\chi_x) \end{bmatrix}$$

and

$$P = \frac{1}{1+\varphi(\chi_x + \kappa\chi_\pi)} \begin{bmatrix} 1 & -\varphi\chi_\pi \\ \kappa & 1+\varphi\chi_x \end{bmatrix}$$

4.1.1 Determinacy

First, consider the question whether the system under rational expectations (RE) possesses a unique stationary RE equilibrium (REE), in which case the model is said to be "determinate." If instead the model is "indeterminate," so that multiple stationary solutions exist, these will include "sunspot solutions"; that is, REE depending on extraneous random variables that influence the economy solely through the expectations of the agents.[14]

It is well known that in this case with two forward-looking variables, the condition for determinacy is that both eigenvalues of the matrix M lie inside the unit circle. It is easy to show that the resulting condition for determinacy for system (13) is given by [15]

$$\chi_\pi + \frac{1-\beta}{\kappa}\chi_x > 1. \tag{14}$$

In the determinate case, the unique stationary solution will be of the MSV form and only a function of the exogenous shocks:

[14] A number of papers have examined whether sunspot solutions are stable under learning. See, for example, Evans and Honkapohja (2003c) and Evans and McGough (2005b).

[15] See Bullard and Mitra (2002).

$$y_t = \bar{c} v_t. \tag{15}$$

Using the method of undetermined coefficients, it is straightforward to show that in the case of serially uncorrelated shocks $\bar{c}$ will be equal to P.

4.1.2 E-stability

Next, we consider system (13) under adaptive learning rather than rational expectations. In line with the MSV solution (15), suppose that agents believe that the solution is of the form:

$$y_t = a + c v_t \tag{16}$$

but that the 2×1 vector a and the 2×2 matrix c are not known but instead estimated by the private agents. In the terminology of Evans and Honkapohja (2001), Eq. (16)) is the perceived law of motion (PLM) of the agents. In (16), we assume that although in the RE equilibrium the intercept vector is zero, in practice the agents will need to estimate the intercept as well as the slope parameters. We also assume that the agents observe the fundamental shocks.

With this PLM and serially uncorrelated shocks, the agents expectations will be given by

$$E_t^* y_{t+1} = a.$$

Inserting these expectations in Eq. (13) and solving for y_t, we get the implied actual law of motion (ALM), which is given by

$$y_t = Ma + P v_t$$

We have now obtained a mapping from the PLM to the ALM given by

$$T(a, c) = (Ma, P) \tag{17}$$

and the REE solution $(0,P)$ is a fixed point of this map.

Under real-time learning, the sequence of events will be as follows. Private agents begin period t with estimates (a_t, c_t) of the PLM parameters computed on the basis of data through $t - 1$. Next, exogenous shocks v_t are realized and private agents form expectations using the PLM (16). The central bank then sets the interest rate, and y_t is generated according to Eq. (13). Finally, at the beginning of the period $t + 1$, agents add the new data point to update their parameter estimates to (a_{t+1}, c_{t+1}) using least-squares and the process continues.[16]

The E-stability principle of Evans and Honkapohja (2001) states that the REE will be stable under least-squares learning; that is, (a_t, c_t) will converge to the REE $(0,P)$,

[16] The learning algorithms, used in this literature, typically assume that the data sample steadily expands as time goes by. As the weight on each sample observation is the same, this implies that the gain from additional observations declines over time. This is contrast with the constant-gain least-squares learning considered in Section 5.

if the REE is locally asymptotically stable under the differential equation defined by the T-map (17):

$$\frac{d}{d\tau}(a, c) = T(a, c) - (a, c)$$

Using the results of Evans and Honkapohja (2001), we need the eigenvalues of M (given by Eq. 13) to have real parts less than 1 for E-stability. As shown by Bullard and Mitra (2002), this will be the case when condition (14) is satisfied. In the basic forward-looking New Keynesian model with a Taylor rule, the condition for stability under adaptive learning (E-stability) is implied by the condition for determinacy. This is, however, not a general result: sometimes E-stability will be a stricter requirement than determinacy and in other cases neither condition implies the other.[17]

Condition (14) is a variant of the Taylor principle, which states that the nominal interest rate should rise by more than current inflation in order to stabilize the economy.[18] In this case, the response could be slightly less than one, as long as there is a sufficiently large response to the output gap. Clarida, Gali and Gertler (2000) argued that the period of high and volatile inflation in the United States before Paul Volker became chairman of the Federal Reserve Board in 1979 can be explained by the Taylor principle being violated. Based on estimated reaction functions for the Federal Reserve, they show that the nominal federal funds rate reacted by less than one for one to expected inflation in the pre-Volker period. As a result, inflationary expectation shocks (sunspot shocks) can become self-fulfilling as they lead to a drop in the real rate and a boost in output and inflation. Using full-system maximum likelihood methods, Lubik and Schorfheide (2004) provide an empirical test of this proposition.

4.1.3 Extensions

Bullard and Mitra (2002) examined the stability of the REE under different variants of the Taylor rule (12) and found that the results are sensitive to whether the instrument rule depends on lagged, current, or future output and inflation. In all cases, the rules result in a stable equilibrium only if certain restrictions are imposed on the policy parameters. The role of learning is that it increases the set of restrictions required for stability and, thus, makes some instrument rules that were stable under RE unstable under learning. The results are of clear and immediate policy relevance. Specifically, they find that the Taylor principle, that is the interest rate, should be adjusted more than one-to-one in response to inflation ($\chi_\pi > 1$ in the previous equation) is crucial for learnability under the whole range of specifications they consider. More precisely

[17] Formal analysis of learning and E-stability for multivariate linear models is provided in Chapter 10 of Evans and Honkapohja (2001).

[18] See also Svensson and Woodford (2010) and Woodford (2003) for an extensive analysis of determinacy under various policy rules in the New Keynesian model.

they find that for $\chi_\pi > 1$ and χ_x sufficiently small, the outcome is determinate under rational expectations and stable under learning. However, for $\chi_\pi > 1$, but χ_x large the system can be indeterminate, but the MSV solution is stable under learning.[19] Overall, the results in Bullard and Mitra (2002) show that even when the system displays a unique and stable equilibrium under rational expectations, the parameters of the policy rule have to be chosen appropriately to ensure stability under learning. Bullard and Mitra (2002) also show that the nonobservability of current inflation and output gap can be circumvented by the use of "nowcasts" $E_t^* y_t$ instead of the actual data. The determinacy and E-stability conditions are not affected by this modification.

Evans and Honkapohja (2003b and 2006) considered the effect of learning on stability when the monetary authorities conduct policy according to the optimal policy rules under discretion and commitment presented in Eqs. (10) and (11). Both papers show that the "fundamental-based" optimal policy rules (11), which depend only on observable exogenous shocks and lagged variables, are consistently unstable under learning and therefore are less desirable as a guide for monetary policy. The authors show that the problem of instability under learning can instead be overcome when the policymaker is able to observe private sector expectations and incorporates them into the interest rate rule as in Eq. (10). The fundamental difference in these monetary policy rules is that they do not assume that private agents have RE but are designed to feed back on private expectations so that they generate convergence to the optimal REE. Also note that the expectations-based rule obeys a form of the Taylor principle since $\delta_\pi > 1$. One practical concern highlighted in Evans and Honkapohja (2008a) is that private sector expectations are not perfectly observed. However, if the measurement error in private sector expectations is small, the E-stability conditions discussed above remain valid. Overall, the importance of responding to private sector expectations for stability under learning is an important result, which will be echoed in Section 5. It provides a clear rationale for the central banks' practice of closely monitoring various measures of private sector inflation expectations and responding to deviations of those expectations from their desired inflation objective. As an aside, it is also curious to note that Evans and Honkapohja (2003d) showed that a Friedman k-percent money growth rule always results in determinacy and E-stability. However, it does not deliver an allocation close to optimal policy.

Following Bullard and Mitra (2002) and Evans and Honkapohja (2003b, 2005), a number of papers analyzed alternative monetary policy rules using different objective functions (Duffy & Xiao, 2007), in open economy settings (Bullard & Schaling, 2010; Bullard & Singh, 2008; Llosa & Tuesta, 2006), using extensions of the New Keynesian model with

[19] There can also exist E-stable sunspot equilibria as was shown by Honkapohja and Mitra (2004), Carlstrom and Fuerst (2004), and Evans and McGough (2005b).

a cost-channel (Kurozumi, 2006; Llosa & Tuesta, 2007), explicit capital accumulation (Duffy & Xiao, 2007; Kurozumi & Van Zandweghe, 2007; Pfajfar & Santoro, 2007) and sticky information (Branch, Carlson, Evans, and McGough, 2007, 2009) and under constant-gain rather than declining-gain least-squares learning (Evans & Honkapohja, 2008b).

4.2 Hyperinflation, deflation and learning

The central message from the literature discussed above to policymakers is very clear: When agents' knowledge is imperfect and they are trying to learn from observations, it is crucial that monetary policy prevents inflation expectations from becoming a source of instability in the economy. Most of the literature discussed earlier considers local stability in linear models. The literature also provides a number of examples of the importance of learning in a nonlinear context, where more than one inflation equilibrium is possible. Two examples stand out: hyperinflation and deflationary spirals.[20]

In an important recent paper, Marcet and Nicolini (2003) tried to explain recurrent hyperinflations experienced by some countries in the 1980s. They remarked that only a combination of orthodox (reduction of the deficit) and heterodox policies (an exchange rate rule) has been able to break the recurrence of hyperinflation. Marcet and Nicolini's model starts from a standard hyperinflation model with learning (as in Evans & Honkapohja, 2001). In this model, the high inflation equilibrium is not stable under adaptive learning. They extended the standard model to the case of a small open economy by considering one purchasing power parity equation and the possibility of following an exchange rate rule. The authors show that, with rational expectations, the model cannot account for the relevant empirical facts. However, under learning, the model simulations look very plausible and are able to account for all the empirical facts that Marcet and Nicolini (2003) document.

The global crisis brought the study of liquidity traps and deflationary spirals back to the center of the policy debate. Evans, Guse, and Honkapohja (2008) and Evans and Honkapohja (2009) considered these issues in the context of a New Keynesian model. They followed an insight by Benhabib, Schmitt-Grohe, and Uribe (2001) who showed that the consideration of the zero lower bound (ZLB) on nominal interest rates implies that the monetary policy rule must be nonlinear. It also implies the existence of a second lower inflation equilibrium (possibly with negative inflation rates).

