DISCUSSION PAPER SERIES

No. 6640

EXPECTATIONS, LEARNING AND MONETARY POLICY: AN OVERVIEW OF RECENT RERSEARCH

George W Evans and Seppo Honkapohja

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP6640.asp**

ISSN 0265-8003

EXPECTATIONS, LEARNING AND MONETARY POLICY: AN OVERVIEW OF RECENT RERSEARCH

George W Evans, University of Oregon and University of St Andrews
Seppo Honkapohja, Bank of Finland, University of Cambridge and CEPR

Discussion Paper No. 6640
January 2008

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 6640

January 2008

ABSTRACT

Expectations, Learning and Monetary Policy: An Overview of Recent Rersearch*

Expectations about the future are central for determination of current macroeconomic outcomes and the formulation of monetary policy. Recent literature has explored ways for supplementing the benchmark of rational expectations with explicit models of expectations formation that rely on econometric learning. Some apparently natural policy rules turn out to imply expectational instability of private agents' learning. We use the standard New Keynesian model to illustrate this problem and survey the key results about interest-rate rules that deliver both uniqueness and stability of equilibrium under econometric learning. We then consider some practical concerns such as measurement errors in private expectations, observability of variables and learning of structural parameters required for policy. We also discuss some recent applications including policy design under perpetual learning, estimated models with learning, recurrent hyperinflations, and macroeconomic policy to combat liquidity traps and deflation.

JEL Classification: D84, E31 and E52
Keywords: determinacy, fluctuations, imperfect knowledge, interest-rate setting, learning and stability

George W Evans
Department of Economics
1285 University of Oregon
Eugene, OR 97403-1285
USA
Email: gevans@oregon.uoregon.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=136065

Seppo Honkapohja
Bank of Finland
P.O.Box 160
FI-00101 Helsinki
FINLAND
Email: seppo.honkapohja@bof.fi

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=101990

* Prepared for Bank of Chile Conference “Monetary Policy under Uncertainty and Learning”, November 2007. We thank Carl Walsh for useful comments. Financial support from National Science Foundation Grant No. SES-0617859 and ESRC grant RES-000-23- 1152 is gratefully acknowledged. Any views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland.

Submitted 02 January 2008

1 Introduction

The conduct of monetary policy in terms of interest rate or other rules has been extensively studied in recent research.[1] This literature gives a central role for forecasts of future inflation and output, and the question of whether monetary policy should be forward-looking has been subject to discussion and debate. Bank of England Inflation Reports, see Bank of England (2007), and the June and December Issues of the Monthly Bulletin of the European Central Bank, see European Central Bank (2007), present private sector forecasts as well as internal macroeconomic projections. Empirical evidence on Germany, Japan and the US since 1979 provided by Clarida, Gali, and Gertler (1998) suggests that central banks are forward-looking in practice.

The rational expectations (RE) hypothesis, the standard benchmark in macroeconomics since the seminal work of Lucas (1976) and Sargent and Wallace (1975), has been employed in most of the research on monetary policy and interest rate rules. The most common formulation of the RE hypothesis is based on the assumption that the private agents and the policy-maker know the "true model of the economy", except for unforecastable random shocks.[2] The RE assumption is excessively strong: neither private agents nor policy-makers have perfect knowledge of the economy. In reality, economists formulate and estimate models that are used to make macroeconomic forecasts and carry out policy analysis. These models are re-estimated and possibly reformulated as new data becomes available. In other words, economists engage in learning processes about the economy as they attempt to improve their knowledge of the economy.

Formal study of these learning processes and their implications for macroeconomic dynamics and policy-making are becoming an increasingly important line of research in macroeconomics.[3] This research is based on a principle of cognitive consistency stating that private agents and policy-makers

[1] Woodford (2003) is a monumental treatise on the subject, while the text of Walsh (2003) provides an accessible graduate-level treatment. For surveys see e.g. Clarida, Gali, and Gertler (1999), and McCallum (1999).

[2] It should be noted that some papers do extend the standard notion of RE equilibrium to an equilibrium with limited information. In many cases such extensions assume that economic agents do not observe some variables but do know the structure of the economy.

[3] Evans and Honkapohja (2001) provide a treatise on the analysis of adaptive learning and its implications in macroeconomics. Evans and Honkapohja (1999), Evans and Honkapohja (1995), Marimon (1997), Sargent (1993) and Sargent (1999) provide surveys of the field.

in the economy behave like applied economists and econometricians. Thus, it is postulated that expectations of macroeconomic variables are formed by using statistical or other formal forecasting models and procedures.

An important policy question is whether the processes of learning create new tasks and constraints for macroeconomic policy. An affirmative answer to this question has been demonstrated by the recent work on learning and monetary policy.[4] This view is also reflected in recent speeches by two prominent Central Bank Governors, see Trichet (2005) and Bernanke (2007). This research has shown that interest-rate setting by monetary policy-makers faces two fundamental problems.

First, some of the proposed interest rate rules may not perform well when the expectations of the agents are out of equilibrium. The consequences of errors in forecasting, and the resulting correction mechanisms, may create instability in the economy. For (usually non-optimal) instrument rules, Bullard and Mitra (2002) consider the stability of rational expectations equilibrium (REE) when monetary policy is conducted using variants of the Taylor rule. These rules work well only under certain parameter restrictions, and Bullard and Mitra suggests that monetary policy-making should take into account the learnability constraints on the parameters of policy behavior. For optimal monetary policy Evans and Honkapohja (2003c) and Evans and Honkapohja (2006) show that certain standard forms of optimal interest rate setting by the Central Bank can lead to expectational instability as economic agents unsuccessfully try to correct their forecast functions over time. Evans and Honkapohja also propose a new rule for implementing optimal policy that always leads to stability under learning.

Second, monetary policy rules, including some formulations for optimal setting of the instrument and some Taylor rules based on forecasts of inflation and output gap, can create multiple equilibria, also called indeterminacy of equilibria.[5] Under indeterminacy there are multiple, even continua of REE and the economy need not settle on the desired REE. The possible rest points have been studied using stability under learning as a selection criterion, see Honkapohja and Mitra (2004) and the further papers by Carlstrom and Fuerst (2004) and Evans and McGough (2005a). We note that

[4] Evans and Honkapohja (2003a) and Bullard (2006) provide surveys of the recent research.

[5] This was first noted by Bernanke and Woodford (1997), Woodford (1999b), and Svensson and Woodford (2005). The problem was systematically explored for Taylor rules by Bullard and Mitra (2002).

indeterminacy is not a critical problem if the fundamental REE is the only stable equilibrium under learning. Moreover, indeterminacy need not arise if the forward-looking interest rate rule is carefully designed, as was shown by Bullard and Mitra (2002), Evans and Honkapohja (2003c), and Evans and Honkapohja (2006).

The central message from these studies is that monetary policy has important new tasks when agents' knowledge is imperfect and agents try to improve their knowledge through learning. Policy should be designed to facilitate learning by private agents so that expectations do not create instability in the economy.

Recently, many further aspects of expectations, learning and monetary policy have been analyzed in the rapidly expanding literature. In this paper we provide a non-technical overview of this research program. In the first part of the paper we begin by reviewing the basic theoretical results, after which we take up some immediate practical concerns that can arise in connection with rules for interest rate setting. These include issues of observability in connection with private forecasts as well as with current output and inflation data. A second concern is knowledge of the structure of the economy that is required for implementation of optimal interest rate policies.

In the second part of the paper we provide an overview of the recent and ongoing developments in the literature. We first summarize research on learnability of REE when the basic New Keynesian model is extended to incorporate further features of the economy. After this we discuss four topics of applied interest in more detail: policy design under perpetual learning, estimated models with learning, recurrent hyperinflations, and macroeconomic policy to combat liquidity traps and deflation.

2 The Model

We conduct our discussion using the New Keynesian model that has become the workhorse in the analysis of monetary policy, and we employ directly its linearized version. The original nonlinear framework is based on a representative consumer and a continuum of firms producing differentiated goods under monopolistic competition. Nominal stickiness of prices arises from constraints of firms on the frequency of price changes, as originally suggested by Calvo (1983).

The behavior of the private sector is summarized by the two equations

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

which is the "IS" curve derived from the Euler equation for consumer optimization, and

$$\pi_t = \lambda x_t + \beta E_t^* \pi_{t+1} + u_t, \tag{2}$$

which is the price setting rule for the monopolistically competitive firms, often called the New Keynesian Phillips or aggregate supply curve.

Here x_t and π_t denote the output gap and inflation rate for period t, respectively. i_t is the nominal interest rate, expressed as the deviation from the steady state real interest rate. The determination of i_t will be discussed below. $E_t^* x_{t+1}$ and $E_t^* \pi_{t+1}$ denote private sector expectations of the output gap and inflation next period. Since our focus is on learning behavior, these expectations need not be rational (E_t without $*$ denotes RE). The parameters φ and λ are positive and β is the discount factor with $0 < \beta < 1$.

For brevity we do not discuss details of the derivation of equations (1) and (2). It should be pointed out that the derivation is based on individual Euler equations under (identical) subjective expectations, together with aggregation and definitions of the variables. The Euler equations for the current period give the decisions as functions of the expected state next period. Rules for forecasting the next period's values of the state variables are the other ingredient in the description of individual behavior. It is assumed that given forecasts, private agents make decisions according to the Euler equations.[6]

The shocks g_t and u_t are assumed to be observable and follow

$$\begin{pmatrix} g_t \\ u_t \end{pmatrix} = F \begin{pmatrix} g_{t-1} \\ u_{t-1} \end{pmatrix} + \begin{pmatrix} \tilde{g}_t \\ \tilde{u}_t \end{pmatrix}, \tag{3}$$

where

$$F = \begin{pmatrix} \mu & 0 \\ 0 & \rho \end{pmatrix},$$

[6] This kind of behavior is boundedly rational but in our view reasonable since agents attempt to meet the margin of optimality between the current and the next period. Other models of bounded rationality are possible. Recently, Bruce Preston has proposed a formulation in which long horizons matter in individual behavior, see Preston (2005) and Preston (2006).

$0 < |\mu| < 1$, $0 < |\rho| < 1$ and $\tilde{g}_t \sim iid(0, \sigma_g^2)$, $\tilde{u}_t \sim iid(0, \sigma_u^2)$ are independent white noise. g_t represents shocks to government purchases and/or potential output. u_t represents any cost push shocks to marginal costs other than those entering through x_t. For simplicity, we assume throughout the paper that μ and ρ are known (if not, they could be estimated).

The model is closed by an equation describing interest rate setting by the Central Bank.[7] One approach examines "instrument rules" under which i_t is directly specified in terms of key macroeconomic variables without explicit policy optimization. A prominent example of this type is the standard Taylor (1993) rule, i.e.,

$$i_t = \pi_t + 0.5(\pi_t - \bar{\pi}) + 0.5x_t,$$

where $\bar{\pi}$ is the target level of inflation and the target level of the output gap is zero. (Recall that i_t is specified net of the real interest rate, which in the standard Taylor rule is usually set at 2%). More generally, Taylor rules are of the form $i_t = \chi_0 + \chi_\pi \pi_t + \chi_x x_t$. For convenience (and without loss of generality) we will take the inflation target to be $\bar{\pi} = 0$ so that this class of rules takes the form

$$i_t = \chi_\pi \pi_t + \chi_x x_t \text{ where } \chi_\pi, \chi_x > 0. \tag{4}$$

Variations of the Taylor rule replace π_t and x_t by lagged values or by forecasts of current or future values.

Alternatively, interest rate policy can be derived explicitly to maximize a policy objective function. This is frequently taken to be of the quadratic loss form, i.e.

$$E_t \sum_{s=0}^{\infty} \beta^s \left[(\pi_{t+s} - \bar{\pi})^2 + \alpha x_{t+s}^2 \right], \tag{5}$$

where $\bar{\pi}$ is the inflation target. This type of optimal policy is often called "flexible inflation targeting" in the current literature, see e.g. Svensson (1999) and Svensson (2003). The policy-maker is assumed to have the same discount factor β as the private sector. α is the relative weight placed by

[7] We follow the common practice of leaving hidden the government budget constraint and the equation for the evolution of government debt. This is acceptable provided fiscal policy appropriately accommodates the consequences of monetary policy for the government budget constraint. The interaction of monetary and fiscal policy can be important for the stability of equilibria under learning, see Evans and Honkapohja (2007a), McCallum (2003) and Evans, Guse, and Honkapohja (2007). We discuss some aspects of the interaction below.

the policy-maker on the output target, and strict inflation targeting would be the case $\alpha = 0$. The loss function (5) can alternatively be viewed as a quadratic approximation to the welfare function of a representative agent; see Rotemberg and Woodford (1999) and Woodford (2003).[8]

The literature on optimal policy under RE distinguishes between optimal discretionary policy, in which the policy-maker is unable to commit to policies for future periods, and optimal policy in which such commitment is possible. Under commitment the policy-maker can do better because of the effect on private expectations, but commitment policy exhibits time inconsistency, in the sense that policy-makers would have an incentive to deviate from the policy in the future. Assuming that the policy has been initiated at some point in the past (the "timeless perspective" described by Woodford (1999a)), and setting $\bar{\pi} = 0$, the first order condition specifies

$$\lambda \pi_t + \alpha(x_t - x_{t-1}) = 0 \tag{6}$$

in every period.

Condition (6) for optimal policy with commitment is not a complete specification of monetary policy, since one must also provide a "reaction function" for i_t that implements the policy. A number of interest rate rules are consistent with the model (1)-(2), the optimality condition (6), and RE. However, some of the ways of implementing "optimal" monetary policy can make the economy vulnerable to either indeterminacy or expectational instability or both, while other implementations are robust to these difficulties.

We will consider "fundamentals-based" and "expectations-based" rules. The basic fundamentals-based rule depends only on the observable exogenous shocks g_t, u_t and also on x_{t-1}

$$i_t = \psi_x x_{t-1} + \psi_g g_t + \psi_u u_t, \tag{7}$$

where the optimal coefficients are determined by the structural parameters and the policy objective function. The coefficients ψ_i are chosen so that the effects of aggregate demand shocks g_t are neutralized and so that for inflation shocks u_t the optimal balance is struck between output and inflation effects. The dependence of i_t on x_{t-1} is optimally chosen to take advantage of the

[8] In this formulation α is a function of various "deep" structural parameters in the fully microfounded version of the model.

effects on expectations of commitment to a rule.[9]

Expectations-based optimal rules are advocated in Evans and Honkapohja (2003c) and Evans and Honkapohja (2006) because, as further discussed below, fundamentals-based optimal rules are often unstable under learning. If private expectations are observable they can be incorporated into the interest rate rule. If this is done appropriately, the REE will be stable under learning and thus optimal policy can be successfully implemented. The essence of these rules is that they do not assume RE on the part of private agents, but are designed to feed back on private expectations in such a way that they generate convergence to the optimal REE under learning. (If expectations are rational, these rules deliver the optimal REE.)

The optimal expectations-based rule under commitment is

$$i_t = \delta_L x_{t-1} + \delta_\pi E_t^* \pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t. \tag{8}$$

The coefficients of (8) are

$$\begin{aligned} \delta_L &= \frac{-\alpha}{\varphi(\alpha+\lambda^2)}, \\ \delta_\pi &= 1 + \frac{\lambda\beta}{\varphi(\alpha+\lambda^2)},\ \delta_x = \varphi^{-1}, \\ \delta_g &= \varphi^{-1},\ \delta_u = \frac{\lambda}{\varphi(\alpha+\lambda^2)}. \end{aligned} \tag{9}$$

This rule is obtained by combining the IS curve (1), the price setting equation (2) and the first order optimality condition (6), treating the private expectations as given.[10]

Interest rate rules based on observations of x_t and π_t that (outside the REE) only approximate the first order optimality condition (6) have been considered by Svensson and Woodford (2005). They suggest a set of "hybrid"rules, the simplest of which would be

$$i_t = \psi_x x_{t-1} + \psi_g g_t + \psi_u u_t + \theta[\pi_t + (\alpha/\lambda)(x_t - x_{t-1})],\ \theta > 0. \tag{10}$$

[9]The coefficients of the interest rate rule (7) are $\psi_x = \bar{b}_x[\varphi^{-1}(\bar{b}_x - 1) + \bar{b}_\pi]$, $\psi_g = \varphi^{-1}$, and $\psi_u = [\bar{b}_\pi + \varphi^{-1}(\bar{b}_x + \rho - 1)]\bar{c}_x + \bar{c}_\pi \rho$. Here $\bar{b}_x = (2\beta)^{-1}[\varsigma - (\varsigma^2 - 4\beta)^{1/2}]$ with $\varsigma = 1 + \beta + \lambda^2/\alpha$, and $\bar{b}_\pi = (\alpha/\lambda)(1 - \bar{b}_x)$, $\bar{c}_x = -[\lambda + \beta\bar{b}_\pi + (1-\beta\rho)(\alpha/\lambda)]^{-1}$, $\bar{c}_\pi = -(\alpha/\lambda)\bar{c}_x$.

[10]Under optimal discretionary policy the first order condition is $\lambda\pi_t + \alpha x_t = 0$ and the coefficients are identical except that $\delta_L = 0$. The discretionary case is analyzed in Evans and Honkapohja (2003c).

This rule combines the fundamentals-based rule (7) with the correction for the first order condition.[11] Note that under RE the rule (10) delivers the optimal equilibrium. Another hybrid rule has been suggested by McCallum and Nelson (2004) taking the form

$$i_t = \pi_t + \theta[\pi_t + (\alpha/\lambda)(x_t - x_{t-1})], \tag{11}$$

where $\theta > 0$.

3 Determinacy and Stability under Learning

Given an interest rate rule we can obtain the reduced form of the model and study its properties under RE. Two basic properties of interest are determinacy of the RE solution and stability under learning of the REE of interest.

Consider the system given by (1), (2), (3) and one of the i_t policy rules (4), (7), (8), (10) or (11). Defining the vectors

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

the reduced form can be written as

$$y_t = ME_t^* y_{t+1} + Ny_{t-1} + Pv_t \tag{12}$$

for appropriate matrices M, N and P. In the case of the rule (4) we have $N = 0$ and thus the simpler system

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

We now briefly describe the concepts of determinacy/indeterminacy and stability under adaptive (least-squares) learning using the general frameworks (12) and (13).

The first issue of concern is whether under RE the system possesses a unique stationary 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", i.e. REE depending on

[11] The model and the interest rate rule analyzed in Svensson and Woodford (2005) incorporate additional information lags.

extraneous random variables that influence the economy solely through the expectations of the agents.[12]

The second issue concerns stability under adaptive learning. In the introduction we stressed the principle of cognitive consistency according to which agents in the model are assumed to behave like econometricians or statisticians when they form their expectations. In the next section this approach is formalized in terms of the "Perceived Law of Motion" (PLM) describing the beliefs of the agents. These beliefs concern the stochastic process followed by the endogenous (and exogenous) variables that need to be forecasted. The parameters of the PLM are updated using an appropriate statistical technique, called an adaptive learning rule, and forecasts are made using the estimated PLM at each moment of time. If private agents follow an adaptive learning rule like recursive least squares to update the parameters of their forecasting model, will the RE solution of interest be stable, i.e. reached asymptotically by the learning process? If not, the REE is unlikely to be attained. This is the focus of the papers by Bullard and Mitra (2002), Bullard and Mitra (2007), Evans and Honkapohja (2003c), Evans and Honkapohja (2006) and many others.

3.1 Digression on Methodology

Consider first the simpler reduced form (13) under RE. It is well-known that the condition for determinacy is that both eigenvalues of the 2×2 matrix M lie inside the unit circle. In the determinate case the unique stationary solution will be of the "minimal state variable" (or MSV) form

$$y_t = \bar{c}v_t,$$

where $\bar{c}$ is a 2×2 matrix that is easily computed. If instead one or both roots lie inside the unit circle then the model is indeterminate. There will still be a solution of the MSV form, but there will also be other stationary solutions.

[12] If the model is indeterminate, one can ask whether the sunspot solutions are stable under learning. For a general discussion see Evans and Honkapohja (2001). In general, different forms of sunspot solutions exist, and stability under learning can depend on the particular representation, see Evans and Honkapohja (2003b) and Evans and McGough (2005b).

Next, we consider the system (13) under learning. Suppose that agents believe that the solution is of the form

$$y_t = a + cv_t, \tag{14}$$

but that the 2×1 vector a and the 2×2 matrix c are not known but instead are estimated by the private agents. (14) is the PLM of the agents. Note that we include an intercept vector because, although for theoretical simplicity we have translated all variables to have zero means, in practice agents will need to estimate intercept as well as slope parameters.[13]

With this PLM and parameter estimates (a, c) agents would form expectations as

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

where either F is known or is also estimated. Inserting these expectations into (13) and solving for y_t we get the implied "Actual Law of Motion" or ALM, i.e. the law that y_t would follow for a fixed PLM (a, c).[14] This is given by

$$y_t = Ma + (P + McF)v_t.$$

We have thus obtained an associated mapping from PLM to ALM given by

$$T(a, c) = (Ma, P + McF),$$

and the RE solution $(0, \bar{c})$ is a fixed point of this map.

Under real-time learning the sequence of events is as follows.[15] 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 $E_t^* y_{t+1} = a_t + c_t F v_t$ (assuming for convenience that F is known). Following, for example, the rule (4) the central bank sets the interest rate i_t, and y_t is generated according to (1) and (2) together with the interest rate rule. This temporary equilibrium is summarized by (13). Then at the beginning of $t + 1$ agents add the new

[13] We remark that private agents and the policymaker are here assumed to observe the shocks v_t. If v_t is not observable then the PLM would be adjusted to reflect relevant available information.

[14] The ALM describes the temporary equilibrium for given expectations as specified by the forecasts from the given PLM.

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

data point to their information set to update their parameter estimates to (a_{t+1}, c_{t+1}), e.g. using least squares, and the process continues. The question of interest is whether $(a_t, c_t) \to (0, \bar{c})$ over time.

It turns out that the answer to this question is given by the E-stability principle, which advises us to look at the differential equation

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

where τ denotes notional time. If the REE $(0, \bar{c})$ is locally asymptotically stable under this differential equation then the REE is stable under least squares and closely related learning rules. Conditions for local stability of this differential equation are known as expectational stability or "E-stability" conditions. We will also refer to these stability conditions as the "conditions for stability under adaptive learning", or the "conditions for stability under learning", or even "learnability" of equilibrium.

For the reduced form (13) it can be shown that the E-stability conditions are that (i) the eigenvalues of M have real parts less than one and (ii) all products of eigenvalues of M times eigenvalues of F have real parts less than one. It follows that for this reduced form the conditions for stability under adaptive learning are implied by determinacy but not vice versa.[16] This is not, however, a general result: sometimes E-stability is a stricter requirement than determinacy and in other cases neither condition implies the other.

Consider next the reduced form (12). Standard techniques are available to determine whether the model is determinate.[17] In the determinate case the unique stationary solution takes the MSV form

$$y_t = a + by_{t-1} + cv_t, \tag{15}$$

for appropriate values $(a, b, c) = (0, \bar{b}, \bar{c})$. In the indeterminate case there are multiple solutions of this form, as well as non-MSV REE.

To examine stability under learning we treat (15) as the PLM of the agents. Under real-time learning agents estimate the coefficients a, b, c of (15). This is a vector autoregression (VAR) with exogenous variables v_t. The estimates (a_t, b_t, c_t) are updated at each point in time by recursive least

[16] See McCallum (2007) for conditions when determinacy implies E-stability.

[17] The procedure is to rewrite the model in first order form and compare the number of non-predetermined variables with the number of roots of the forward-looking matrix that lie inside the unit circle.

squares. Once again it can be shown that the E-stability principle gives the conditions for local convergence of real-time learning.

For E-stability we compute the mapping from the PLM to the ALM as follows. The expectations corresponding to (15) are given by

$$E_t^* y_{t+1} = a + b(a + by_{t-1} + cv_t) + cFv_t, \tag{16}$$

where we are treating the information set available to the agents, when forming expectations, as including v_t and y_{t-1} but not y_t. (Alternative information assumptions would be straightforward to consider.) This leads to the mapping from PLM to ALM given by

$$T(a, b, c) = \left(M(I + b)a, Mb^2 + N, M(bc + cF) + P\right), \tag{17}$$

E-stability is again determined by the differential equation

$$\frac{d}{d\tau}(a, b, c) = T(a, b, c) - (a, b, c), \tag{18}$$

and the E-stability conditions govern stability under least-squares learning.

3.2 Results for Monetary Policy

We now describe the determinacy and stability results for the interest rate rules described in Section 2.

3.2.1 Taylor Rules

Bullard and Mitra (2002) consider Taylor-type rules and find that the results are sensitive to whether the i_t rule conditions on current, lagged or expected future output and inflation. In addition to assuming that $\chi_\pi, \chi_x \geq 0$, they assume that the serial correlation parameters in F are nonnegative. For the rule (4) the results are particularly straightforward and natural.[18] Bullard and Mitra (2002) show that the REE is determinate and stable under learning if and only if (using our notation)

$$\lambda(\chi_\pi - 1) + (1 - \beta)\chi_x > 0.$$

[18] Throughout we will assume that we are not exactly on the border of the regions of determinacy or stability.