Evans, Guse, and Honkapohja (2008) assumed a global Taylor rule and conventional Ricardian fiscal policy with exogenous public expenditures. They showed that the higher inflation equilibrium is locally stable under learning, but the lower inflation equilibrium is not. Around the latter equilibrium there is the possibility of deflationary

[20] See Sections 7 and 8 in Evans and Honkapohja (2008a) for a more extensive discussion of the papers discussed in this section.

spirals under learning. Interestingly, they showed that the possibility of deflationary spirals can be excluded by aggressive monetary and fiscal policy at some low threshold for the inflation rate.

5. OPTIMAL MONETARY POLICY UNDER ADAPTIVE LEARNING

In Section 4, we discussed the large literature that analyzes how various simple monetary policy rules affect the stability and determinacy of macroeconomic equilibria under adaptive learning. In this section, we analyze the optimal monetary policy response to shocks and the associated macroeconomic outcomes when the central bank minimizes an explicit loss function and has full information about the structure of the economy, including the precise mechanism generating private sector's expectations.[21] The basic model used is again the New Keynesian Phillips curve presented in Eq. (1). Different from most of the literature discussed in Section 4, we assume constant-gain (or perpetual) learning, which provides a more robust learning mechanism in the presence of structural change. Another difference is that for simplicity we will not explicitly take into account the IS curve, assuming instead that central banks can directly control the output gap. The next subsection analyzes the purely forward-looking New Keynesian Phillips curve and considers constant-gain learning about the inflation target as in Molnar and Santoro (2006). Section 5.2 analyzes the hybrid Phillips curve in Eq. (1) with indexation and considers constant-gain learning about the persistence of inflation as in Gaspar, Smets, and Vestin (2006a).

5.1 Adaptive learning about the inflation target in the forward-looking New Keynesian model

Following Molnar and Santoro (2006), this section first analyzes monetary policy in a purely forward-looking Phillips curve and when the private sector uses a simple adaptive learning mechanism about average inflation to form next period's inflation expectations. Under rational expectations and discretion, private sector expectations of next period's inflation will be zero (or equal to the inflation target in the non-log-linear version of the model). Under adaptive leaning, we hypothesize that the private sector calculates a weighted average of past inflation rates and expects that next period's inflation will be the same as such past average inflation. In particular,

$$E_t^* \pi_{t+1} \equiv a_t = a_{t-1} + \phi(\pi_{t-1} - a_{t-1}) \tag{18}$$

The advantage of analyzing this simple model is that the optimal policy problem is linear quadratic so it can be solved analytically. Private agents use a constant gain (similar to using a fixed sample length) to guard against structural changes. This example will

[21] See footnote 5.

also be useful to develop some of the intuition of the optimal policy response in the more complicated case in the next section.

In this case, the central bank problem can be stated as minimizing the expected present discounted value of the period loss function (2) with respect to $\{\pi_t, x_t, a_{t+1}\}$ subject to Eqs. (17) and (18). The first-order conditions are

$$2\pi_t - \lambda_{1,t} + \phi\lambda_{2,t} = 0 \tag{19}$$

$$2\lambda x_t + \kappa\lambda_{1,t} = 0 \tag{20}$$

$$E_t[\beta^2\lambda_{1,t+1} + \beta(1-\phi)\lambda_{2,t+1} - \lambda_{2,t}] = 0 \tag{21}$$

where $\lambda_{1,t}$ and $\lambda_{2,t}$ are the Lagrange multipliers associated with Eqs. (17) and (18), respectively.

Combining Eqs. (19) and (20) yields:

$$x_t = -\frac{\kappa}{\lambda}\left(\pi_t + \frac{\phi}{2}\lambda_{2,t}\right) \tag{22}$$

Assuming for simplicity that $\beta = 1$ and using Eq. (20), we can solve for $\lambda_{2,t}$ as a function of future output gaps by iterating Eq. (21) forward, yielding:

$$\lambda_{2,t} = -2\frac{\lambda}{\kappa}E_t\sum_{i=0}^{\infty}(1-\phi)^i x_{t+1+i} \tag{23}$$

Combining Eqs. (22) and (23) yields:

$$\pi_t = -\frac{\lambda}{\kappa}\left(x_t - \phi E_t\sum_{i=0}^{\infty}(1-\phi)^i x_{t+1+i}\right) \tag{24}$$

When there is no learning ($\phi = 0$), we are back to the discretionary RE solution as in Eq. (5). The central bank cannot affect inflation expectations and is left with managing the *intra*-temporal trade-off between stabilizing current output and current inflation in the presence of cost-push shocks. With learning ($\phi > 0$), there is also an *inter*-temporal trade-off given by the second term in Eq. (24). By allowing inflation to be affected to smooth current output, the central bank will affect future inflation expectations according to Eq. (18), which will create a trade-off between stabilizing inflation and the output gap in the future. The cost of this future trade-off is given by the second term in Eq. (24). By keeping current inflation closer to its target than suggested by the intratemporal trade-off, the central bank can stabilize future inflation expectations and improve the intertemporal trade-off. A first important result worth highlighting is that under optimal policy the central bank should act more aggressively toward inflation than what a rational expectations model under discretion would suggest. This is consistent with the work by Ferrero (2007), and Orphanides and Williams (2005a,b). A second important feature of optimal policy is that it is time consistent and qualitatively resembles the commitment solution under rational expectations because the optimal policy will be persistent and less willing to accommodate the effect of cost-push shocks on inflation.

This model can also be used to analyze the impact of a decreasing gain. Molnar and Santoro (2006) showed that following a structural break (e.g., a decrease in the inflation target), optimal policy should be more aggressive in containing inflation expectations because early on agents put more weight on the most recent inflation outcomes. Finally, Molnar and Santoro (2006) also investigated the robustness of their results when there is uncertainty about how the private sector forms its expectations and showed that the optimal policy under learning is robust to misperceptions about the expectation formation process.

5.2 Adaptive learning about inflation persistence in the hybrid New Keynesian model

In the more general New Keynesian model of Eq. (1), the equilibrium dynamics of inflation under rational expectations and discretionary optimal monetary policy will follow a first-order autoregressive process as shown in Eq. (4):

$$\pi_t = \rho\pi_{t-1} + \widetilde{u}_t \tag{4'}$$

In this case, we assume that under adaptive learning the private sector believes the inflation process is well approximated by such an AR(1) process. However, as the private agents do not know the underlying parameters, they estimate the equation recursively, using a "constant-gain" least-squares algorithm, implying perpetual learning. Thus, the agents estimate the following reduced-form equation for inflation,[22,23]

$$\pi_t = c_t\pi_{t-1} + \varepsilon_t \tag{25}$$

Agents are bounded rational because they do not take into account the fact that the parameter c varies over time. The c parameter captures the estimated, or perceived, inflation persistence. The following equations describe the recursive updating of the parameters estimated by the private sector.

$$c_t = c_{t-1} + \phi R_t^{-1}\pi_{t-1}(\pi_t - \pi_{t-1}c_{t-1}) \tag{26}$$

$$R_t = R_{t-1} + \phi(\pi_{t-1}^2 - R_{t-1}), \tag{27}$$

[22] In contrast to Section 5.1., we assume that the private sector knows the inflation target (equal to zero). While it would be useful to also analyze the case where the private sector learns about both the constant and the inflation target (as in Orphanides & Williams, 2005b), this is currently computationally infeasible.

[23] Alternatively, we could also assume that the private sector assumes that also lagged output gap affects inflation as in the case of commitment (Eq. 8). However, this would introduce three additional state variables in the nonlinear optimal control problem making it computationally infeasible to numerically solve the model. In this chapter, we therefore stick to the simpler univariate AR(1) case.

where ϕ is again the constant gain. Note that due to the learning dynamics the number of state variables is expanded to four: u_t, π_{t-1}, c_{t-1}, R_t, The last two variables are predetermined and known by the central bank at the time they set policy at time t.[24]

A further consideration regarding the updating process concerns the information the private sector uses when updating its estimates and forming its forecast for next period's inflation. We assume that agents use current inflation when they forecast future inflation, but not in updating the parameters. This implies that inflation expectations, in period t, for period t + 1 may be written simply as:

$$E_t^* \pi_{t+1} = c_{t-1}\pi_t \tag{28}$$

Generally, there is a double simultaneity problem in forward-looking models with learning. In Eq. (1), current inflation is determined, in part, by future expected inflation. However, according to Eq. (28), expected future inflation is not determined until current inflation is determined. Moreover, in the general case also the estimated parameter, c, will depend on current inflation. The literature has taken (at least) three approaches to this problem. The first is to lag the information set such that agents use only $t-1$ inflation when forecasting inflation at $t+1$, which was the assumption used in Gaspar and Smets (2004). A different and more common route is to look for the fixed point that reconciles both the forecast and actual inflation, but not to allow agents to update the coefficients using current information (i.e., just substitute Eq. 28 into Eq. 1 and solve for inflation). This keeps the deviation from the standard model as small as possible (also the rational expectations equilibrium changes if one lags the information set), while keeping the fixed-point problem relatively simple. At an intuitive level, it can also be justified by the assumption that it takes more time to re-estimate a forecasting model than to apply an existing model. Finally, a third approach is to also let the coefficients be updated with current information. This results in a more complicated fixed-point problem.

Substituting Eq. (28) into the New-Keynesian Phillips curve (1) we obtain:

$$\pi_t = \frac{1}{1+\beta(\gamma - c_{t-1})}(\gamma\pi_{t-1} + \kappa x_t + u_t). \tag{29}$$

5.2.1 Solution method for optimal monetary policy

We want to distinguish between the case where the central bank follows a simple rule (specifically the rules given in Eqs. 3 and 7) and fully optimal policy under the loss function (2). In the first case, the simple rule (Eqs. 3 or 7), the Phillips curve (1), and Eqs. (26)–(28) determine the dynamics of the system. Standard questions, in the

[24] Note that although agents are bounded rational, the forecast errors are close to serially independent and it would therefore be very difficult to detect systematic errors. In the benchmark case discussed later, the correlation between the forecast and actual inflation is 0.35. The serial correlation of the forecast error is 0.0036.

adaptive learning literature as discussed in Section 4, are then whether a given equilibrium is learnable and which policy rules lead to convergence to rational expectations equilibrium. By focusing on optimal policy, we aim at a different question. Suppose the central bank knows fully the structure of the model including that agents behave in line with adaptive learning. What is the optimal policy response? How will the economy behave? In this case, the central banker is well aware that policy actions influence expectations formation and inflation dynamics. To emphasize that we assume the central bank knows everything about the expectations' formation mechanism, Gaspar et al. (2006a) have labeled this extreme case "sophisticated" central banking. This implies solving the full dynamic optimization problem, where the parameters associated with the estimation process are also state variables.