In particular, if policy obeys the "Taylor principle" that $\chi_\pi > 1$, so that nominal interest rates respond at least one for one with inflation, then determinacy and stability are guaranteed.

If lagged or forward-looking Taylor rules are used the situation is more complicated and full analytical results are not available. For the lagged variable case they find that for $\chi_\pi > 1$ and $\chi_x > 0$ sufficiently small the policy leads to an REE that is determinate and stable under learning. For $\chi_\pi > 1$ but χ_x too large the system is explosive.

Bullard and Mitra (2002) also look at forward-looking versions of the Taylor rule, taking the form

$$i_t = \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1} \text{ where } \chi_\pi, \chi_x > 0, \tag{19}$$

where we can interpret $E_t^* \pi_{t+1}$ and $E_t^* x_{t+1}$ as identical one step ahead forecasts, based on least-squares updating, used by both private agents and policy-makers. They find that for $\chi_\pi > 1$ and $\chi_x > 0$ sufficiently small the policy leads to an REE that is determinate and stable under learning. Now for $\chi_\pi > 1$ and χ_x large the system is indeterminate, yet the MSV solution is stable under learning. However, there can also exist E-stable sunspot equilibria as was shown by Honkapohja and Mitra (2004) and discussed further by Carlstrom and Fuerst (2004) and Evans and McGough (2005a).

The Bullard and Mitra (2002) results emphasize the importance of the Taylor principle in obtaining stable and determinate interest rate rules. At the same time their results show that stability under learning must not be taken for granted, even when the system is determinate so that a unique stationary solution exists. The parameters of the policy rule χ_π, χ_x must be appropriately selected by the policy-maker when an instrument rule describes policy. Stability under learning provides a constraint for this choice.

3.2.2 Optimal Monetary Policy

Evans and Honkapohja (2006) focus on optimal monetary policy under commitment. It turns out that under the fundamentals-based policy rule (7), the economy is invariably unstable under learning. This is the case even though with this rule there are regions in which the optimal REE is determinate.[19]

[19] It can be noted that the learning stability results are sensitive to the detailed information assumptions. With PLM (15) if agents can make forecasts conditional also on y_t then under the fundamentals-based rule there are both regions of stability and instability, depending on the structural parameters.

The basic intuition for this result can be seen from the reduced form

$$\begin{pmatrix} x_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta + \lambda\varphi \end{pmatrix} \begin{pmatrix} E_t^* x_{t+1} \\ E_t^* \pi_{t+1} \end{pmatrix} + \begin{pmatrix} -\varphi\psi_x & 0 \\ -\lambda\varphi\psi_x & 0 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \end{pmatrix} + \begin{pmatrix} -\varphi\psi_u \\ 1 - \lambda\varphi\psi_u \end{pmatrix} u_t. \quad (20)$$

Since typically $\beta + \lambda\varphi > 1$, say, upward mistakes in $E_t^* \pi_{t+1}$ lead to higher π_t, both directly and indirectly through lower ex ante real interest rates, which under learning sets off a cumulative movement away from REE. The feedback from x_{t-1} under the fundamentals-based i_t rule with commitment (7) does not stabilize the economy. Figure 1 shows how divergence from the optimal REE occurs under the rule (7).[20]

The instability of the fundamentals-based rules, designed after all to obtain optimal policy, serves as a strong warning to policy-makers not to automatically assume that RE will be attained. It is necessary to examine explicitly the robustness of contemplated policy rules to private agent learning.

In Evans and Honkapohja (2003c) and Evans and Honkapohja (2006) we show how the problems of instability and indeterminacy can be overcome if private agents' expectations are observable, so that interest rate rules can be in part conditioned on these expectations. In Evans and Honkapohja (2006) we show that under the rule (8) the economy is determinate and the optimal REE is stable under private agent learning for all possible structural parameter values. The key to the stability results can be seen from the reduced form

$$\begin{pmatrix} x_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} 0 & -\frac{\lambda\beta}{\alpha+\lambda^2} \\ 0 & \frac{\alpha\beta}{\alpha+\lambda^2} \end{pmatrix} \begin{pmatrix} E_t^* x_{t+1} \\ E_t^* \pi_{t+1} \end{pmatrix} + \begin{pmatrix} \frac{\alpha}{\alpha+\lambda^2} & 0 \\ \frac{\alpha\lambda}{\alpha+\lambda^2} & 0 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \end{pmatrix} + \begin{pmatrix} -\frac{\lambda}{\alpha+\lambda^2} \\ \frac{\alpha}{\alpha+\lambda^2} \end{pmatrix} u_t. \quad (21)$$

In (21) the feedback from inflation expectations to actual inflation is stabilizing since the coefficient $\frac{\alpha\beta}{\alpha+\lambda^2}$ is less than one and the influence of x_{t-1} is also weak. Thus, deviations from RE are offset by policy and in such a way that

[20] Figures 1 and 2 are based on the calibration by McCallum and Nelson (1999). Using other calibrations would yield similar results.

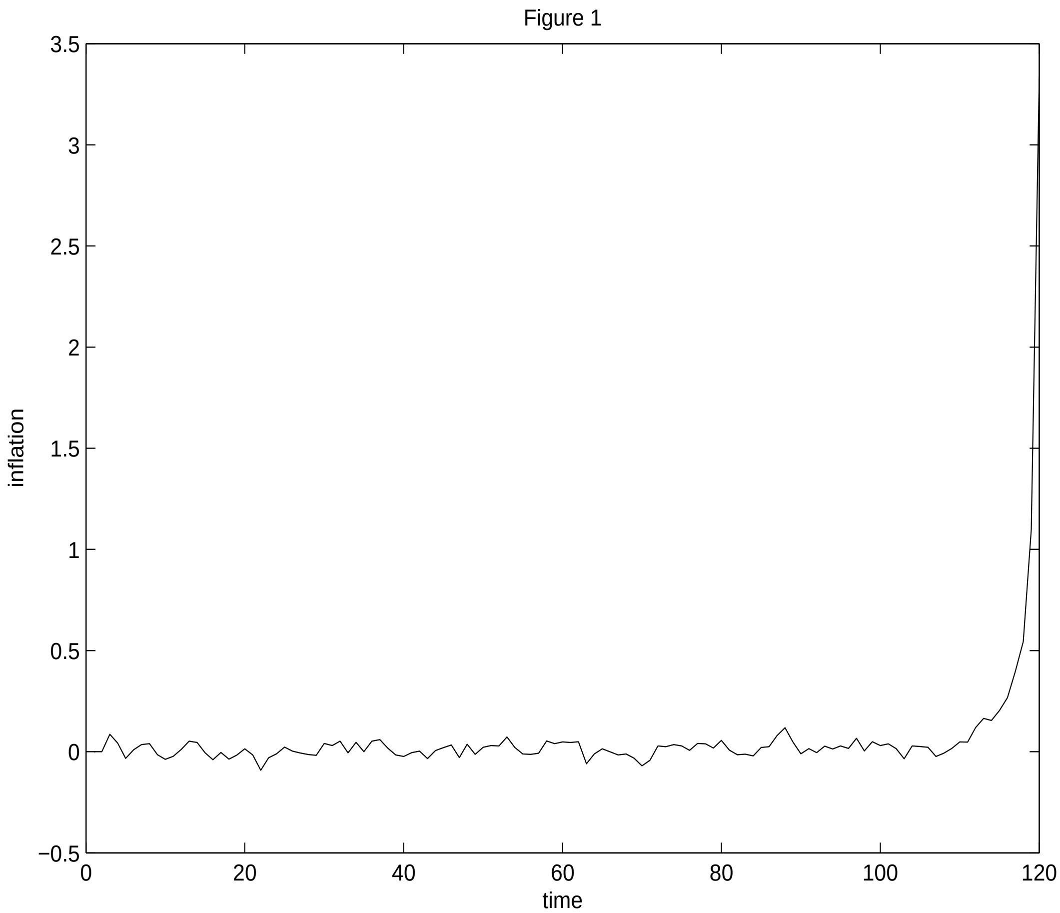

Figure 1: Instability with fundamentals-based rule

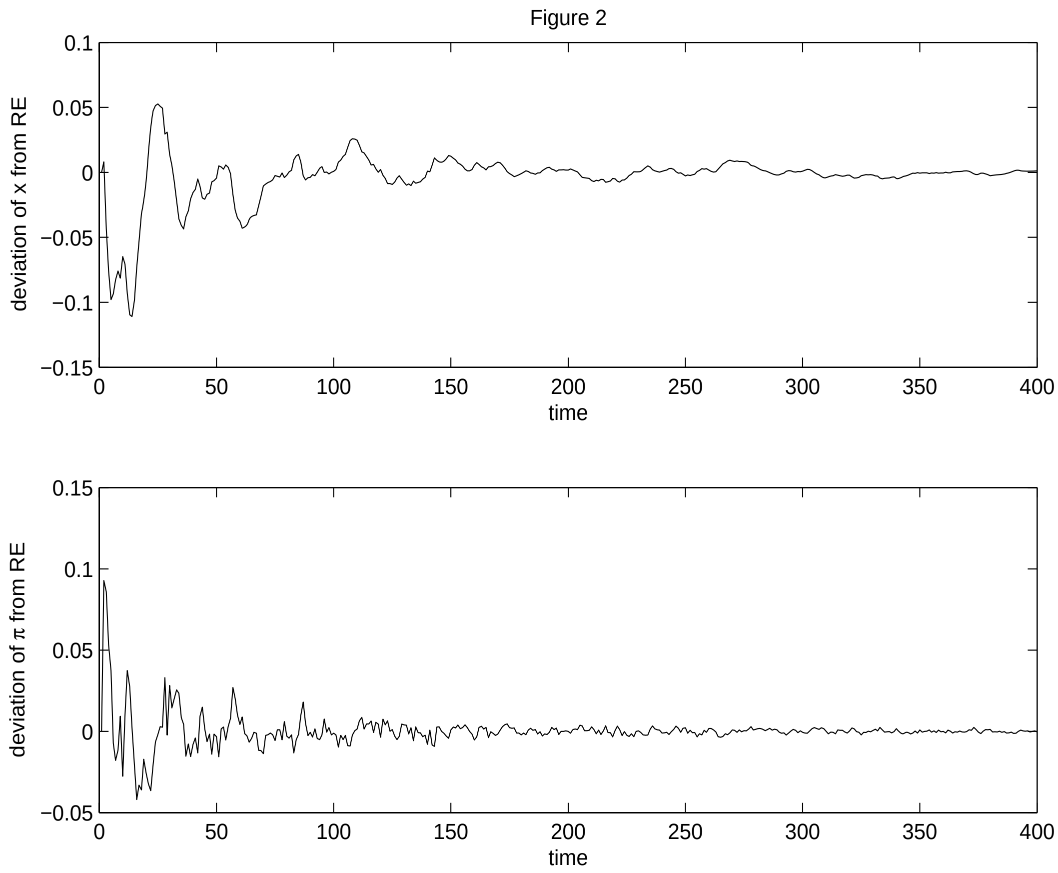

Figure 2: Stability with expectations-based rule

under learning private agents are guided over time to form expectations consistent with the optimal REE. Our expectations-based rule obeys a form of the Taylor principle since $\delta_\pi > 1$. Figure 2 illustrates convergence of learning under the rule (8).

Note that our optimal policy rule conditions on both private expectations and observable exogenous shocks, as well as lagged output. We also remark that, when computing the optimal expectations-based rule, it is important for the central bank to use the correct structural model of the IS and price setting relationships, which in turn depend on the specific form of boundedly rational individual behavior. For example, the form of the optimal expectations-based rule would be different if one adopted the long-horizon decision rules advocated by Preston (2005) and Preston (2006).

There are some cases in which variations of fundamentals-based rules can perform well, at least for a relevant region of structural parameter values. For the "hybrid"rule suggested by Svensson and Woodford (2005) numerical analysis shows that, in calibrated models, the rule (10) yields both determinacy and stability under learning for sufficiently high values of θ. Similarly, the hybrid rule suggested by McCallum and Nelson (2004) appears to deliver E-stability of the REE. Another favorable case emerges if the policy objective (5) is extended to include a motive for interest rate stabilization. Duffy and Xiao (2007b) show that in this case the fundamentals-based interest rules can deliver determinacy and E-stability for a region of parameter values that includes the usual calibrations used in the literature. However, see our comments below concerning stability with constant-gain learning for operational versions of these rules.

Finally, we remark that other formulations of monetary policy than interest rate rules could be analyzed. For example, policy could be formulated as a money supply rule, where a prominent case is the Friedman proposal for $k-$percent money growth. Evans and Honkapohja (2003d) show that Friedman's rule always delivers determinacy and E-stability in the standard New Keynesian model. However, it does not perform well in terms of the policy objective function.

3.3 Some Practical Concerns

Many of the i_t rules discussed above have the potential difficulty that they may not be operational, as discussed in McCallum (1999). For example, McCallum and Nelson (2004) note that it may be unrealistic to assume that policy-makers can condition policy on current x_t and π_t. Similarly, one could question whether accurate observations on private expectations are available. We consider these points in the reverse order. In the subsequent discussion we focus on the expectations-based rule (8), the Taylor rule (4) and the hybrid rules (10) and (11).

3.3.1 Observability of Private Expectations

The expectations-based rule (8) requires observations of current private expectations of future variables. While survey data on private forecasts of future inflation and various measures of future output exist, there are concerns about the accuracy of this data. If observations of expectations are subject to a white noise measurement error then our stability and determinacy results are unaffected. Furthermore, if measurement errors are small then the policy will be close to optimal. However, if measurement errors are large then this will lead to a substantial deterioration in performance. In this case one might consider substituting a proxy for such observations. Since we are assuming that private agents forecast by running VARs, the most natural proxy is for the Central Bank to estimate corresponding VARs and use these in (8).

Suppose now that agents and the Central Bank begin with different initial estimates, possibly have different learning rules and/or use data sets with different initial dates. When the private agents and the Central Bank are separately estimating and forecasting using VARs, we must distinguish between their expectations. An extended E-stability analysis for economies with heterogenous expectations gives the conditions for convergence of heterogeneous learning, as shown in Honkapohja and Mitra (2006). For the case of optimal discretionary policy and expectations-based interest rate rules this issue was analyzed in Honkapohja and Mitra (2005b). In Evans and Honkapohja (2003a) it was shown that using VAR proxies can also achieve convergence to the optimal REE with commitment.

We remark that the form of the extended E-stability conditions for heterogeneous learning depends on the nature of heterogeneity among agents. If the heterogeneities are transient (in the sense described in Honkapohja and Mitra (2006)), then the standard E-stability conditions directly apply. In cases of persistent heterogeneity the learning stability conditions are somewhat sensitive to the detailed assumptions. Additional restrictions are required for stability in some cases, e.g. if private agents estimate parameters using stochastic gradient techniques while the Central Bank uses least squares.

3.3.2 Non-Availability of Current Data

A difficulty with the standard Taylor rule (4) as well as some other rules, including the hybrid rules of Svensson and Woodford (2005) and McCallum and Nelson (2004), is that they presuppose that the policy-maker can observe both current output gap and inflation when setting the interest rate. McCallum (1999) has criticized such policy rules as not being "operational".

In the case of the Taylor rule, Bullard and Mitra (2002) show that this problem of non-observability can be avoided by the use of "nowcasts" $E_t^* y_t$

in place of the actual data y_t. Determinacy and E-stability conditions are not affected by this modification.

For the hybrid rules performance depends on the rule. Numerical analysis suggests that E-stability can still be achieved for Svensson-Woodford rule under standard values of the parameters. In contrast, for the McCallum-Nelson rule the situation is more complex. McCallum and Nelson (2004) suggest using forward expectations in place of actual data. If this is done, determinacy and stability under learning are no longer guaranteed, and sufficiently large values of the policy parameter θ induce both instability under learning and indeterminacy. This is unfortunate since large values of θ are needed to achieve a close approximation to optimal policy. Evans and Honkapohja (2003a) argue that the loss in welfare relative to the optimum is significant if θ is required to satisfy the constraints of E-stability and determinacy.

There is an additional issue with stability under learning that arises when current data are not observable for the policy-maker. If private agents are using constant-gain learning (see Section 5.2 for details), the stability conditions are more demanding. As discussed in Evans and Honkapohja (2007b), both hybrid rules suggested by Svensson and Woodford (2005) and McCallum and Nelson (2004), as well as the Taylor-type optimal rule of Duffy and Xiao (2007b), are subject to the problem of instability under constant-gain learning for many realistic gain parameter values.

3.3.3 Imperfect Knowledge of Structural Parameters

A third practical concern is that the use of optimal rules requires knowledge of the true values of the structural parameters on the part of the Central Bank. Evans and Honkapohja (2003c) and Evans and Honkapohja (2003a) extend the basic analysis to a situation where the Central Bank estimates the structural parameters φ and λ in equations (1)-(2) and in each period uses the current estimates in its optimal interest rate rule.[21] The basic results concerning optimal interest rate rules extend naturally to this situation. The fundamental-based rules under commitment and discretion are not learnable, while the corresponding expectations-based rules deliver convergence of simultaneous learning by the private agents and the Central Bank.

Since optimal monetary policy depends on structural parameters, uncertainty about their values is an issue, even if asymptotically their values can be

[21]It is natural to assume that the discount factor β and the policy weight α are known to the Central Bank.

learned by the Central Bank. Evans and McGough (2007) examine optimal Taylor-type rules based on Bayesian model averaging, where determinacy and stability under learning are imposed across all plausible structural parameter values.

The importance of structural uncertainty is also stressed in Orphanides and Williams (2007). Their model incorporates both imperfect knowledge about the natural rates of interest and unemployment and constant-gain learning by private agents. They emphasize monetary policy rules that are robust in all of these dimensions.

4 Further Developments

There has been a great deal of recent work that extends the results on monetary policy and learning. Several of these are discussed in some detail below.

One of the more significant issues, from an applied point of view, is the issue of "constant" gain or "perpetual" learning, in which private agents update estimates using least squares, but discount past data. Consequently, under learning agents' expectations never fully converge to the REE, but, provided the REE is stable, have expectations that are (asymptotically) in a neighborhood of the REE. Several papers discuss the issue of optimal policy when the learning process itself is incorporated into the optimal policy problem, either during the learning transition or under perpetual learning. The main papers are Orphanides and Williams (2005b), Molnar and Santoro (2006), Gaspar, Smets, and Vestin (2006), Gaspar, Smets, and Vestin (2005) and Orphanides and Williams (2007). A related issue studied by Ferrero (2007) concerns speed of convergence of learning for alternative policy rules. Arifovic, Bullard, and Kostyshyna (2007) consider the implications of social learning for monetary policy rules.

Extensions of the learning stability results to open economy and multi-country settings have been made by Llosa and Tuesta (2006), Bullard and Schaling (2006), Bullard and Singh (2006), Zanna (2006), and Wang (2006) among others. These papers examine both Taylor-type rules and interest-rate rules that target real exchange rates.

In the standard New Keynesian model, monetary policy works entirely via the demand side. Kurozumi (2006) and Llosa and Tuesta (2007) consider how determinacy and learning conditions are altered when monetary policy has direct effects on inflation. Kurozumi and Van Zandweghe (2007), Duffy and

Xiao (2007a) and Pfajfar and Santoro (2007a) have examined in detail how the learning stability conditions for Taylor rules are modified when capital is incorporated into the New Keynesian model. The results for models with capital depend on precisely how capital is modeled, i.e. on whether or not adjustment costs are included and on whether there is firm-specific capital or a rental market for capital. One result that emerges in some of these settings is that determinacy and E-stability requires the interest-rate rule to have a positive response to the output gap.

Detailed policy issues arise in which learning plays a key role. Some central banks often set monetary policy based on the constant interest rate that is expected to deliver a target inflation rate over a specified horizon. How this affects stability under learning is studied in Honkapohja and Mitra (2005a). Transparency and communication of targets and rules are further considered by Berardi and Duffy (2007) and Eusepi and Preston (2007).

While the New Keynesian model is based on a linearized set-up under Calvo-type pricing, nonlinear settings based on quadratic costs of price adjustments suggested by Rotemberg (1982) have been useful for studying the possibility of liquidity trap equilibria.[22] This issue was investigated under perfect foresight by Benhabib, Schmitt-Grohe, and Uribe (2001). This set-up was investigated under learning for the case of flexible prices by Evans and Honkapohja (2005) and in a sticky-price version by Evans, Guse, and Honkapohja (2007). The latter paper is discussed further below. Sticky-information models that incorporate learning have also been developed. See Branch, Carlson, Evans, and McGough (2006b) and Branch, Carlson, Evans, and McGough (2006a).

A number of theoretical learning topics have recently been pursued that have a bearing on monetary policy issues. Forward-looking Taylor rules can generate indeterminacy for some choices of parameters. In these cases can stationary sunspot equilibria be stable under learning? For the New Keynesian setting this issue has been examined by Honkapohja and Mitra (2004), Carlstrom and Fuerst (2004), and Evans and McGough (2005a), where conditions for stable sunspots are obtained in linearized models, and by Eusepi (2007), who looks at the question in a nonlinear setting. Evans, Honkapohja, and Marimon (2007) show that stable sunspot equilibria can arise in a cash-in-advance framework in which part of the government deficit is financed by

[22] Using a linearized New Keynesian model, the possibility of liquidity traps under learning was studied by Bullard and Cho (2005).

seigniorage.

Constant-gain learning raises the issue of the appropriate choice of gain parameter (see Evans and Honkapohja (1993), Chapter 14 of Evans and Honkapohja (2001) and Marcet and Nicolini (2003)). This issue is considered by Evans and Ramey (2006) in a simple monetary set-up in which private agents face an unknown regime switching process. This paper shows how the Lucas Critique, based on RE, can carry over to learning dynamics in which agents have misspecified models.

Monetary policy with near-rational expectations has been studied by Woodford (2005) and Bullard, Evans, and Honkapohja (2007). The Woodford paper develops a (minmax) concept of policy robustness in which policy-makers protect against agents' expectations being distorted away from RE within some class of near rational expectations. Bullard, Evans, and Honkapohja (2007) consider the possibility that "expert" judgement based on extraneous factors believed to be present can become almost self-fulfilling. They show how to alter monetary policy to protect against these near-rational "exuberance equilibria."

Heterogeneous expectations is another area increasingly receiving attention. Theoretical work on monetary policy that allows for learning heterogeneity across private agents, or between policy-makers and private agents, includes Evans, Honkapohja, and Marimon (2001), Giannitsarou (2003), Honkapohja and Mitra (2005b), and Honkapohja and Mitra (2006). A related line of thought (see Brock and Hommes (1997) and Branch and Evans (2006a)) emphasizes that private agents may have different types of predictors, with the proportions of agents using the different forecast methods changing over time according to relative forecast performance. For an application to monetary inflation models and monetary policy see Branch and Evans (2007) and Brazier, Harrison, King, and Yates (2006).

A number of empirical applications of learning to macroeconomics and monetary policy have recently been developed. Bullard and Eusepi (2005) and Orphanides and Williams (2005c) look at estimated models that focus on the explanation of the large increase in inflation rates in the 1970s. Incorporating learning as a way to explain persistence in New Keynesian models has been examined, using US data, by Milani (2005) and Milani (2007). First attempts to incorporate learning to applied stochastic general equilibrium (DSGE) models have most recently been made, see Slobodyan and Wouters (2007) and Murray (2007). Using least-squares learning models and/or dynamic predictors to explain expectations data has been studied in Branch

(2004), Branch and Evans (2006b), Orphanides and Williams (2005a), Basdevant (2005), Pfajfar (2007), and Pfajfar and Santoro (2007b).

Other important empirical learning papers include Marcet and Nicolini (2003), which studies hyperinflation in South American countries (this paper is discussed in detail below), Cogley and Sargent (2005), Sargent, Williams, and Zha (2006), Ellison and Yates (2007), and Carboni and Ellison (2007). The latter papers emphasize the importance of policy-maker model uncertainty and the role of central bank learning in explaining the historical evolution of inflation and unemployment in the post 1950 period.

In the next sections we discuss four recent topics that address important applied questions. Learning plays a crucial role in these analyzes, but the main focus in each case goes well beyond stability of REE under learning.

5 Perpetual Learning and Persistence

In the preceding sections our concern has been the stability of the REE under least-squares (LS) learning. That is, we used LS learning to assess whether an REE is attainable if we model agents as econometricians. Orphanides and Williams (2005b) (OW) show that taking the further step of replacing ("decreasing gain") LS learning with constant-gain learning has important implications for monetary policy, even if the REE is stable under learning.

OW work with a simple two-equation macro model. The first equation is a New Classical expectations-augmented Phillips curve with inertia:

$$\pi_{t+1} = \phi \pi^e_{t+1} + (1-\phi)\pi_t + \alpha y_{t+1} + e_{t+1}, \tag{22}$$

where π_{t+1} is the rate of inflation between period t and period $t+1$, π^e_{t+1} is the rate of inflation over this period expected at time t, y_{t+1} is the level of the output gap in $t+1$ and e_{t+1} is a white noise inflation shock. $(1-\phi)\pi_t$ represents intrinsic inflation persistence. We assume $0 < \phi < 1$.

The second equation is an aggregate demand relation that embodies a lagged policy effect,

$$y_{t+1} = x_t + u_{t+1}.$$

x_t is set by monetary policy at t and u_{t+1} is white noise. Through monetary policy it is assumed that policy-makers are able one period ahead to control aggregate output up to the unpredictable random disturbance u_{t+1}. This equation basically replaces the IS and LM curves. It is convenient for the task at hand, but of course suppresses issues of monetary control.