Specifically, the central bank solves the following dynamic programming problem:

$$V(u_t,\pi_{t-1},c_{t-1},R_t) = \max_{x_t} \left\{ -\frac{(\pi_t - \gamma\pi_{t-1})^2 + \lambda x_t^2}{2} + \beta \mathrm{E}_t V(u_{t+1},\pi_t,c_t,R_{t+1}) \right\}, \quad (30)$$

subject to Eq. (29) and the recursive parameter updating Eqs. (26) and (27).[25]

The solution characterizes optimal policy as a function of the states and parameters in the model, which may be written simply as:

$$x_t = \psi(u_t,\pi_{t-1},c_{t-1},R_t). \quad (31)$$

As in this case the value function will not be linear-quadratic in the states, we employ the collocation-methods described in Judd (1998) and Miranda and Fackler (2002) to solve the model numerically. This amounts to approximating the value function with a combination of cubic splines and translates into a root-finding exercise. Further information on the numerical simulation procedure is outlined in Gaspar, Smets, and Vestin (2010).

5.2.2 Calibration of the baseline model

To study the dynamics of inflation under adaptive learning, we need to make specific assumptions about the key parameters in the model. In the simulations, we use the set of parameters shown in Table 1 as a benchmark. Coupled with additional assumptions on the intertemporal elasticity of substitution of consumption and the elasticity of labor supply, these structural parameters imply that $\kappa = 0.019$ and $\lambda = 0.002$.[26] γ is chosen such that there is some inflation persistence in the benchmark calibration. A value of 0.5 for γ is frequently found in empirically estimated New Keynesian Phillips curves (Smets, 2004; Gali & Gertler, 1999). $\theta = 10$ corresponds to a markup of about 10%.

[25] The value function is defined as $V(.) = \max_{\{x_j\}} \left\{ -\sum_j \beta^j \left[(\pi_j - \gamma\pi_{j-1})^2 + \lambda x_j^2 \right] s.t.(1),(26),(27)\ and\ (28) \right\}$, that is as maximizing the negative of the loss. It is important to bear this in mind when interpreting first-order conditions.

[26] Here we follow the discussion in Woodford (2003). See especially pages 187 and 214–215.

Table 1 Relevant Parameters for the Benchmark Case

β	γ	λ	θ	α	ϕ	κ	σ
0.99	0.5	0.002	10	0.66	0.02	0.019	0.004

Table 2 Summary of Macroeconomic Outcomes

	Rational expectations		Adaptive learning		
	Discretion	**Commitment**	**Discretion rule**	**Commitment rule**	**Optimal**
Corr(x_t, x_{t-1})	0	0.66	0	0.66	0.54
Corr(π_t, π_{t-1})	0.50	0.24	0.56	0.34	0.34
Var(x_t)	0.95	1	0.95	1	1.02
Var(π_t)	1.85	1	2.18	1.27	1.23
Var(π_t-$\gamma\pi_{\tau-1}$)	1.38	1	1.49	1.14	1.11
E[L_t]	1.29	1	1.37	1.11	1.09

Notes: Var(x_t), Var(π_t-$\gamma\pi_{\tau-1}$) and E[L_t] are measured as ratios relative to commitment.

1-α measures the proportion of firms allowed to change prices optimally each period. α is chosen such that the average duration of prices is three quarters, which is consistent with evidence from the United States. The constant gain, ϕ, is calibrated at 0.02. Orphanides and Williams (2005c) found that a value in the range of 0.01 to 0.04 is needed to match the resulting model-based inflation expectations with the Survey of Professional Forecasters. A value of 0.02 corresponds to an average sample length of about 25 years.[27] In the limiting case, when the gain approaches zero, the influence of policy on the estimated inflation persistence goes to zero and hence plays no role in the policy problem.

5.2.3 Macro-economic performance and persistence under optimal policy

In this section, we discuss the macroeconomic performance under adaptive learning. We compare the outcomes under rational and adaptive expectations for both optimal monetary policy and the simple policy rules given by Eqs. (3) and (7). Table 2 compares, for our benchmark calibration, five cases: two under rational expectations and three under adaptive learning. Under rational expectations we compare the discretionary and commitment policy; under adaptive learning we compare the optimal policy with the discretion and commitment rules (Eqs. 3 and 7, respectively) that would be optimal under rational expectations.

[27] See Orphanides & Williams (2005c). Similarly, Milani (2007) estimated the gain parameter to be 0.03 using a Bayesian estimation methodology.

It is instructive to first compare the well-known outcomes under commitment and discretion, under rational expectations. For such a case, we have shown in Section 3 (Clarida et al. 1999; Woodford, 2003) that commitment implies a long-lasting response to cost-push shocks persisting well after the shock has vanished from the economy. As previously stated, intuitively by generating expectations of a reduction in the price level in the face of a positive cost-push shock, optimal policy reduces the immediate impact of the shock and spreads it over time. With optimal policy under commitment, inflation expectations operate as automatic stabilizers in the face of cost-push shocks. Such intuition is clearly present in the results presented in Table 2. Clearly, the output gap is not persistent under the simple rule (under the assumption that cost-push shocks are i.i.d.). In contrast, under commitment the output gap becomes very persistent with autocorrelation of 0.66. The reverse is true for inflation. Inflation persistence, under discretion, is equal to the assumed intrinsic persistence parameter at 0.5. Under commitment it comes down to less than half of that at 0.24. The inflation variance is about 85% higher under discretion and the variance of the quasi-difference of inflation is about 37% higher. At the same time, output gap volatility is only about 5% lower. The reduction in output gap volatility illustrates the stabilization bias under optimal discretionary monetary policy. Overall, the loss is about 28% higher under discretion.

Following Orphanides and Williams (2002), it is also useful to compare the outcomes under rational expectations and adaptive learning for the case of the discretion and commitment rules (comparing the first and second columns with the third and fourth in Table 2). This comparison confirms the findings of Orphanides and Williams (2002). Clearly, the autocorrelation and the volatility of the output gap remain unchanged in both cases, under the simple rules the output gap only responds to the exogenous cost-push shock and (in the commitment case) its own lag. Nevertheless, under adaptive learning, the autocorrelation of inflation increases from 0.5 to about 0.56 in the discretion case and from 0.24 to 0.34 in the commitment case. As a result, the loss increases by about 8 percentage points under discretion and 11 percentage points under commitment. Intuitively, under adaptive learning, inflation expectations operate as an additional channel magnifying the immediate impact of cost-push shocks contributing to the persistence of their propagation in the economy. The increase in persistence and volatility are intertwined with dynamics induced by the learning process.

How does optimal monetary policy perform under adaptive learning (last column of Table 2)? As expected, it is able to improve macroeconomic performance relative to the simple linear rules that were optimal under rational expectations. Interestingly, it leads to similar outcomes as the commitment cases. Optimal policy induces considerable persistence in the output gap sharply reducing the persistence of inflation to about 0.34 (the same as under the commitment rule). As before, this is linked with a significant decline in inflation volatility relative to the discretionary outcomes. Inflation variance declines by 95 percentage points to only 23% more than in case of commitment under rational expectations. The variance of the quasi-difference of inflation also falls by about 38 percentage points. At the same time, the volatility of the output gap is slightly higher than under the discretion rules. On balance, the expected welfare loss falls significantly, by about 28 percentage points, when optimal policy replaces the simple discretionary rule.

Overall, it appears that optimal policy under adaptive learning brings the loss close to the one under commitment and rational expectations, as we can see from a comparison between the second and the last column in Table 2. Moreover, in both cases the output gap exhibits significant persistence and inflation is much less persistent than under the discretion rule. Nevertheless, it is still the case that even under optimal policy, adaptive learning makes inflation more persistent and the economy less stable than under rational expectations and the commitment rule. A second important conclusion to highlight is that the simple commitment rule, in which the output gap only responds to the cost-push shock and its own lag, does surprisingly well under adaptive learning. It delivers results very close to full optimal policy. The remarkable performance of the simple commitment rule under adaptive learning suggests that the ability of the central bank to adapt its response to cost-push shocks, depending on the state of the economy (e.g., lagged inflation and the perceived inflation persistence), is only of second-order importance relative to its ability to bring the perceived persistence of the inflation process down through a persistent response to cost-push shocks.

Figure 4 provides some additional detail concerning the distribution of the endogenous variables — the estimated persistence, output gap, inflation, quasi-difference of inflation, and the moment matrix — under optimal policy and the simple rules. First, panel (a) shows that the average of the estimated persistence parameter is significantly lower under the optimal policy and the simple commitment rule, and that the distribution is more concentrated around the mean. It is important to note that, under optimal policy, the perceived inflation parameter never goes close to one, contrary to what happens under the simple discretion rule. In fact, the combination of the simple discretion rule and private sector's perpetual learning at times gives rise to explosive dynamics, when perceived inflation persistence exceeds unity.[28] To portray the long-run distributions, we have excluded explosive paths by assuming (following Orphanides & Williams, 2005a) that when perceived inflation reaches unity the updating stops, until the updating pushes the estimated parameter downwards again. Naturally, this assumption leads to underestimating the risks of instability under the discretion rule. Gaspar, Smets, and Vestin (2006b) looked at the transition from an economy, regulated by the discretion rule, taking off on an explosive path to optimal policy leading gradually to the anchoring of inflation. Optimal monetary policy under adaptive learning succeeds in excluding such explosive dynamics.

[28] Similar results, for the case of a Taylor rule, are reported by Orphanides and Williams (2005a).