5.1 Optimal Policy under Rational Expectations

At time t the only state variable is π_t. Policy-makers have a target inflation rate π^* and care about the deviation of π_t from π^*. Their instrument is x_t and they are assumed to follow a rule of the form

$$x_t = -\theta(\pi_t - \pi^*). \tag{23}$$

Policy-makers also care about the output gap y_{t+1}. Since stable inflation requires $Ey_t = 0$, policy-makers are assumed to choose θ to minimize

$$\mathcal{L} = (1-\omega)Ey_t^2 + \omega E(\pi_t - \pi^*)^2.$$

This is a standard quadratic loss function. We can think of ω as reflecting policy-makers preferences, which may (or may not) be derived from the preferences of the representative agent.

Under RE, $\pi^e_{t+1} = E_t\pi_{t+1}$ and it follows that

$$\pi^e_{t+1} = \pi_t + \frac{\alpha}{1-\phi}x_t.$$

Substituting into (22) yields

$$\pi_{t+1} = \pi_t + \frac{\alpha}{1-\phi}x_t + \alpha u_{t+1} + e_{t+1}.$$

Substituting in the policy rule (23) yields

$$\tilde{\pi}_{t+1} = \left(\frac{1-\phi-\alpha\theta}{1-\phi}\right)\tilde{\pi}_t + \alpha u_{t+1} + e_{t+1},$$

where $\tilde{\pi}_t = \pi_t - \pi^*$.

Computing $E\tilde{\pi}_t^2$ and Ey_t^2 it is straightforward to minimize $\mathcal{L}$ over θ to get θ^P, the optimal choice of θ under RE. OW show that

$$\theta^P = \theta^P(\omega, (1-\phi)/\alpha),$$

and that θ^P is increasing in both ω and in the degree of inertia, $1-\phi$. Varying ω leads to an efficiency frontier, described by a familiar trade-off between σ_π and σ_y, sometimes called the "Taylor curve."

For this choice of feedback parameter, in the REE inflation follows the process

$$\begin{aligned}\pi_t &= c_0^P + c_1^P \pi_{t-1} + \text{noise}_t \\ E_t \pi_{t+1} &= c_0^P + c_1^P \pi_t,\end{aligned}$$

where

$$c_0^P = \frac{\alpha\theta^P}{1-\phi} \text{and } c_1^P = 1 - \frac{\alpha\theta^P}{1-\phi}.$$

Here noise_t is white noise. The superscript "P" refers to "perfect knowledge," which OW use as a synonym for RE.

Thus under RE the problem is quite straightforward. How "aggressive" policy should be with respect to deviations of inflation from target depends in a natural way on the structural parameters ϕ, α and the policy-maker preferences as described by ω.

5.2 Least-Squares Learning

We now make the crucial step of backing away from RE. Instead of assuming that agents are endowed *a priori* with RE, we model the agents as forecasting in the same way that an econometrician might: by assuming a simple time series model for the variable of interest, and by estimating its parameters and using the estimated model to forecast. Specifically, suppose private agents believe that inflation follows an AR(1) process, as it does in an REE, but that they do not know c_0^P, c_1^P. Instead they estimate the parameters of

$$\pi_t = c_0 + c_1 \pi_{t-1} + v_t$$

by a least-squares-type regression, and at time t forecast

$$\pi_{t+1}^e = c_{0,t} + c_{1,t}\pi_t.$$

Over time the estimates $c_{0,t}, c_{1,t}$ are updated as new data become available. We consider two cases for this updating.

First, suppose that agents literally do least squares using all the data. We assume that policy-makers do not explicitly take account of private agent learning and follow the feedback rule with $\theta = \theta^P$. Then, with "infinite memory" (no discounting of observations), one can show that

$$c_{0,t}, c_{1,t} \rightarrow c_0^P, c_1^P \text{ w.p.1.}$$

Asymptotically, we get the optimal REE.

OW make a small but significant change to the standard least-squares updating formula. With regular LS each data point counts equally. When expressed in terms of a recursive algorithm ("recursive least squares" or "RLS") the coefficients estimates $c_{0,t}, c_{1,t}$ are updated in response to the most recent data point with a weight proportion to the sample size $1/t$. We often say that RLS has a "decreasing gain" since the "gain" or weight on each data point is $\kappa_t = 1/t$, which declines towards 0 as $t \to \infty$. OW instead consider "constant gain" RLS in which past data is discounted. In terms of the RLS algorithm, this is accomplished technically by setting the gain, the weight on the most recent observation used to update estimates, to a small constant i.e. setting $\kappa_t = \kappa$ (e.g. 0.05). This is equivalent to using weighted least squares with weights declining geometrically in time as we move backwards from the current date.

Why would it be natural for agents to use a constant rather than decreasing gain? The main rationale for this procedure is that it allows estimates to remain alert to structural shifts. As economists, and as econometricians, we tend to believe that structural changes occasionally occur, and we might therefore assume that private agents also recognize and allow for this. Although in principle one might attempt to model the process of structural change, this tends to unduly strain the amount of knowledge we have about the economic structure. A reasonable alternative is to adjust parameter estimators to reflect the fact that recent observations convey more accurate information on the economy's law of motion than do data further in the past, and "constant gain" estimators are one very natural way of accomplishing this down-weighting of past data. Another possibility that is sometimes used in practice is to use a rolling data-window of finite length.[23]

5.3 Implications of Constant-Gain Least Squares

With constant-gain procedures, estimates no longer fully converge to the REE. The estimators $c_{0,t}, c_{1,t}$ converge instead to a stochastic process. Because of this OW use the term "perpetual learning" to refer to the constant gain case.

If the gain parameter κ is very small, then estimators will be close to

[23] Honkapohja and Mitra (2003) discuss the implications of bounded memory as a model of learning.

the REE values for most of the time with high probability, and output and inflation will be near their REE paths. Nonetheless, small plausible values like $\kappa = 0.05$ can lead to very different outcomes in the calibrations OW consider. They analyze the results using simulations, with $\phi = 0.75$ and $\alpha = 0.25$. They consider $\theta \in \{0.1, 0.6, 1.0\}$, which corresponds to weights $\omega = 0.01, 0.5$ and 1, respectively, under RE.

Their main findings are: (i) the standard deviations of $c_{0,t}$ and $c_{1,t}$ are large even though forecast performance remains good, (ii) there is a substantial increase in the persistence of inflation, compared to the REE, as measured by the AR(1) coefficient for π_t, and (iii) the policy frontier shifts out very substantially and sometimes in a non-monotonic way.

5.4 Policy Implications

Under perpetual learning by private agents, if policy-makers keep to the same class of rules

$$x_t = -\theta^S(\pi_t - \pi^*),$$

then they should choose a different θ than under RE. Here the notation θ^S is meant to indicate that we restrict policy-makers to choose from the same "simple" class of policy rules. There are four main implications for policy in the context of constant-gain (perpetual) learning by private agents. First, the "naive" policy choice, i.e. the policy that assumes RE ("perfect knowledge") on the part of agents, can be strictly inefficient when in fact the agents are following perpetual learning with $\kappa > 0$: there are cases in which increasing θ^S above θ^P would decrease the standard deviations of *both* inflation and output. Second, in general policy should be more hawkish, i.e. under perpetual learning the monetary authorities should pick a larger θ^S than if agents had RE.

Third, following a sequence of unanticipated inflation shocks, inflation doves (i.e. policy-makers with low θ reflecting a low ω) can do very poorly, as these shocks can lead expectations to temporarily but persistently deviate substantially from RE. Finally, if the inflation target π^* is known to private agents, so that they need estimate only the slope parameter c_1 using the PLM

$$\pi_{t+1} - \pi^* = c_1(\pi_t - \pi^*) + v_{t+1}$$

then the policy frontier is more favorable than when the intercept c_0 is not known. One way to interpret this is that central bank transparency is useful.

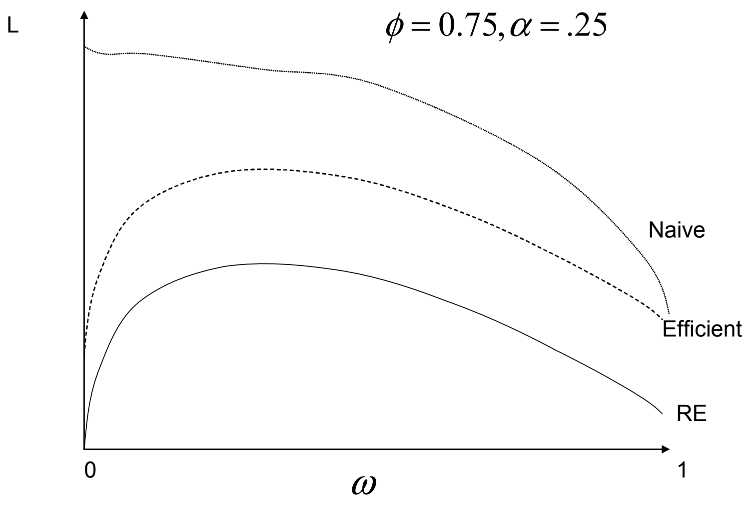

Figure 3: Policy-maker's loss

Figure 3 indicates how the performance of policy depends on expectations formation and what the policy-maker assumes about it. The middle curve is the efficient policy under learning, while "naive" refers to the case in which policy presumes RE while agents are in fact learning with gain $\kappa = 0.05$.

Thus "perpetual learning" turns out to have major policy implications for policy, even when the deviation from REE might be thought not too large. The main policy implication is that with perpetual learning, there should be a policy bias towards "hawkishness." The intuition for this result is intuitive: a more hawkish (high θ) policy helps to keep inflation expectations π^e_{t+1} "in line," i.e. closer to RE values. This qualitative result also emerges in the more general setting in Orphanides and Williams (2007).

6 Estimated Models with Learning

The OW results suggest another implication of learning that goes beyond policy, namely that learning itself can be a source of persistence in macroeconomic dynamics. This line of thought has been pursued by Milani (2005), Milani (2007). The starting point is that inflation persistence in the data is much higher than arises from the basic New Keynesian model. For a good empirical fit to the data, a backward-looking component is needed in the New Keynesian Phillips curve under the RE assumption. The source of the backward-looking component used in these "hybrid" models, is, however, controversial. Milani (2005) considers the question of whether learning dynamics can provide some or all of the persistence needed to fit the data.

To investigate this, consider the most frequently used modification to the basic New Keynesian model, namely adding indexation to Calvo price setting; that is, firms that do not optimize in any given period set prices that are indexed to past inflation. This yields

$$\pi_t = \frac{\gamma}{1+\beta\gamma}\pi_{t-1} + \frac{\beta}{1+\beta\gamma}E_t^*\pi_{t+1} + \frac{\delta}{1+\beta\gamma}x_t + u_t.$$

where x_t is the output gap and γ measures the degree of indexation. Earlier work under RE empirically finds values of γ that are close to 1.

For expectations we assume a PLM of the form

$$\pi_t = \phi_0 + \phi_1\pi_{t-1} + \varepsilon_t,$$

and agents at t are assumed to use data $\{1, \pi_i\}_0^{t-1}$ to estimate ϕ_0, ϕ_1 using constant-gain least squares. For time t estimates $\phi_{0,t}, \phi_{1,t}$ the agents' forecasts are given by

$$\begin{aligned} E_t^*\pi_{t+1} &= \phi_{0,t} + \phi_{1,t}E_t^*\pi_t \\ &= \phi_{0,t} + \phi_{1,t}(\phi_{0,t} + \phi_{1,t}\pi_{t-1}), \end{aligned}$$

where we assume that the aggregate inflation rate π_t is not included in the agents information set at the time of their forecasts.

The implied ALM is

$$\pi_t = \frac{\beta\phi_{0,t}(1+\phi_{1,t})}{1+\beta\gamma} + \frac{\gamma+\beta\phi_{1,t}^2}{1+\beta\gamma}\pi_{t-1} + \frac{\delta}{1+\beta\gamma}x_t + u_t.$$

Alternatively, Milani (2005) also considers using real marginal cost as the driving variable in place of output gap x_t. To estimate the model for the US, Milani computes inflation from the GDP deflator and output gap as detrended GDP, while real marginal cost is proxied by deviation of labor income share from 1960:01 to 2003:04. Agents' initial parameter estimates are obtained by using pre-sample data 1951-1959.

A two-step procedure is used. First, the PLM is estimated from constant-gain learning using an assumed constant gain of $\kappa = 0.015$. This is in line with earlier empirical estimates. Then Milani estimates the ALM using nonlinear least squares. This procedure allows us to estimate the structural source of persistence, γ, taking into account the learning effects. The PLM parameter estimates show the following pattern:

(i) $\phi_{1,t}$ was initially low in 1950s and 60s, then higher (up to 0.958), and then declined somewhat to values above 0.8.

(ii) $\phi_{0,t}$ was initially low, then became much higher and then gradually declined after 1980.

The ALM structural estimates in particular generate a degree of indexation of $\gamma = 0.139$ (with the output gap). The results are fairly robust to other choices of gain κ that appear appropriate based on Schwartz' BIC model fit criterion. The estimate of γ is not significantly different from zero, and is in sharp contrast to the high levels of γ found under the RE assumption. Thus it appears that the data are consistent with the "learning" interpretation of the sources of persistence for inflation.

Milani (2007) estimates the full New Keynesian model under learning. He finds that also the degree of habit persistence is low in IS curve. This is in contrast with the usual extension of the New Keynesian model under RE that is often employed to improve the empirical fit of the model. Milani's work can be seen as a starting point for the very recent attempts by Slobodyan and Wouters (2007) and Murray (2007) to incorporate learning into DSGE models.

7 Recurrent Hyperinflations

The paper by Marcet and Nicolini (2003) starts from the standard hyperinflation model with learning and extends it to an open economy setting. The aim is to provide a unified theory to explain the recurrent hyperinflations experienced by many countries in the 1980s.

7.1 The Basic Hyperinflation Model

The starting point is the theoretical model sometimes known as the seigniorage model of inflation. (e.g. see Chapter 11 of Evans and Honkapohja (2001)). The Cagan model is based on the linear money demand equation

$$M_t^d/P_t = \phi - \phi\gamma(P_{t+1}^e/P_t) \text{ if } 1-\gamma(P_{t+1}^e/P_t) > 0 \text{ and 0 otherwise,}$$

which can be obtained from an overlapping generations (OG) endowment economy with log utility. This equation is combined with exogenous government purchases $d_t > 0$ that are entirely financed by seigniorage:

$$M_t = M_{t-1} + d_t P_t.$$

Rewriting this as $M_t/P_t = (M_{t-1}/P_{t-1})(P_{t-1}/P_t) + d$, setting $M_t^d = M_t$ and assuming $d_t = d$ we get

$$\frac{P_t}{P_{t-1}} = \frac{1 - \gamma(P_t^e/P_{t-1})}{1 - \gamma(P_{t+1}^e/P_t) - d/\phi}.$$

Under perfect foresight, i.e. $P_{t+1}^e/P_t = P_{t+1}/P_t$, there are two steady states, $\beta_L < \beta_H$, provided $d \geq 0$ is not too large, while if d is above a critical value then there are no perfect foresight steady states. There is also a continuum of perfect foresight paths converging to β_H. Some early theorists suggested that these paths might provide an explanation for actual hyperinflation episodes.

Consider now the situation under adaptive learning. Suppose the PLM is that the inflation process is perceived to be a steady state, i.e. $P_{t+1}/P_t = \beta + \eta_t$, where η_t is perceived white noise. Then PLM expectations are

$$\left(\frac{P_{t+1}}{P_t}\right)^e = \beta,$$

and the corresponding ALM is

$$\frac{P_t}{P_{t-1}} = \frac{1 - \gamma\beta}{1 - \gamma\beta - d/\phi} \equiv T(\beta; d).$$

Under steady-state learning, agents estimate β based on past average inflation, i.e. $(P_{t+1}/P_t)^e = \beta_t$ where

$$\beta_t = \beta_{t-1} + t^{-1}(P_{t-1}/P_{t-2} - \beta_{t-1}).$$

This is simply a recursive algorithm for the average inflation rate, which is equivalent to a least-squares regression on a constant.[24] It can be shown

[24] One can consider more general classes of PLM. Adam, Evans, and Honkapohja (2006) study the circumstances in which autoregressive PLMs can converge to hyperinflation paths.

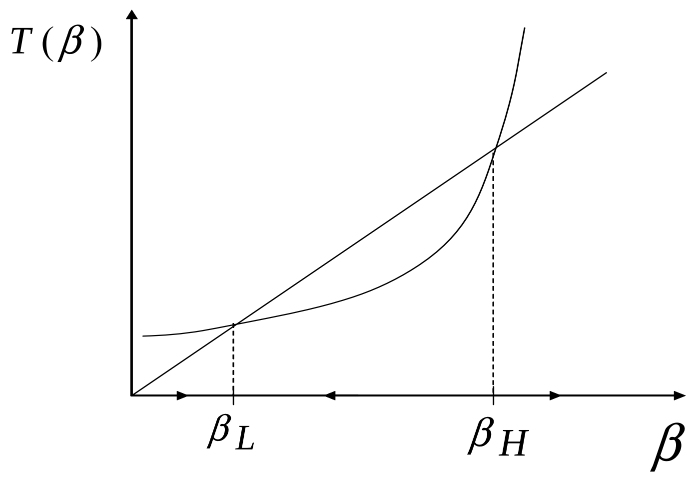

Figure 4: Figure 4: Steady state learning in the hyperinflation model

that stability of this learning rule is governed by the E-stability differential equation

$$d\beta/d\tau = T(\beta; d) - \beta,$$

where d is a fixed parameter. Since $0 < T'(\beta_L) < 1$ and $T'(\beta_H) > 1$, β_L is E-stable, and therefore locally stable under learning, while β_H is not. This is illustrated in Figure 4.

We remark that an increase in d shifts $T(\beta)$ up, so the comparative statics of β_L are natural but those of β_H are counterintuitive. This, together with the fact that the steady state β_H is not stable under learning suggests problems with the RE version of this model as a theoretical explanation for hyperinflations.

7.2 Empirical Background

Marcet and Nicolini (MN) list four stylized facts about hyperinflation episodes during the 1980s in a number of South American countries (as well as some episodes in other places and at other times):

1. Recurrence of hyperinflation episodes.

2. ERR (exchange rate rules) stop hyperinflations, though new hyperinflations eventually occur.

3. During a hyperinflation, seigniorage and inflation are not highly correlated.

4. Average inflation and seigniorage are strongly positively correlated across countries. Hyperinflations only occur in countries where seigniorage is on average high.

Stabilization plans to deal with hyperinflation have been based either on *heterodox* policy (ERR) or *orthodox* policy (permanently reducing the deficit). Policies that combine both elements appear to have been successful in stopping hyperinflations permanently.

7.3 The Marcet-Nicolini Model

MN use an open economy version of the overlapping-generations hyperinflation model. This is a flexible price model with PPP, so that

$$P_t^f e_t = P_t,$$

where P_t^f is the foreign price of goods, assumed exogenous. There is a cash-in-advance constraint for local currency on net purchases of consumption. This generates the demand by young agents for the local currency. Hence we continue to have the money demand equation as in the basic model. Government expenditure d_t is assumed to be *iid*.

There are two exchange rate regimes. In the floating regime the government does not buy or sell foreign exchange, and its budget constraint is as in the basic model. There is no foreign trade, and the economy behaves just like the closed economy model, with PPP determining the price of foreign currency by $e_t = P_t/P_t^f$.

In the ERR (exchange rate rule) regime, the government buys or sells foreign exchange R_t as needed to meet a target exchange rate e_t. Sales of foreign exchange generate revenue in addition to seigniorage that the government can use to finance government purchases, i.e. $(M_t - M_{t-1})/P_t = d_t + ((R_t - R_{t-1})e_t)/P_t$. In equilibrium, any increase in reserves must be matched by a trade surplus, i.e. $(R_t - R_{t-1})e_t = TB_t \cdot P_t$, where TB_t is total endowment minus total private consumption minus d_t.

The key question is the form of the ERR. When an ERR is adopted it is assumed that the object is to stabilize inflation at a targeted rate $\bar{\beta}$. This is accomplished by setting e_t to satisfy

$$\frac{P_t^f}{P_{t-1}^f}\frac{e_t}{e_{t-1}} = \bar{\beta},$$

which by PPP guarantees

$$\frac{P_t}{P_{t-1}} = \bar{\beta}.$$

Under ERR this last equation determines P_t. Given expectations, money demand determines M_t. Reserves R_t must then adjust to satisfy the flow government budget constraint.

The remaining question is how the government chooses exchange rate regimes. We assume there is a maximum inflation rate tolerated, β_U. ERR is imposed only in periods when inflation would otherwise exceed this bound (or if no positive P_t would otherwise clear the market).

7.4 Learning

MN argue that under RE the model cannot properly explain the stylized facts of hyperinflation outlined above. An adaptive learning formulation will be more successful. MN use a variation of the simple (decreasing gain) steady-state learning rule, given above, in which the gain is made state contingent:

$$\beta_t = \beta_{t-1} + \frac{1}{\alpha_t}\left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}\right),$$

with given β_0. Here $1/\alpha_t = \kappa_t$ is what we have called the gain, and $\alpha_t = \alpha_{t-1} + 1$ corresponds to decreasing gain learning, while $\alpha_t = \bar{\alpha} > 1$ is a constant-gain algorithm. (α_t can also be thought of as the "effective sample size"). MN consider a version in which agents switch between decreasing and constant gain according to recent performance. Specifically, $\alpha_t = \alpha_{t-1} + 1$ if $\left|\left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}\right)/\beta_{t-1}\right|$ falls below some bound υ and otherwise $\alpha_t = \bar{\alpha}$.

The qualitative features of the model are approximated by the system

$$\frac{P_t}{P_{t-1}} = h(\beta_{t-1}, d_t) \text{ where} \tag{24}$$

$$h(\beta, d) = \begin{cases} T(\beta; d) \text{ if } 0 < T(\beta; d) < \beta_U \\ \bar{\beta} \text{ otherwise} \end{cases}.$$

Figure 5 describes the dynamics of system (24).

There is a stable region consisting of values of β below the "unstable" high inflation steady state β_H and an unstable region that lies above it. Here we set $\bar{\beta} = \beta_L$, the low inflation steady state. β_U is set at a value above β_H.

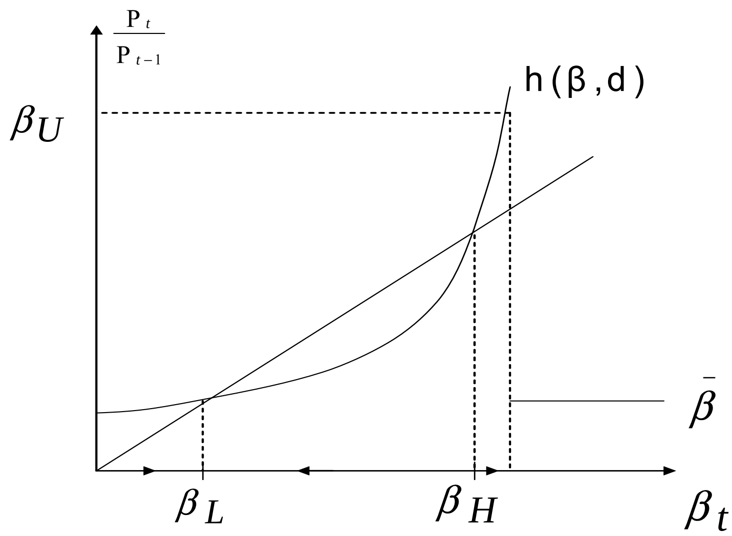

Figure 5: Figure 5: Inflation as a function of expected inflation

This gives rise to very natural recurring hyperinflation dynamics: Starting from β_L a sequence of random shocks may push β_t into the unstable region, at which point the gain is revised upward to $1/\bar{\alpha}$ and inflation follows an explosive path until it is stabilized by ERR. Then the process begins again. The model with learning has the following features:

(i) There may be eventual convergence to RE. This can occur if the random shocks/learning dynamics do not push β_t into the unstable region for a long time. Then decreasing gain may lead to asymptotic convergence to β_L.

(ii) A higher $E(d_t)$ makes average inflation higher and the frequency of hyperinflations greater. Orthodox combined with heterodox policies make sense as a way to end hyperinflations.

(iii) All four stylized facts listed above can be matched using this model, and simulations of a calibrated model look very plausible.

Overall this appears to be a very successful application of boundedly rational learning to a major empirical issue.

8 Liquidity Traps and Deflationary Spirals

Deflation and liquidity traps have been a concern in recent times. The paper by Evans, Guse, and Honkapohja (2007) considers issues of liquidity traps and deflationary spirals under learning in a New Keynesian model. As we have seen, contemporaneous Taylor-type interest-rate rules should respond to the inflation rate more than one for one in order to ensure determinacy and stability under learning. However, as emphasized by Benhabib, Schmitt-Grohe, and Uribe (2001), if one considers the interest-rate rule globally, not just in a neighborhood of the target inflation rate, the requirement that net nominal interest rates must be nonnegative implies that the rule must be nonlinear and also, for any continuous rule, that there exists a second steady state at a lower (possibly negative) inflation rate. This is illustrated in Figure 6, which shows the interest-rate policy $R = 1 + f(\pi)$ as a function of π.[25] The straight line in the figure is the Fisher equation $R = \pi/\beta$, which is obtained from the usual Euler equation for consumption in a steady state.