Second, panels (b), (c), and (d) confirm the results reported in Table 2. Under the optimal policy and the simple commitment rule, the distributions of inflation (panel c) and of the quasi-difference of inflation (panel d) become more concentrated. At the same time, the distributions of the output gap, in panel (d), are very similar confirming the result that the variances of the output gap under the two regimes are identical.

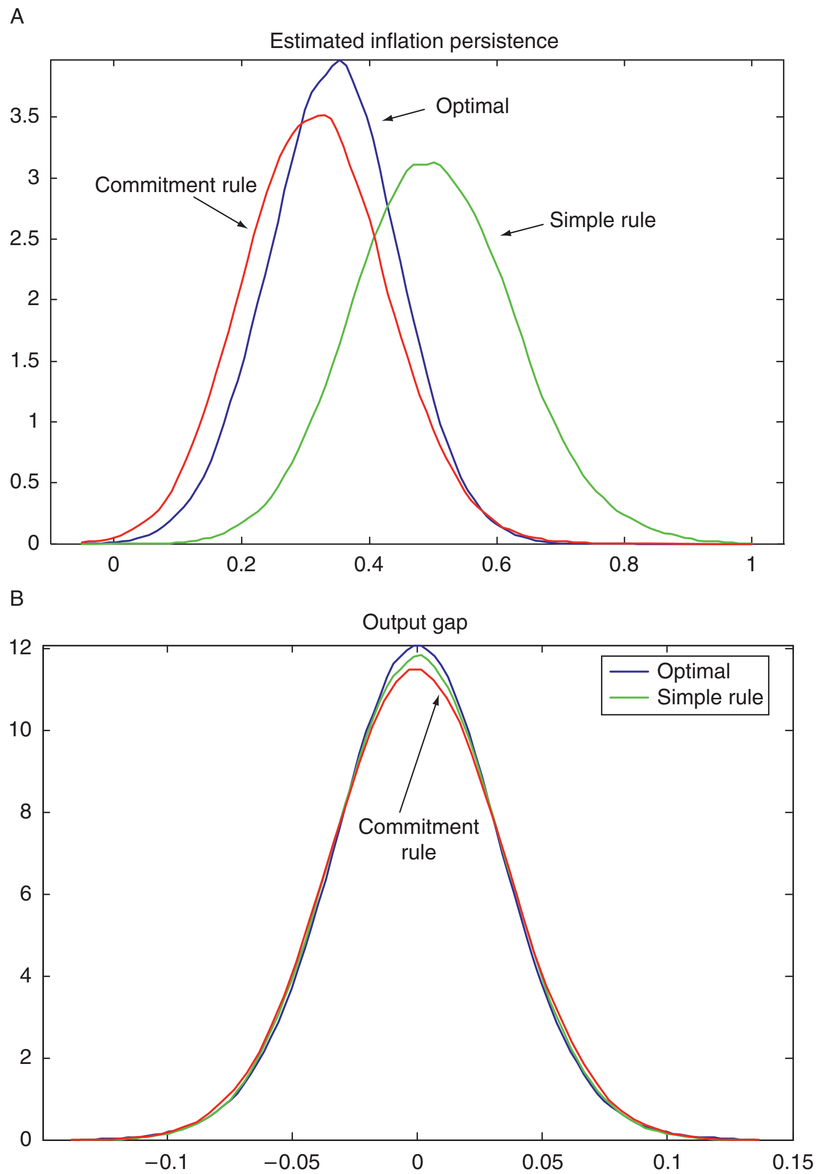

Figure 4 (Continued)

Finally, the distribution of the R matrix (panel e) also shifts to the left and becomes more concentrated under optimal policy, reflecting the fact that the variance of inflation falls relative to the simple discretion rule.

Overall, optimal monetary policy under adaptive learning shares some of the features of optimal monetary policy under commitment. To repeat, in both cases persistent responses to cost-push shocks induce a significant positive autocorrelation in the output gap, leading

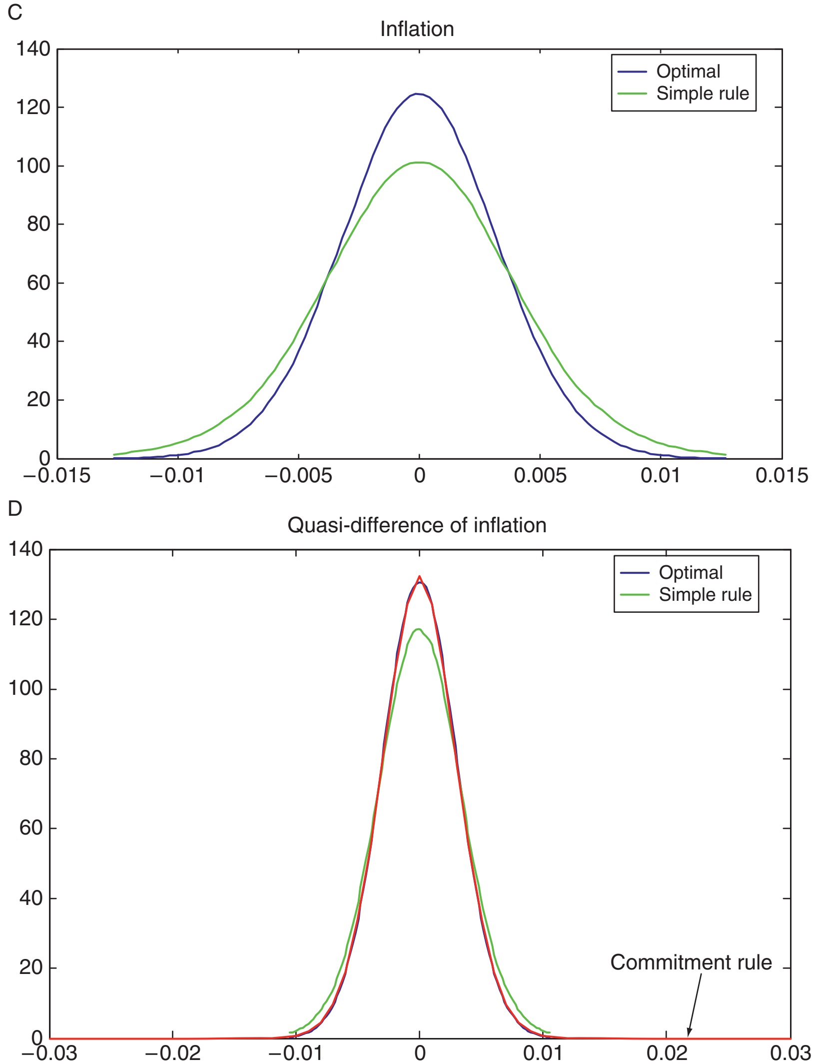

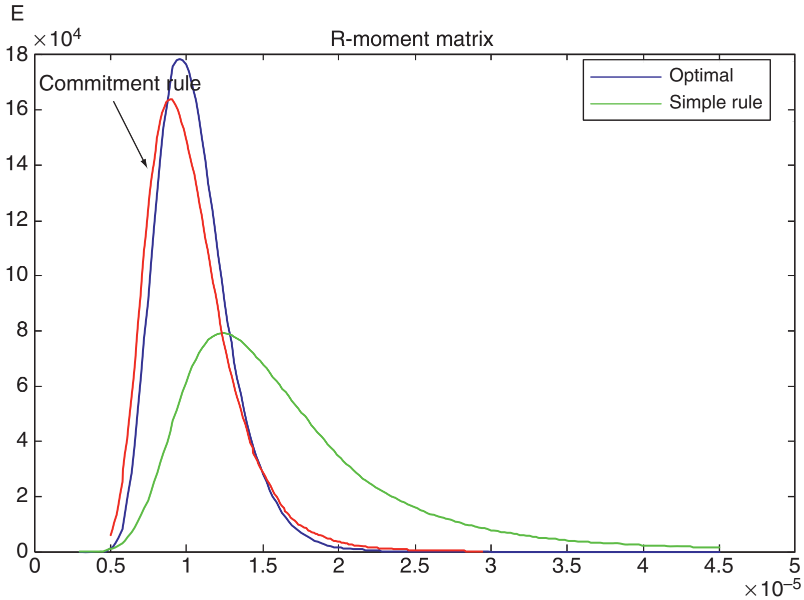

Figure 4 The distribution of (a) the estimated inflation persistence, (b) output gap, (c) inflation, (d) quasi-difference of inflation, and (e) the moment matrix.

to lower inflation persistence and volatility, through stable inflation expectations. Nevertheless, the details of the mechanism, leading to these outcomes, must be substantially different. As we have seen, under rational expectations commitment works through the impact of *future* policy actions on current outcomes. Under adaptive learning, the announcement of *future* policy moves is, by assumption, not relevant.

5.2.3 Optimal monetary policy under adaptive learning: How does it work?

Optimal monetary policy can be characterized by looking at the shape of the policy function and mean dynamic impulse responses following a cost-push shock. As discussed previously, optimal policy may be characterized as a function of the four state variables in the model: $(u_t, \pi_{t-1}, c_{t-1}, R_t)$. Gaspar et al. (2010) showed that Eq. (31) can implicitly be written as:

$$x_t = -\frac{\kappa\delta_t}{\kappa^2\delta_t + \lambda\chi_t^2}u_t + \frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1}E_tV_c}{\kappa^2\delta_t + \lambda\chi_t^2}\pi_{t-1} + \beta\frac{\kappa\chi_t}{\kappa^2\delta_t + \lambda\chi_t^2}E_tV_\pi \quad (32)$$

where $\delta_t = 1 - 2\beta\phi E_t V_R$, $\chi_t = 1 + \beta(\gamma - c_{t-1})$ and V_c, V_π and V_R denote the partial derivatives of the value function with respect to the variables indicated in the subscript. When interpreting Eq. (32) there are two important points to bear in mind. First, the partial derivatives V_c, V_π and V_R depend on the vector of states $(u_{t+1}, \pi_t, c_t, R_{t+1})$. The last three states, in turn, depend on the history of shocks and policy responses. Second, the value function is defined in terms of a maximization problem. In such a case, a positive partial derivative means that an increase in the state contributes favorably to our criterion. Or, more explicitly, that it contributes to a *reduction* in the loss.