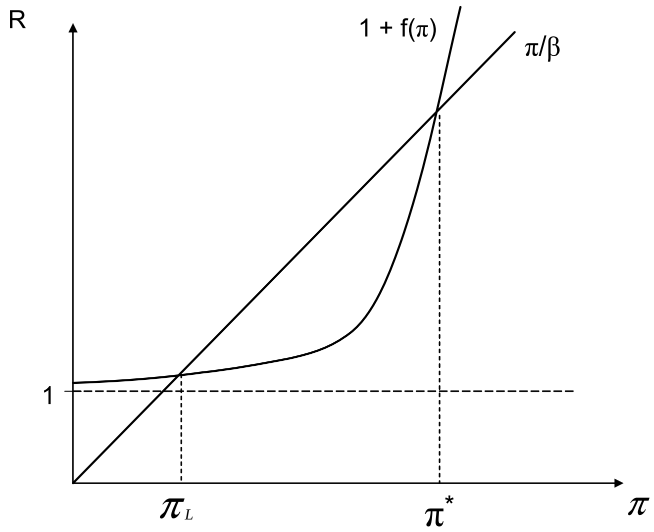

Figure 6: Multiple steady states with global Taylor rule

Here we are now using R to stand for the interest rate factor (so that the net interest rate is $R - 1$), and $\pi_t = P_t/P_{t-1}$ is the inflation factor, so that

[25] Of course, Taylor rules usually also include a dependence on aggregate output, which we omit for simplicity.

$\pi - 1$ is the net inflation rate. In the Figure π^* denotes the intended steady state, at which the "Taylor principle" of a more than one-for-one response is satisfied, and π_L is the unintended steady state. π_L may correspond to either a very low positive inflation rate or to a negative net inflation rate, i.e. deflation. The zero lower bound corresponds to $R = 1$. Benhabib, Schmitt-Grohe, and Uribe (2001) show that under RE, there is a continuum of "liquidity trap" paths that converge on π_L. The pure RE analysis thus suggests a serious risk of the economy following these "liquidity trap" paths.

What happens under learning? In Evans and Honkapohja (2005) we analyzed a flexible-price perfect competition model. We showed that deflationary paths are possible, but that the real risk, under learning, were paths in which inflation slipped below π_L and then continued to fall further. For this flexible-price model we showed that this could be avoided by a change in monetary policy at low inflation rates. The required policy is to switch to an aggressive money supply rule at some inflation rate between π_L and π^*. Such a policy would successfully avoid liquidity traps and deflationary paths.

Evans, Guse, and Honkapohja (2007) reconsider the issues in a model that allows for sticky prices and deviations of output from flexible-price levels. They consider a representative-agent infinite-horizon dynamic stochastic general equilibrium model with (i) monopolistic competition (ii) price-adjustment costs. Monetary policy follows a global Taylor-rule as above. Fiscal policy is standard: exogenous government purchases g_t and Ricardian tax policy that depends on real debt level. The model is essentially a New Keynesian model, except that, in line with Benhabib, Schmitt-Grohe, and Uribe (2001), it has Rotemberg (1982) costs of price adjustment as the friction rather than Calvo pricing. The model equations are nonlinear, and the nonlinearity in its analysis under learning is retained.

The key equations are

$$\begin{aligned} \frac{\alpha\gamma}{\nu}\left(\pi_t - 1\right)\pi_t &= \beta\frac{\alpha\gamma}{\nu}\left(\pi^e_{t+1} - 1\right)\pi^e_{t+1} \\ &\quad + (c_t + g_t)^{(1+\varepsilon)/\alpha} - \alpha\left(1 - \frac{1}{\nu}\right)(c_t + g_t)c_t^{-\sigma_1} \\ c_t &= c^e_{t+1}(\pi^e_{t+1}/\beta R_t)^{\sigma_1}, \end{aligned}$$

The first equation is the New Keynesian Phillips curve, relating π_t positively to π^e_{t+1} and to measures of aggregate activity. The second equation is the New Keynesian IS curve, obtained from the usual household Euler equation. When linearized around a steady state, both of these equations are identical in form to the standard New Keynesian equations. There are also money and debt evolution equations.

It is easily established that there are two stochastic steady states at π_L and π_H. If the random shocks are *iid* then "steady-state" learning is appropriate for both c^e and π^e, i.e.

$$\begin{aligned} \pi^e_{t+1} &= \pi^e_t + \phi_t(\pi_{t-1} - \pi^e_t) \\ c^e_{t+1} &= c^e_t + \phi_t(c_{t-1} - c^e_t), \end{aligned}$$

where ϕ_t is the gain sequence. The main findings are that while the intended steady state at π^* is locally stable under learning, the unintended steady state at π_L is unstable under learning. The key observation is that π_L is a saddlepoint, which implies the existence of deflationary spirals under learning. In particular, an expectational shock can lead to sufficiently pessimistic expectations, and c^e, π^e will follow paths leading to deflation and stagnation. This is illustrated in Figure 7, based on E-stability dynamics.

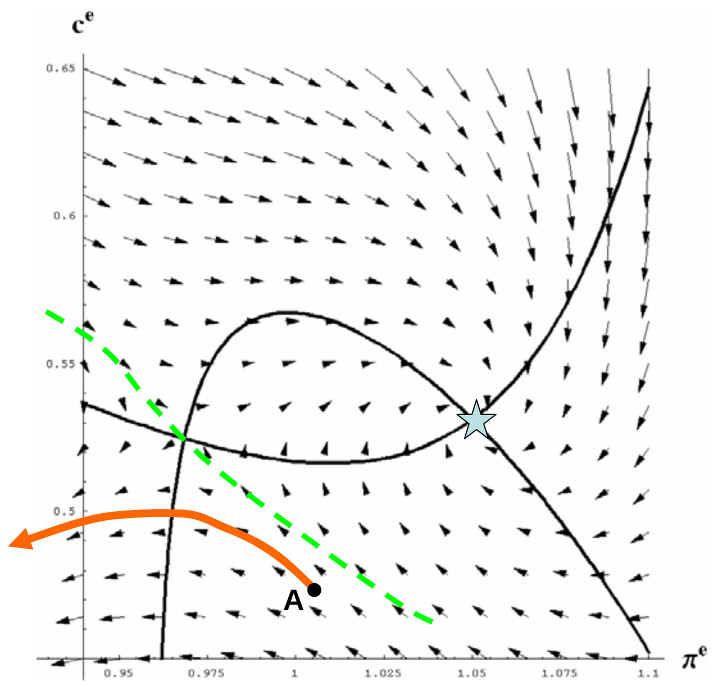

Figure 7: π^e and c^e dynamics under normal policy

The intuition for the result can be seen by supposing that we are initially near the π_L steady state and considering a small drop in π^e. With fixed R this would lead through the IS curve to lower c and thus, through the Phillips curve, to lower π. A sufficient reduction in R would be able to avoid the reductions in c and π , but since we are close to the zero lower bound this is not possible and the global Taylor rule here dictates only small reductions in R. The falls in realized c and π then under learning leads to reductions in c^e and π^e, and this sets in motion the deflationary spiral.

Thus, under normal policy the intended steady state is not globally stable under learning. Large adverse shocks to expectations or structural changes can set in motion unstable downward paths. Can policy be altered to avoid deflationary spiral? Evans, Guse, and Honkapohja (2007) show that it can. The recommended policy is to set a minimum inflation threshold $\tilde{\pi}$, where $\pi_L < \tilde{\pi} < \pi^*$. For example, if the global Taylor rule is chosen so that π_L corresponds to deflation, then a convenient choice for the threshold would be zero net inflation, i.e. $\tilde{\pi} = 1$. The authorities would follow normal monetary and fiscal policy provided this delivers $\pi_t > \tilde{\pi}$. However, if π_t threatens to fall below $\tilde{\pi}$ under normal policy, then aggressive policies would be implemented to ensure that $\pi_t = \tilde{\pi}$: interest rates would be reduced, if necessary to near the zero lower bound $R = 1$, and if this is not sufficient, then government purchases g_t would be increased as required.

Evans, Guse, and Honkapohja (2007) show that these policies can indeed ensure $\pi_t \geq \tilde{\pi}$ always under learning, and that incorporating aggressive monetary and fiscal policies triggered by an inflation threshold $\tilde{\pi}$ leads to global stability of the intended steady state at π^*. Perhaps surprisingly, it is also shown that it is essential to use an *inflation* threshold. They show that using instead an output threshold to trigger aggressive polices will not always avoid deflationary spirals.

9 Conclusions

Expectations play a large role in modern macroeconomics. While the RE assumption is the natural benchmark, it is implausibly demanding. Realistically, it should be assumed that people are smart, but boundedly rational. How should we model bounded rationality? We recommend the "principle of cognitive consistency": economic agents should be about as smart as (good) economists. Since when economists need to make forecasts, they do so using econometric models, a particularly natural choice is to model agents as econometricians.

In many economic models, with an appropriate econometric perceived law of motion, convergence to RE is possible. However, stability of REE under private agent learning is not automatic. Our central message is that monetary policy must be designed to ensure both determinacy and stability under learning. This observation leads to particular choices of interest-rate rules, whether we are considering standard classes of instrument rules or designing optimal monetary policy. Instrument rules that respond appropriately to "nowcasts" perform well in this respect, but implementing optimal policy appears to require an appropriate response to private sector expectations about the future.

More generally, policy-makers need to use policy to guide expectations, and the recent literature provides several important illustrations. If under learning there are persistent deviations from fully rational expectations, then monetary policy may need to respond more aggressively to inflation in order to stabilize expectations. The learning literature has also shown how to guide the economy under extreme threats of either hyperinflation or deflationary spirals. As we have illustrated, appropriate monetary and fiscal policy design can minimize these risks.

References

Adam, K., G. W. Evans, and S. Honkapohja (2006): "Are Hyperinflation Paths Learnable?," *Journal of Economic Dynamics and Control*, 30, 2725–2748.

Arifovic, J., J. Bullard, and O. Kostyshyna (2007): "Social Learning and Monetary Policy Rules," Working paper 2007-007a, Federal Reserve Bank of St Louis.

Bank of England, U. K. (2007): *Inflation Report*. Bank of England, www.bankofengland.co.uk/inflationreport/index.htm.

Basdevant, O. (2005): "Learning Processes and Rational Expectations: an Analysis Using a Small Macro-econometric Model for New Zealand," *Economic Modelling*, 22, 1074–1089.

Benhabib, J., S. Schmitt-Grohe, and M. Uribe (2001): "The Perils of Taylor Rules," *Journal of Economic Theory*, 96, 40–69.

Berardi, M., and J. Duffy (2007): "The Value of Central Bank Transparency When Agents are Learning," *European Journal of Political Economy*, 23, 9–29.

Bernanke, B. (2007): "Inflation Expectations and Inflation Forecasting," Remarks at NBER Summer Institute, Cambridge, MA, http://www.federalreserve.gov/boarddocs/speeches/2007/.

Bernanke, B., and M. Woodford (1997): "Inflation Forecasts and Monetary Policy," *Journal of Money, Credit, and Banking*, 24, 653–684.

——— (eds.) (2005): *The Inflation-Targeting Debate*. University of Chicago Press, Chicago, USA.

Branch, W., J. Carlson, G. W. Evans, and B. McGough (2006a): "Adaptive Learning, Endogenous Inattention and Changes in Monetary Policy," Working paper, University of Oregon.

——— (2006b): "Monetary Policy, Endogenous Inattention and the Volatility Trade-Off," *Economic Journal,* forthcoming.

Branch, W., and G. W. Evans (2006a): "Intrinsic Heterogeneity in Expectation Formation," *Journal of Economic Theory*, 127, 264–295.

——— (2006b): "A Simple Recursive Forecasting Model," *Economic Letters*, 91, 158–166.

——— (2007): "Model Uncertainty and Endogenous Volatility," *Review of Economic Dynamics*, 10, 207–237.

Branch, W. A. (2004): "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *Economic Journal*, 114, 592–621.

Brazier, A., R. Harrison, M. King, and T. Yates (2006): "The Danger of Inflating Expectations of Macroeconomic Stability: Heuristic Switching in an Overlapping Generations Monetary Model," Working paper nr.303, Bank of England.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65, 1059–1095.

Brunner, K., and A. Meltzer (eds.) (1976): *The Phillips Curve and the Labor Market. Carnegie-Rochester Conferences in Public Policy, Vol. 1.* North-Holland, Amsterdam.

Bullard, J. (2006): "The Learnability Criterion and Monetary Policy," *Federal Reserve Bank of St. Louis Review*, 88, 203–217.

Bullard, J., and I.-K. Cho (2005): "Escapist Policy Rules," *Journal of Economic Dynamics and Control*, 29, 1841–1866.

Bullard, J., and S. Eusepi (2005): "Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?," *Review of Economic Dynamics*, 8, 324–359.

Bullard, J., G. W. Evans, and S. Honkapohja (2007): "Monetary Policy, Judgement and Near-Rational Exuberance," *American Economic Review, forthcoming.*

Bullard, J., and K. Mitra (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49, 1105–1129.

——— (2007): "Determinacy, Learnability and Monetary Policy Inertia," *Journal of Money Credit and Banking*, 39, 1177–1212.

Bullard, J., and E. Schaling (2006): "Monetary Policy, Determinacy, Learnability in a Two-Block World Economy," Working paper 2006-038a, Federal Reserve Bank of St Louis, www.stls.frb.org/research/econ/bullard.

Bullard, J., and A. Singh (2006): "Worldwide Macroeconomic Stability and Monetary Policy Rules," Working paper 2006-040b, Federal Reserve Bank of St Louis, www.stls.frb.org/research/econ/bullard.

Calvo, G. A. (1983): "Staggered Pricing in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

Carboni, G., and M. Ellison (2007): "The Great Inflation and the Greenbook," mimeo.

Carlstrom, C. T., and T. S. Fuerst (2004): "Learning and the Central Bank," *Journal of Monetary Economics*, 51, 327–338.

Clarida, R., J. Gali, and M. Gertler (1998): "Monetary Policy Rules in Practice: Some International Evidence," *European Economic Review*, 42, 1033–1067.

——— (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Cogley, T., and T. J. Sargent (2005): "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty," *Review of Economic Dynamics*, 8, 528–563.

Duffy, J., and W. Xiao (2007a): "Investment and Monetary Policy: Learning and Determinacy of Equilibria," Working paper.

——— (2007b): "The Value of Interest Rate Stabilization Policies When Agents are Learning," *Journal of Money, Credit and Banking,* forthcoming.

Ellison, M., and T. Yates (2007): "Escaping Volatile Inflation," *Journal of Money, Credit and Banking*, 39, 981–993.

European Central Bank, E. U. (2007): *Monthly Bulletin.* European Central Bank, www.ecb.int.

Eusepi, S. (2007): "Learnability and Monetary Policy: A Global Perspective," *Journal of Monetary Economics*, 54, 1115–1131.

Eusepi, S., and B. Preston (2007): "Central Bank Communication and Expectations Stabilization," Working paper nr.13259, NBER.

Evans, G. W., E. Guse, and S. Honkapohja (2007): "Liquidity Traps, Learning and Stagnation," Working paper nr.732, University of Cambridge.

Evans, G. W., and S. Honkapohja (1993): "Adaptive Forecasts, Hysteresis and Endogenous Fluctuations," *Federal Reserve Bank of San Francisco Economic Review*, 1993(1), 3–13.

——— (1995): "Adaptive Learning and Expectational Stability: An Introduction," in Kirman and Salmon (1995), chap. 4, pp. 102–126.

——— (1999): "Learning Dynamics," in Taylor and Woodford (1999), chap. 7, pp. 449–542.

——— (2001): *Learning and Expectations in Macroeconomics.* Princeton University Press, Princeton, New Jersey.

——— (2003a): "Adaptive Learning and Monetary Policy Design," *Journal of Money, Credit and Banking*, 35, 1045–1072.

——— (2003b): "Expectational Stability of Stationary Sunspot Equilibria in a Forward-looking Model," *Journal of Economic Dynamics and Control*, 28, 171–181.

——— (2003c): "Expectations and the Stability Problem for Optimal Monetary Policies," *Review of Economic Studies*, 70, 807–824.

——— (2003d): "Friedman's Money Supply Rule versus Optimal Interest Rate Policy," *Scottish Journal of Political Economy*, 50, 550–566.

——— (2005): "Policy Interaction, Expectations and the Liquidity Trap," *Review of Economic Dynamics*, 8, 303–323.

——— (2006): "Monetary Policy, Expectations and Commitment," *Scandinavian Journal of Economics*, 108, 15–38.

——— (2007a): "Policy Interaction, Learning and the Fiscal Theory of Prices," *Macroeconomic Dynamics*, 11, 665–690.

——— (2007b): "Robust Learning Stability with Operational Monetary Policy Rules," mimeo.

Evans, G. W., S. Honkapohja, and R. Marimon (2001): "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules," *Macroeconomic Dynamics*, 5, 1–31.

——— (2007): "Stable Sunspot Equilibria in a Cash-in-Advance Economy," *The B.E. Journal of Macroeconomics* (Advances), 7, Iss. 1, Article 3.

Evans, G. W., and B. McGough (2005a): "Monetary Policy, Indeterminacy and Learning," *Journal of Economic Dynamics and Control*, 29, 1809–1840.

——— (2005b): "Stable Sunspot Solutions in Models with Predetermined Variables," *Journal of Economic Dynamics and Control*, 29, 601–625.

——— (2007): "Optimal Constrained Interest-Rate Rules," *Journal of Money, Credit and Banking*, 39, 1335–1356.

Evans, G. W., and G. Ramey (2006): "Adaptive Expectations, Underparameterization and the Lucas Critique," *Journal of Monetary Economics*, 53, 249–264.

Federal Reserve Bank, K. C. (1999): *New Challenges for Monetary Policy*. Federal Reserve Bank of Kansas City, Kansas City.

Ferrero, G. (2007): "Monetary Policy, Learning and the Speed of Convergence," *Journal of Economic Dynamics and Control*, 31, 3006–3041.

Gaspar, V., F. Smets, and D. Vestin (2005): "Optimal Monetary Policy Under Adaptive Learning," mimeo.

——— (2006): "Adaptive Learning, Persistence and Optimal Monetary Policy," *Journal of the European Economic Association*, 4, 376–385.

GIANNITSAROU, C. (2003): “Heterogenous Learning,” *Review of Economic Dynamics*, 6, 885–906.

HONKAPOHJA, S., AND K. MITRA (2003): “Learning with Bounded Memory in Stochastic Models,” *Journal of Economic Dynamics and Control*, 27, 1437–1457.

——— (2004): “Are Non-Fundamental Equilibria Learnable in Models of Monetary Policy?,” *Journal of Monetary Economics*, 51, 1743–1770.

——— (2005a): “Performance of Inflation Targeting Based on Constant Interest Rate Projections,” *Journal of Economic Dynamics and Control*, 29, 1867–1892.

——— (2005b): “Performance of Monetary Policy with Internal Central Bank Forecasting,” *Journal of Economic Dynamics and Control*, 29, 627–658.

——— (2006): “Learning Stability in Economies with Heterogeneous Agents,” *Review of Economic Dynamics*, 9, 284–309.

KIRMAN, A., AND M. SALMON (eds.) (1995): *Learning and Rationality in Economics*. Basil Blackwell, Oxford.

KREPS, D., AND K. WALLIS (eds.) (1997): *Advances in Economics and Econometrics: Theory and Applications, Volume I.* Cambridge University Press, Cambridge.

KUROZUMI, T. (2006): “Determinacy and Expectational Stability of Equilibrium in a Monetary Sticky-price Model with Taylor Rule,” *Journal of Monetary Economics*, 53, 827–846.

KUROZUMI, T., AND W. VAN ZANDWEGHE (2007): “Investment, Interest Rate Policy and Equilibrium Stability,” *Journal of Economic Dynamics and Control,* forthcoming.

LLOSA, G., AND V. TUESTA (2006): “Determinacy and Learnability of Monetary Policy Rules in Small Open Economies,” Working paper nr.576, I-ADB.

——— (2007): “E-Stability of Monetary Policy When the Cost Channel Matters,” mimeo.

Lucas, Jr., R. E. (1976): "Econometric Policy Evaluation: a Critique," in Brunner and Meltzer (1976), pp. 19–46.

Marcet, A., and J. P. Nicolini (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93, 1476–1498.

Marimon, R. (1997): "Learning from Learning in Economics," in Kreps and Wallis (1997), chap. 9, pp. 278–315.

McCallum, B. T. (1999): "Issues in the Design of Monetary Policy Rules," in Taylor and Woodford (1999), chap. 23, pp. 1483–1530.

——— (2003): "Is the Fiscal Theory of the Price Level Learnable?," *Scottish Journal of Political Economy*, 50, 634–649.

——— (2007): "E-stability vis-à-vis Determinacy Results for a Broad Class of Linear Rational Expectations Models," *Journal of Economic Dynamics and Control*, 31, 1376–1391.

McCallum, B. T., and E. Nelson (1999): "Performance of Operational Policy Rules in an Estimated Semi-Classical Model," in Taylor (1999), pp. 15–45.

——— (2004): "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models," *Federal Reserve Bank of St.Louis Review*, 86, 43–56.

Milani, F. (2005): "Adaptive Learning and Inflation Persistence," Working paper, University of California, Irvine.

——— (2007): "Expectations, Learning and Macroeconomic Persistence," *Journal of Monetary Economics*, 54, 2065–2082.

Molnar, K., and S. Santoro (2006): "Optimal Monetary Policy When Agents are Learning," Mt-dp working paper 2006/1.

Murray, J. (2007): "Empirical Significance of Learning in New Keynesian Model with Firm-Specific Capital," mimeo.

Orphanides, A., and J. C. Williams (2005a): "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations," *Journal of Economic Dynamics and Control*, 29, 1927–1950.

——— (2005b): "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in Bernanke and Woodford (2005), chap. 5, pp. 201–234.

——— (2005c): "Inflation Scares and Forecast-Based Monetary Policy," *Review of Economic Dynamics*, 8, 498–527.

——— (2007): "Robust Monetary Policy with Imperfect Knowledge," *Journal of Monetary Economics*, 54, 1406–1435.

Pfajfar, D. (2007): "Formation of Rationally Heterogeneous Expectations," mimeo.

Pfajfar, D., and E. Santoro (2007a): "Credit Market Distortions, Asset Prices and Monetary Policy," mimeo.

——— (2007b): "Heterogeneity and Learning in Inflation Expectation Formation: An Empirical Assessment," mimeo.

Preston, B. (2005): "Learning about Monetary Policy Rules when Long-Horizon Expectations Matter," *International Journal of Central Banking*, 1, 81–126.

——— (2006): "Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy," *Journal of Monetary Economics*, 53, 507–535.

Rotemberg, J. J. (1982): "Sticky Prices in the United States," *Journal of Political Economy*, 90, 1187–1211.

Rotemberg, J. J., and M. Woodford (1999): "Interest Rate rules in an Estimated Sticky-Price Model," in Taylor (1999), pp. 57–119.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.

——— (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.

Sargent, T. J., and N. Wallace (1975): "'Rational Expectations', the Optimal Monetary Instrument and the Optimal Money Supply Rule," *Journal of Political Economy*, 83, 241–254.

Sargent, T. J., N. Williams, and T. Zha (2006): "Shocks and Government Beliefs: The Rise and Fall of American Inflation," *American Economic Review*, 96, 1193–1224.

Slobodyan, S., and R. Wouters (2007): "Learning in an Estimated DSGE Model," mimeo.

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

——— (2003): "What is Wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules," *Journal of Economic Literature*, 41, 426–477.

Svensson, L. E. O., and M. Woodford (2005): "Implementing Optimal Policy through Inflation-Forecast Targeting," in Bernanke and Woodford (2005), pp. 19–83.

Taylor, J. (1993): "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series in Public Policy*, 39, 195–214.

——— (ed.) (1999): *Monetary Policy Rules*. University of Chicago Press, Chicago.

Taylor, J., and M. Woodford (eds.) (1999): *Handbook of Macroeconomics, Volume 1*. Elsevier, Amsterdam.

Trichet, J.-C. (2005): "Monetary Policy and Private Expectations," Zolotas Lecture at the Bank of Greece, Athens, February, http://www.ecb.int/press/key/date/2005/html/.

Walsh, C. (2003): *Monetary Theory and Policy*. MIT Press, Cambridge, Massachusetts, second edn.

Wang, Q. (2006): "Learning Stability for Monetary Policy Rules in a Two-Country Model," Working paper nr.659, University of Cambridge.

Woodford, M. (1999a): "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?," in Federal Reserve Bank (1999), pp. 277–316.

——— (1999b): "Optimal Monetary Policy Inertia," *The Manchester School, Supplement*, 67, 1–35.

——— (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.

——— (2005): "Robustly Optimal Monetary Policy with Near-Rational Expectations," mimeo.

ZANNA, L.-F. (2006): "PPP Rules, Macroeconomic (In)stability and Learning," mimeo.