To discuss some of the intuition behind the optimal policy reaction function, it is useful to consider a number of special cases. In particular, in the discussion that follows, we assume that $E_t V_R$ is zero, so that the expected marginal impact of changes in the moment matrix on the value function is zero. Such assumption provides a reasonable starting point for the discussion for reasons made clear in Gaspar et al. (2010). If $E_t V_R$ is zero, then $\delta_t = 1$, which makes Eq. (32) much simpler.

5.2.4 The intra-temporal trade-off ($\pi_{t-1} = 0$)

If lagged inflation is equal to zero, $\pi_{t-1} = 0$, the optimal monetary policy reaction (24) can be reduced to a simple response to the current cost-push shock:

$$\mathrm{x_t} = -\frac{\kappa}{\kappa^2 + \lambda \chi_{\mathrm{t}}^2} \mathrm{u_t}. \tag{33}$$

This is the case because clearly the second term on the right-hand side of Eq. (24) is zero; moreover, it can be shown that for $\pi_{t-1} = 0$, $E_t V_\pi$ is zero.

If, in addition, $c_{t-1} = \gamma$ and as a result $\chi_t^2 = \chi_t = 1$, Eq. (33) reduces to the simple rule derived under rational expectations and discretion given by Eq. (3). In other words, when lagged inflation is zero and the estimated inflation persistence is equal to the degree of intrinsic persistence, the immediate optimal monetary policy response to a shock under adaptive learning coincides with the optimal response under discretion and rational expectations.[29] The reason for this finding is quite intuitive. From Eq. (26), it is clear that, when lagged inflation is zero, the estimated persistence parameter is not going to change irrespective of current policy actions. As a result, no benefit can possibly materialize from trying to affect the perceived persistence parameter. The same intuition holds true to explain why when the constant gain parameter is zero (ϕ=0) the solution under fully optimal policy coincides with Eq. (3), meaning that the simple discretion rule would lead to full optimal policy. In this case, only the intra-temporal trade-off between output and inflation stabilization plays a role. However, different from the discretionary policy under rational expectations, the optimal response under adaptive learning will generally depend on the perceived degree of inflation persistence. For example, when the estimated persistence is lower than the degree of intrinsic persistence, $\gamma > c_{\mathrm{t}-1}$, the immediate response to a cost-push shock will be less, $\frac{\kappa}{\kappa^2+\lambda\chi_{\mathrm{t}}^2} < \frac{\kappa}{\kappa^2+\lambda}$, than under the simple discretion rule. The reason is again intuitive. As shown in Eq. (29), the smaller the degree of perceived inflation

[29] However, it is clear from Figure 5 that the policy response under optimal policy will persist contrary to the simple discretion rule.

persistence, the smaller the impact of a given cost-push shock on inflation, all other things constant. As a result, when balancing inflation and output gap stabilization, it is optimal for the central bank to mute its immediate response to the cost-push shock. This clearly illustrates the first-order benefits of anchoring inflation expectations. Conversely, when perceived inflation persistence is relatively high, the response of optimal policy to cost-push shocks becomes stronger on impact than under the simple rule.

In Figure 5, we illustrate this response by showing the mean dynamic response of the output gap, inflation and estimated persistence to a one-standard deviation (positive) cost-push shock, taking lagged inflation to be initially zero, for different initial

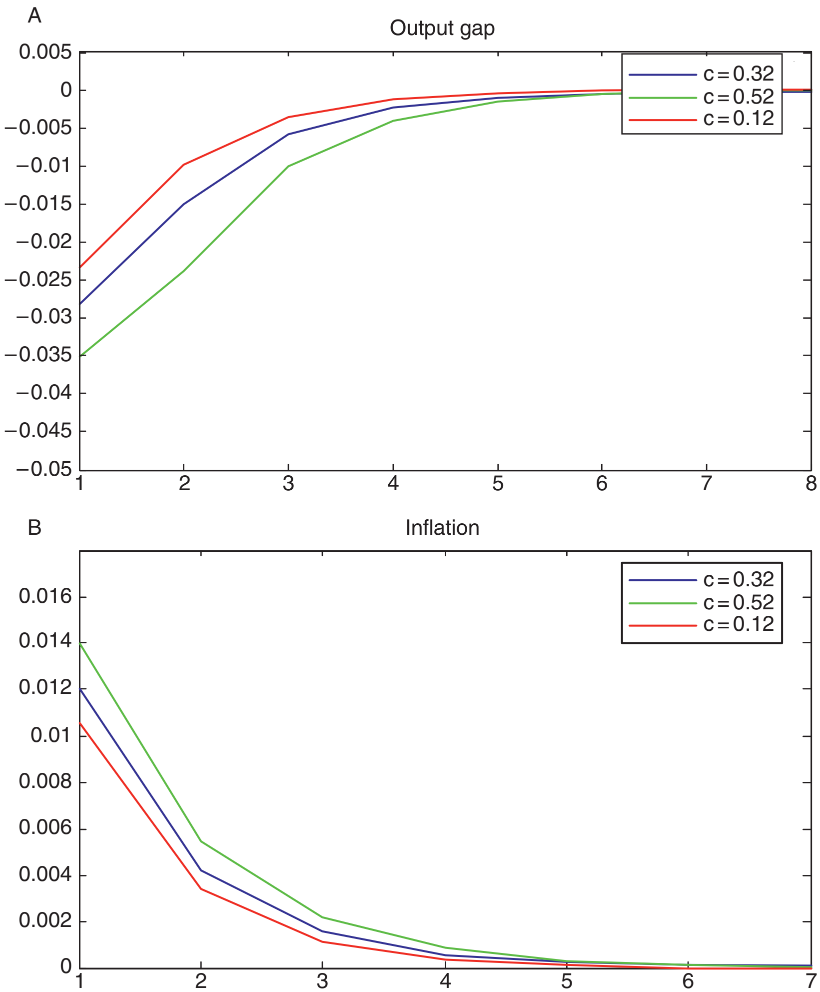

Figure 5 (Continued)

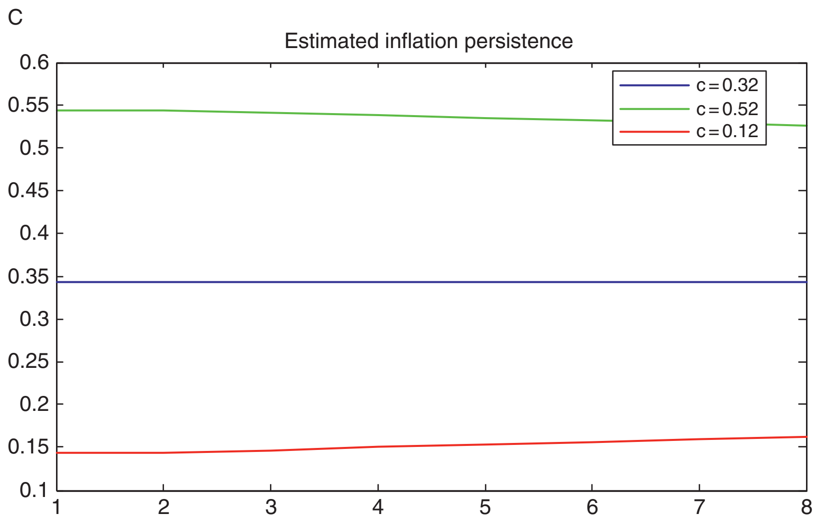

Figure 5 The mean dynamics of the output gap, inflation, and the estimated inflation persistence following a one-standard deviation cost-push shock: (a) output gap, (b) inflation, and (c) estimated inflation persistence.

levels of perceived (or estimated) inflation persistence on the side of the private sector. Panel (a) confirms the finding discussed earlier that as estimated persistence increases so does the output gap response (in absolute value). The stronger policy reaction helps mitigate the inflation response, although it is still the case (from panel b) that inflation increases by more when estimated inflation persistence is higher. This illustrates the worse trade-off the central bank is facing when estimated persistence is higher. Finally, from panel (c) it is apparent that the estimated persistent parameter adjusts gradually to its equilibrium value, which is lower than the degree of intrinsic persistence.

5.2.5 The intertemporal trade-off ($u_t = 0$)

Returning to Eq. (32) and departing from the assumption that $\pi_{t-1} = 0$, we can discuss the second term, on the right-hand side, which captures part of the optimal response to lagged inflation.

$$x_t = + \frac{\kappa\gamma(\chi_t - \delta_t) + \beta\kappa\chi_t\phi R_t^{-1} E_t V_c}{\kappa^2\delta_t + \lambda\chi_t^2}\pi_{t-1} + \ldots$$

Note that the first term in the numerator is zero when $\gamma = c_{t-1}$ (still using the simplifying assumption that $\delta_t = 1$). In such a case, inflation expectations adjust to past inflation just in line with the partial adjustment of inflation due to its intrinsic persistence (Eq. 16). Given the loss function, this is a desirable outcome. In the absence of any further shock, inflation will move exactly enough so that the quasi-difference of inflation will be zero. Note that when $\gamma > c_{t-1}$ or $\chi_t > 1$, the response of the output gap to past inflation, according to this effect, is positive. Hence, past inflation justifies expansionary policy. At first sight, this is counterintuitive. However, the reason is clear; when estimated persistence is below intrinsic persistence, past inflation does not feed enough into inflation expectations to stabilize the quasi-difference of inflation. To approach such a situation an expansionary policy must be followed. This factor is important because it shows that, in the context of this model, there is a cost associated with pushing the estimated persistence parameter too low.

However, another important point to make is that, in general, the second term in the numerator of the reaction coefficient will be negative and dominate the first term ensuring a negative response of the output gap to inflation. This term reflects the intertemporal trade-off the central bank is facing between stabilizing the output gap and steering the perceived degree of inflation persistence by inducing forecast errors. In our simulations it turns out that the expected marginal cost (the marginal impact on the expected present discounted value of all future losses) of letting estimated inflation persistence increase is always positive, that is, $V_c < 0$ and large. Intuitively, as discussed earlier, a lower degree of perceived persistence will lead to a much smaller impact of future cost-push shocks on inflation, which tends to stabilize inflation, its quasi-difference, and the output gap. As a result, under optimal policy the central bank will try to lower the perceived degree of inflation persistence. As is clear by updating Eq. (26) by the private sector, it can do so by engineering unexpectedly low inflation when past inflation is positive and conversely by unexpectedly reducing the degree of deflation when past inflation is negative. In other words, to reap the future benefits of lowering the degree of perceived inflation persistence, monetary policy will tighten if past inflation is positive and will ease if past inflation is negative. Overall, this effect justifies a counterveiling response to lagged inflation, certainly in the case of $\gamma = c_{t-1}$, when the first term in the numerator is zero.

Finally, the third term in Eq. (32) is also interesting. We have already noticed that when $\pi_{t-1} = 0$, $E_t(V_\pi) = 0$ and this term plays no role. Now, if $\pi_{t-1} > 0$, and $u_t = 0$, then $E_t(V_\pi) < 0$ and this will reinforce the negative effect of inflation on the output gap previously discussed. More explicitly, if lagged inflation is positive, this term will contribute to a negative output gap — tight monetary policy — even in the absence of a contemporary shock. This effect will contribute to stabilizing inflation close to zero. In the case $\pi_{t-1} < 0$, and $u_t = 0$, in contrast $E_t(V_\pi) > 0$. Thus, when lagged inflation is negative, this term will contribute to a positive output gap — loose monetary policy — even in the absence of a contemporary shock. Again this effect contributes to stabilizing inflation close to zero.

Figures 6a and 6b summarize some of the important features of the shape of the policy function (32) in the calibrated model. Figure 6a plots the output gap (on the vertical axis) as a function of lagged inflation and the perceived degree of inflation persistence for a zero cost-push shock and assuming that the moment matrix R equals its average for a particular realization of *c*. A number of features are worth repeating. First, when lagged inflation and the cost-push shock are zero, the output gap is also zero irrespective of the estimated degree of inflation persistence. Second, when the shock is zero, the response to inflation and deflation is symmetric. Third, as the estimated persistence of inflation increases, the output gap response to inflation (and deflation) rises.

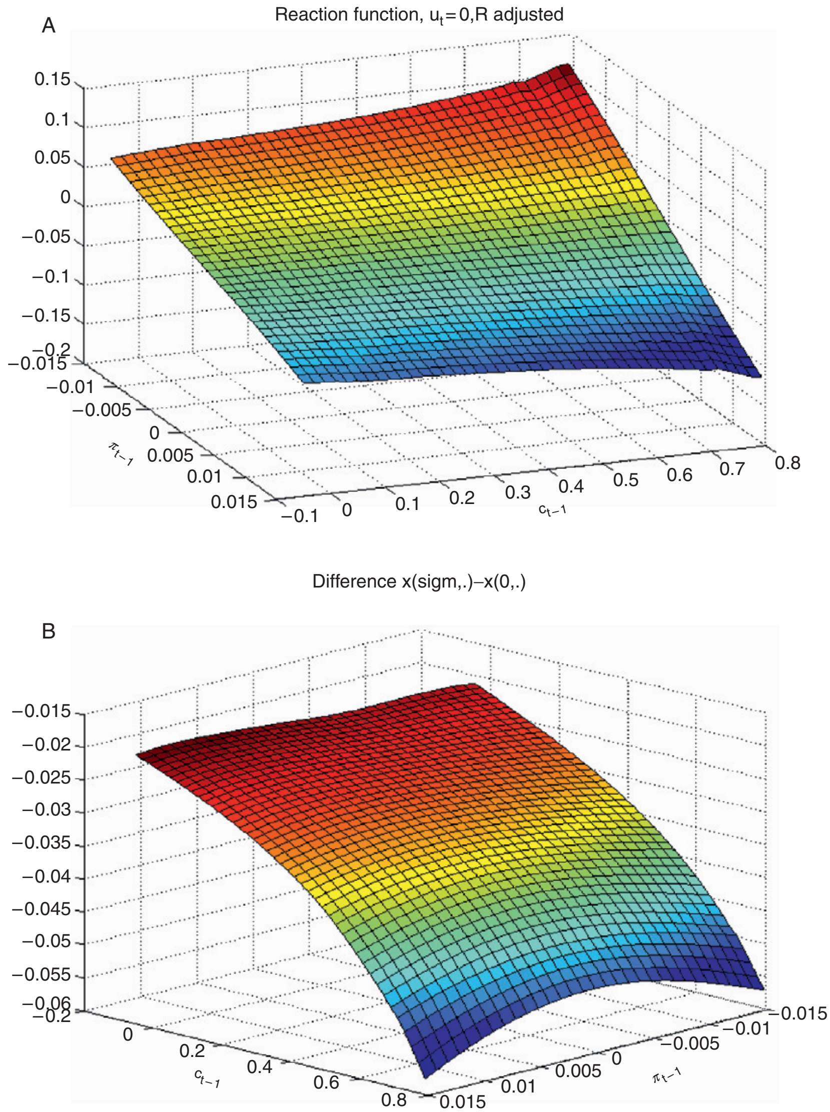

Figure 6 The policy function output gap as a function of lagged inflation and the estimated degree of inflation persistence. (a) reaction function, $u_t = 0$, R adjusted and (b) difference x(sigm,.)-x(0,.).

It is then interesting to see how the output gap response differs when a positive cost-push shock hits the economy. This is shown in Figure 6b, which plots the differences in output gap response to a positive one-standard deviation cost-push shock and zero cost-push shock as a function of lagged inflation and the perceived persistence parameter. The output gap response is always negative and increases with the estimated degree of inflation persistence. This figure also shows the nonlinear interaction with lagged inflation. In particular, the output gap response becomes stronger when inflation is already positive.

5.2.6 Some sensitivity analysis

How do the results depend on some of the calibrated parameters? First, we investigate how the results change with a different gain and a different degree of price stickiness. Second, we look at the impact of increasing the weight on output gap stabilization in the central bank's loss function.

Figure 7 plots the realization of the average perceived inflation persistence in economies with different gains and two different degrees of price stickiness ($\alpha = 0.66$, corresponding to our baseline calibration and a higher degree of price stickiness, $\alpha = 0.75$). Remember that $(1 - \alpha)$ measures the proportion of firms changing prices optimally each period. The other parameters are as in the calibration reported in Table 1. We focus on the perceived degree of persistence because this gives an idea

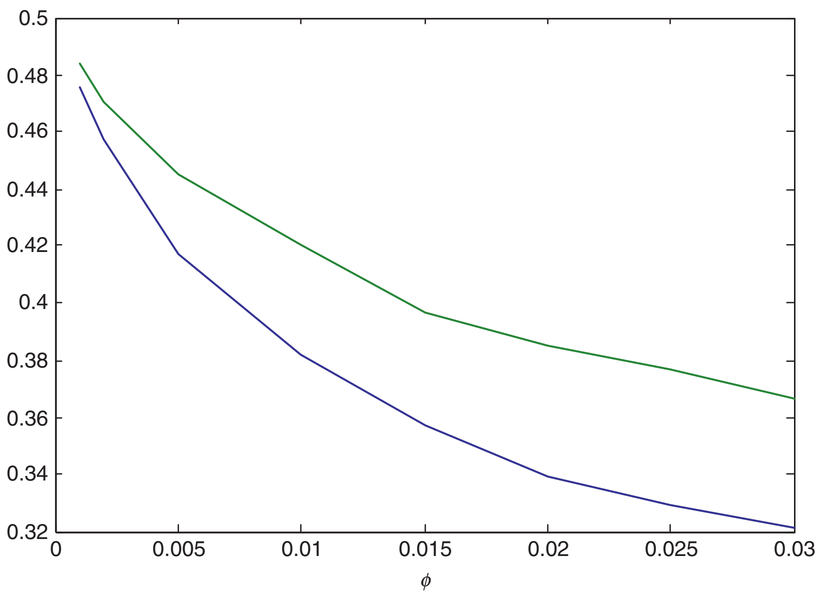

Figure 7 Sensitivity analysis: average estimated persistence in function of the gain and the degree of price stickiness.

about how the trade-off between lowering inflation persistence and stabilizing the output gap changes as those parameters change. As discussed earlier, when the gain is zero, the optimal policy converges to the simple discretion rule and the estimated degree of persistence equals the degree of intrinsic persistence in the economy (0.5 in the benchmark case). In this case, the central bank can no longer steer inflation expectations and the resulting equilibrium outcome is the same as under rational expectations. Figure 7 shows that an increasing gain leads to a fall in the average perceived degree of inflation persistence. With a higher gain, agents update their estimates more strongly in response to unexpected inflation developments. As a result, the monetary authority can more easily affect the degree of perceived persistence, which affects the trade-off in favor of lower inflation persistence. Figure 7 also shows that a higher degree of price stickiness increases the degree of inflation persistence. Again the intuition is straightforward. With higher price stickiness, it is more costly in terms of variation in the output gap to affect the degree of inflation persistence through unexpected inflation.

Finally, we look at the impact of increasing the weight on output gap stabilization in the central bank's loss function. Figure 8 shows that increasing the weight λ from 0.002 to 0.012 shifts the distribution of the estimated degree of inflation persistence to the right. The mean increases from 0.33 to 0.45. A higher weight on output gap stabilization makes it more costly to affect the private sector's estimation of the degree of inflation persistence leading to a higher average degree of inflation persistence.

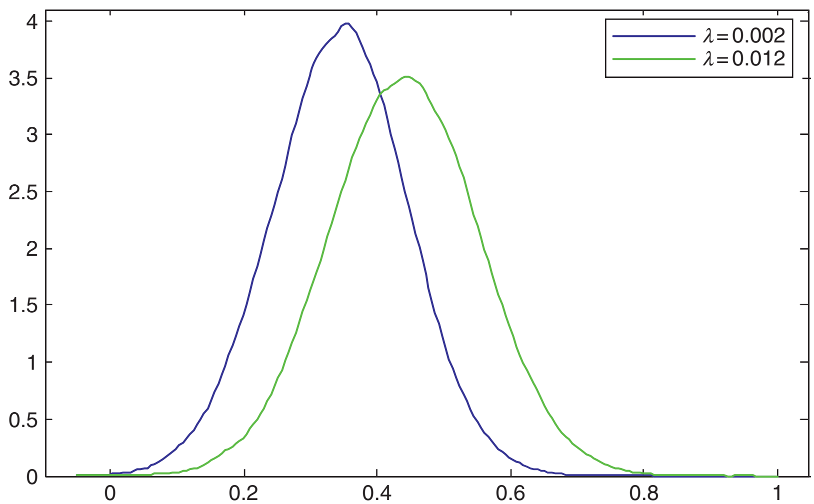

Figure 8 Distribution of estimated inflation persistence as a function of the weight on output gap stabilization.

Overall, the analysis in this section is closely related to the work of Orphanides and Williams. For example in Orphanides and Williams (2005a) they showed that, for the case of linear feedback rules, inflation persistence increases when adaptive learning is substituted for rational expectations. They also showed that a stronger response to inflation helps limit the increase in inflation persistence and that, in such a context, a strategy of stricter inflation control helps to reduce both inflation and output gap volatility. Gaspar et al. (2006a) found that under adaptive learning optimal policy responds persistently to cost-push shocks. Such a persistent response to shocks allows central banks to stabilize inflation expectations, and reduce inflation persistence and inflation variance at little cost in terms of output gap volatility. Persistent policy responses and well-anchored inflation expectations resemble optimal monetary policy under commitment and rational expectations. However, as explained previously, the mechanisms are very different. In the case of rational expectations, it operates through expectations of *future* policy. In the case of adaptive learning, it operates through a reduction in inflation persistence, as perceived by economic agents, given the past history determined by shocks and policy responses. There is no dichotomy between the two mechanisms anchoring inflation expectations. On the contrary, the central bank's ability to influence expectations about the future course of policy rates and its track record in preserving stability are complements.

6. SOME FURTHER REFLECTIONS

Before concluding it is worth making a few additional reflections. First, the analysis in most of this chapter differs from the large literature on monetary policy making under uncertainty; that is, when the central bank faces uncertainty about the data, the shocks, the model, or the way agents form expectations.[30] A few papers have studied the interaction between learning on behalf of the private agents and the uncertainty faced by the central bank. For example, Orphanides and Williams (2005a, 2007) assumed the central bank has imperfect knowledge about the natural interest rate and unemployment and show how the interaction with the constant gain learning by private agents further constrains the actions of the central bank. In particular, it puts a premium on responding relatively more to inflation rather than an imperfectly measured output gap. Evans and Honkapohja (2003a,b) found that expectations-based rules continue to ensure converge to the rational expectations equilibrium in a model where both the private sector and the central bank are learning.[31] A highly relevant paper in this context is Woodford (2010), which develops a concept of policy robustness where policymakers set monetary policy so that agents' expectations are distorted away from

[30] See the literature referenced in Hansen and Sargent (Chapter 20 in this volume) and Taylor and Williams (Chapter 15, this volume).

[31] Other papers are Dennis and Ravenna (2008) and Evans and McGough (2007).

rational expectations within some class of near rational expectations. In line with the results presented in Section 5.2, Woodford (2010) found that the principles of monetary policy under rational expectations are robust to these types of deviations from rational expectations by the private agents.

Second, in most of the chapter we have focused on the monetary policy implications of constant-gain or declining-gain least-squares learning in the formation of expectations. A number of papers have analyzed alternative types of learning. For example, Branch and Evans (2007) and Brazier, Harrison, King, and Yates (2008) assumed that private agents may use different forecast methods with the proportion of agents using specific forecast methods changing over time according to relative forecast performance. Similarly, Arifovic, Bullard, and Kostyshyna (2007) and De Grauwe (2008) used social learning where agents copy better forecasting methods and discard less successful techniques. Bullard, Evans, and Honkapohja (2008) analyzed a case where "expert" judgment, resulting from the perceived presence of extraneous factors, becomes almost self-fulfilling. The authors show how to adjust monetary policy to prevent these near-rational "exuberance equilibria." Overall, the introduction of these alternative ways of learning by private agents strengthens the case for managing inflation expectations by responding more aggressively to inflationary shocks.

Finally, it is also important to cover the issue of structural (including policy regime) change. In the presence of structural change it is only natural to acknowledge that it will take time for economic agents to learn about the new environment. More generally, real-time analysis of economic developments is made difficult by pervasive and fast economic change and by imperfect knowledge about the true structure of the economy. Adaptive learning provides a way to model explicitly the transition dynamics associated with structural change. In doing so models with adaptive learning go beyond credibility as a binary variable, a characteristic of standard rational expectations models (see Section 1).

Ferrero (2007) reasonably argued that, in addition to determinacy and E-stability of equilibrium, it is important to consider the characteristics of the transition to equilibrium and, in particular, how fast agents' beliefs approach rational expectations. Using the baseline model described previously and a forward-looking version of the Taylor rule as in

$$i_t = \gamma + \gamma_\pi E_t^* \pi_{t+1} + \gamma_x E_t^* x_{t+1} + \gamma_g g_t,$$

he showed that, by responding strongly to expected inflation, the monetary authority can shorten transition and increase the speed of convergence. Ferrero (2007) showed that, in the absence of a bias in inflation expectations, fast learning improves social welfare. In the presence of expectations' bias important qualifications apply, illustrating the importance of accurate monitoring of inflation expectations in the actual conduct of monetary policy.

Gaspar et al. (2010) also look at a question related to transition between monetary policy regimes. Specifically they considered the transition from an inflation targeting regime to a price level path stability regime. They showed that the speed of convergence depends on the speed of learning. They also found that the ex ante desirability of regime change depends on the speed of learning. Very slow learning implies very slow transition and the costs of switching may outweigh the permanent benefits from regime shift. Nevertheless, they argued that for empirically reasonable learning algorithms regime switching would be worthwhile for the example they considered. Earlier, Gaspar et al. (2006b) discussed transition dynamics associated with disinflation. They argued that the patterns observed are in line with the facts of the United States disinflation in the 1980s.

7. CONCLUSIONS

This chapter looks at the monetary policy implications when private sector expectations are determined in accordance to adaptive learning. As in Orphanides and Williams (2005b) and Woodford (2010), our main conclusion is that the fundamental policy prescriptions under model consistent expectations continue to hold, or are even strengthened, by limited departures from rational expectations. Specifically, when expectations are formed in accordance with adaptive learning, the gains from anchoring inflation and inflation expectations increase significantly. Optimal policy under adaptive learning stabilizes inflation and inflation expectations mainly through persistent responses to cost-push shocks. The previous remark explains why, in our numerical examples, the simple commitment rule performs well under adaptive learning. By responding persistently to cost-push shocks, the simple commitment rule is able to significantly lower the degree of estimated inflation persistence relative to the simple discretion rule. It is worthwhile stressing that the simple commitment rule is able to approximate quite closely the outcomes that could be obtained under full optimal policy.

In our setup, monetary policy actions have intra- and intertemporal effects. For example, we have seen that monetary policy responds relatively strongly to lagged inflation and to inflation shocks when the estimated persistence parameter is high. In such a case the central bank, facing positive inflation, will push down estimated persistence by generating unexpectedly low inflation (in the case of deflation by generating unexpectedly high inflation). In our model simulations the intertemporal, long-term considerations dominate optimal policy when trade-offs between intra- and intertemporal considerations arise. The importance of intertemporal considerations helps to explain why optimal policy under adaptive learning pushes down the estimated persistence parameter to values well below intrinsic inflation persistence and the equilibrium value under the simple rule. By behaving in this way, optimal monetary policy provides an anchor for inflation and inflation expectations, thus contributing to the overall stability of the economy and to better macroeconomic outcomes as evaluated by the social loss function. We view optimal monetary policy under adaptive learning as illustrating (once more) why medium term price stability and anchoring inflation expectations is key in environments characterized by endogenous inflation expectations.

We have also found that, even in the context of a simple model, the characterization of optimal policy becomes very involved. It is easy to imagine how much more difficult such a characterization would become if we would try to reckon the complexity of actual policy choices and the prevalence of economic change. Such considerations clearly limit the possibility of using our framework in a *prescriptive* way. However, the results in this chapter suggest that Woodford's (2003) case for emphasizing central banking as management of expectations comes out even stronger when adaptive learning substitutes for model consistent expectations.

REFERENCES

Arifovic, J., Bullard, J., Kostyshyna, O., 2007. Social learning and monetary policy rules. Federal Reserve Bank of St. Louis Working paper.

Basdevant, O., 2005. Learning processes and rational expectations: An analysis using a small macro-econometric model for New Zealand. Econ. Model. 22, 1074–1089.

Beechey, M.J., Johannsen, B.K., Levin, A., 2007. Are long-run inflation expectations anchored more firmly in the Euro area than in the United States?. CEPR Discussion Paper 6536.

Benhabib, J., Schmitt-Grohe, S., Uribe, M., 2001. The perils of Taylor rules. J. Econ. Theory 96, 40–69.

Bernanke, B., 2007. Inflation expectations and inflation forecasting. NBER Summer Institute, Cambridge, MA Remarks at.

Bernanke, B., Woodford, M., 1997. Inflation forecasts and monetary policy. J. Money Credit Bank. 24, 653–684.

Blinder, A.S., Ehrmann, M., Fratzscher, M., De Haan, J., Jansen, D.J., 2008. Central bank communications and monetary policy: A survey of theory and evidence. J. Econ. Lit. 46 (4), 910–945.

Boivin, J., Kiley, M., Mishkin, F., 2010. How has the monetary policy transmission mechanism evolved over time? In: Friedman, B., Woodford, M. (Eds.), Handbook of monetary economics. vol. 3A. Elsevier, Amsterdam.

Branch, W.A., 2004. The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. Econ. J. 114, 592–621.

Branch, W., Carlson, J., Evans, G.W., McGough, B., 2007. Adaptive learning, endogenous inattention and changes in monetary policy. University of Oregon Working paper.

Branch, W., Carlson, J., Evans, G.W., McGough, B., 2009. Monetary policy, endogenous inattention and the volatility trade-off. Econ. J. 119, 123–157.

Branch, W., Evans, G.W., 2006. A simple recursive forecasting model. Econ. Lett. 91, 158–166.

Branch, W., Evans, G.W., 2007. Model uncertainty and endogenous volatility. Rev. Econ. Dyn. 10, 207–237.

Brazier, A., Harrison, R., King, M., Yates, T., 2008. The danger of inflating expectations of macroeconomic stability: Heuristic switching in an overlapping generations monetary model. International Journal of Central Banking 4 (2), 219–254.

Bullard, J., Mitra, K., 2002. Learning about monetary policy rules. J. Monet. Econ. 49, 1105–1129.

Bullard, J., Schaling, E., 2009. Monetary policy, indeterminacy and learnability in a two-block world economy. J. Money Credit and Banking 41 (8), 1585–1612.

Bullard, J., Singh, A., 2008. Worldwide macroeconomic stability and monetary policy rules. Monet. Econ 55 (1), s34–s47.
Bullard, J., Evans, G.W., Honkapohja, S., 2008. Monetary policy, judgment and near-rational exuberance. Am. Econ. Rev. 98 (3), 1163–1177.
Carlstrom, C.T., Fuerst, T.S., 2004. Learning and the central bank. J. Monet. Econ. 51, 327–338.
Castelnuovo, E., Nicoletti-Altimari, S., Palenzuela, D.R., 2003. Definition of price stability, range and point inflation targets: The anchoring of long-term inflation expectations. In: Issing, O. (Ed.), Background studies for the ECB's evaluation of its monetary policy strategy. European Central Bank, Frankfurt, pp. 43–90.
Clarida, R., Galí, J., Gertler, M., 1999. The science of monetary policy: A new Keynesian perspective. J. Econ. Lit. 37, 1661–1707.
Clarida, R., Gali, J., Gertler, M., 2000. Monetary policy and macroeconomic stability: Evidence and some theory. Q. J. Econ. CXV (1), 147–180.
De Grauwe, P., 2008. DSGE modeling when agents are imperfectly informed. European Central Bank, Working Paper 897.
Dennis, R., Ravenna, F., 2008. Learning and optimal monetary policy. J. Econ. Dyn. Control 32 (6), 1964–1994.
Dixit, A.K., Stiglitz, J.E, 1977. Monopolistic competition and optimum product diversit. Am. Econ. Rev. 67, 297–308.
Duffy, J., Xiao, W., 2007. Investment and monetary policy: Learning and determinacy of equilibria. University of Pittsburgh, Working paper 324.
Evans, G.W., Honkapohja, S., 1998. Economic dynamics with learning: New stability results. Rev. Econ. Stud. 6, 23–44.
Evans, G.W., Honkapohja, S., 2001. Learning and expectations in macroeconomics. Princeton University Press, Princeton, NJ.
Evans, G.W., Honkapohja, S., 2003a. Adaptive learning and monetary policy design. J. Money Credit Bank. 35, 1045–1072.
Evans, G.W., Honkapohja, S., 2003b. Expectations and the stability problem for optimal monetary policies. Rev. Econ. Stud. 70, 807–824.
Evans, G.W., Honkapohja, S., 2003c. Expectational stability of stationary sunspot equilibria in a forward-looking model. J. Econ. Dyn. Control 28, 171–181.
Evans, G.W., Honkapohja, S., 2003d. Friendman's money supply rule versus optimal interest rate policy. Scott. J. Polit. Econ. 50, 550–566.
Evans, G.W., Honkapohja, S., 2005. Policy interaction, expectations and the liquidity trap. Rev. Econ. Dyn. 8, 303–323.
Evans, G.W., Honkapohja, S., 2006. Monetary policy, expectations and commitment. Scand. J. Econ. 108, 15–38.
Evans, G.W., Honkapohja, S., 2008a. Expectations, learning and monetary policy: An overview of recent research. In: Schmidt-Hebbel, K., Walsh, C. (Eds.), Monetary policy under uncertainty and learning 2009, Central Bank of Chile, Santiago, pp. 27–76.
Evans, G.W., Honkapohja, S., 2008b. Robust learning stability with operational monetary policy rules. In: Schmidt-Hebbel, K., Walsh, C. (Eds.), Monetary policy under uncertainty and learning 2009, Central Bank of Chile, Santiago, pp. 145–170.
Evans, G.W., Honkapohja, S., 2009. Expectations, deflation traps and macroeconomic policy. Bank of Finland, Research Discussion Papers 24.
Evans, G.W., McGough, B., 2005a. Monetary policy, indeterminacy and learning. J. Econ. Dyn. Control 29, 1809–1840.
Evans, G.W., McGough, B., 2005b. Stable sunspot solutions in models with predetermined variables. J. Econ. Dyn. Control 29, 601–625.
Evans, G.W., McGough, B., 2007. Optimal constrained interest-rate rules. J. Money Credit Bank. 39, 1335–1356.
Evans, G.W., Guse, E., Honkapohja, S., 2008. Liquidity traps, learning and stagnation. Eur. Econ. Rev. 52, 1438–1463.

Ferrero, G., 2007. Monetary policy, learning and the speed of convergence. J. Econ. Dyn. Control 31, 3006–3041.
Gali, J., Gertler, M., 1999. Inflation dynamics: A structural econometric analysis. J. Monet. Econ. 44 (2), 195–222.
Gaspar, V., Smets, F., 2002. Monetary policy, price stability and output gap stabilisation. International Finance 5, 193–202.
Gaspar, V., Smets, F., Vestin, D., 2006a. Adaptive learning, persistence and optimal monetary policy. J. Eur. Econ. Assoc. 4, 376–385.
Gaspar, V., Smets, F., Vestin, D., 2006b. Monetary policy over time. Macroecon. Dyn 1–23.
Gaspar, V., Smets, F., Vestin, D., 2010. Is time ripe for price level path stability? In: Siklos, P., Bohl, M., Wohar, M. (Eds.), The challenges of central banking. Cambridge University Press, Cambridge, UK.
Hansen, L., Sargent, T., 2010. Wanting robustness to misspecification. In: Friedman, B., Woodford, M. (Eds.), Handbook of monetary economics. Elsevier, Amsterdam.
Honkapohja, S., Mitra, K., 2004. Are non-fundamental equilibria learnable in models of monetary policy? J. Monet. Econ. 51, 1743–1770.
Judd, K.L., 1998. Numerical methods in economics. MIT Press, Cambridge, MA.
Kurozumi, T., 2006. Determinacy and expectational stability of equilibrium in a monetary sticky-price model with Taylor rule. J. Monet. Econ 53, 827–846.
Kurozumi, T., Van Zandweghe, W., 2007. Investment, interest rate policy and equilibrium stability. J. Econ. Dyn. Control 32 (5), 1489–1516.
Llosa, G., Tuesta, V., 2006. Determinacy and learnability of monetary policy rules in small open economies. IADB Working Paper 576.
Llosa, G., Tuesta, V., 2007. E-stability of monetary policy when the cost channel matters. Mimeo.
Lubik, T., Schorfheide, F., 2004. Testing for indeterminacy: An application to U. S. monetary policy. Am. Econ. Rev. 94 (1), 190–217.
Marcet, A., Nicolini, P., 2003. Recurrent hyperinflations and learning. Am. Econ. Rev. 93, 1476–1498.
Marcet, A., Sargent, T., 1989. Convergence of least-squares learning mechanisms in self-referential linear stochastic models. J. Econ. Theory 48, 337–368.
Milani, F., 2006. A Bayesian DSGE model with infinite-horizon learning: Do "mechanical" sources of persistence become superfluous? International Journal of Central Banking 2 (3), 87–106.
Milani, F., 2007. Expectations, learning and macroeconomic persistence. J. Monet. Econ. 54, 2065–2082.
Milani, F., 2009. Adaptive learning and macroeconometric inertia in the Euro area. J. Common Mark. Stud. 47 (3), 579–599.
Miranda, M.J., Fackler, P., 2002. Applied computational economics and finance. MIT Press, Cambridge, MA.
Molnar, K., Santoro, S., 2006. Optimal monetary policy when agents are learning. Working Paper 1.
Murray, J., 2007. Empirical significance of learning in new Keynesian model with firm-specific capital. Mimeo.
Muth, J., 1961. Rational expectations and the theory of price movements. Econometrica 29, 315–335.
Orphanides, A., 2005. Comment on: The incredible Volcker disinflation. J. Monet. Econ. 52 (5), 1017–1023.
Orphanides, A., Williams, J.C., 2002. Robust Monetary Policy Rules with Unknown Natural Rates. Brookings Pap. Econ. Act. 63–118.
Orphanides, A., Williams, J.C., 2005a. The decline of activist stabilization policy: Natural rate misperceptions, learning and expectations. J. Econ. Dyn. Control 29, 1927–1950.
Orphanides, A., Williams, J.C., 2005b. The Inflation-Targetting Debate. In: Bernanke, B., Woodford, M. (Eds.), Imperfect knowledge, inflation expectations, and monetary policy. The University of Chicago Press, Chicago, pp. 201–234 (Chapter 5).
Orphanides, A., Williams, J.C., 2005c. Inflation scares and forecast-based monetary policy. Rev. Econ. Dyn. 8, 498–527.
Orphanides, A., Williams, J.C., 2007. Robust monetary policy with imperfect knowledge. J. Monet. Econ. 54, 1406–1435.
Pfajfar, D., Santoro, E., 2007. Credit market distortions, asset prices and monetary policy. Mimeo.

Pfajfar, D., Santoro, E., 2009. Heterogeneity, learning and information stickiness in inflation expectations. University of Tilburg, Mimeo.
Preston, B., 2005. Learning about monetary policy rules when long horizon expectations matter. International Journal of Central Banking 1 (2), 81–126.
Slobodyan, S., Wouters, R., 2009. Estimating a medium-scaled DSGE model with expectations based on small forecasting models. National Bank of Belgium, Mimeo.
Smets, F., 2004. Maintaining price stability: How long is the medium term? J. Monet. Econ. 50, 1293–1309.
Svensson, L., 2003. What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules. J. Econ. Lit. 41, 426–477.
Svensson, L., Woodford, M., 2005. Implementing optimal monetary policy through inflation forecast targeting. In: Bernanke, B., Woodford, M. (Eds.), The inflation-targeting debate. University of Chicago Press, Chicago, IL.
Taylor, J.B., Williams, J.C., 2010. Simple and robust rules for monetary policy. In: Friedman, B., Woodford, M. (Eds.), Handbook of monetary economics. Elsevier, Amsterdam.
Trichet, J.C., 2009. The ECB's enhanced credit support. Keynote Address. University of Munich.
Walsh, C.E., 2009. Inflation targeting: What have we learned? International Finance 12 (2), 195–233.
Woodford, M., 2003. Interest and prices: foundations of a theory of monetary policy. Princeton University Press, Princeton, NJ.
Woodford, M., 2010. Robustly optimal monetary policy with near rational expectations. Am. Econ. Rev. 100 (1), 274–303.