Expectations, Learning, and Monetary Policy: An Overview of Recent Research

George W. Evans
University of Oregon and University of St Andrews

Seppo Honkapohja
Bank of Finland and University of Cambridge

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 to forecasts of future inflation and output, and the question of whether monetary policy should be forward-looking has been subject to discussion and debate. The Bank of England and the European Central Bank include private sector forecasts and internal macroeconomic projections in their periodic reports (Bank of England, 2007; European Central Bank, 2007). Empirical evidence on Germany, Japan, and the United States since 1979 similarly suggests that central banks are forward-looking in practice (Clarida, Galí, and Gertler, 1998).

The rational expectations 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 rational expectations hypothesis is based on the assumption that both private agents and the policymaker know the true model of the economy, except for unforecastable random

1. Woodford (2003) is a monumental treatise on the subject, while Walsh (2003) provides an accessible graduate-level treatment. For surveys, see Clarida, Galí, and Gertler (1999) and McCallum (1999).

shocks.[2] The rational expectations assumption is excessively strong: neither private agents nor policymakers 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 reestimated and possibly reformulated as new data become 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 policymaking 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 policymakers in the economy behave like applied economists and econometricians. It is thus 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 learning processes 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; Bernanke, 2007). This research shows that interest rate setting by monetary policymakers faces two fundamental problems. First, some of the proposed interest rate rules may not perform well when agents' expectations are out of equilibrium. The consequences of errors in forecasting, and the resulting correction mechanisms, may create instability in the economy. For (usually suboptimal) instrument rules, Bullard and Mitra (2002) consider the stability of the rational expectations equilibrium when monetary policy is conducted using variants of the Taylor rule. These rules work well only under certain parameter restrictions, and Bullard and Mitra suggest that monetary policymaking should take into account the learnability constraints on the parameters of policy behavior. For

---

2. Some papers do extend the standard notion of rational expectations equilibrium to an equilibrium with limited information. These extensions often assume that economic agents do not observe some variables but 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 (1995, 1999), Marimon (1997), and Sargent (1993, 1999) provide surveys of the field.

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

optimal monetary policy, Evans and Honkapohja (2003c, 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 the output gap, can create multiple equilibria, also called indeterminacy of equilibria.[5] Under indeterminacy there are multiple, even continua of rational expectations equilibria and the economy need not settle on the desired equilibrium. The possible rest points have been studied using stability under learning as a selection criterion (see Honkapohja and Mitra, 2004; Carlstrom and Fuerst, 2004; Evans and McGough, 2005a). Indeterminacy is not a critical problem if the fundamental rational expectations equilibrium is the only stable equilibrium under learning. Moreover, indeterminacy need not arise if the forward-looking interest rate rule is carefully designed, as shown by Bullard and Mitra (2002) and Evans and Honkapohja (2003c, 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 nontechnical overview of this research program. The first part of the paper reviews the basic theoretical results. We then take up some immediate practical concerns that can arise in connection with rules for interest rate setting, including issues of observability in connection with private forecasts and with current output and inflation data. A second concern is the knowledge of the structure of the economy that is required to implement optimal interest rate policies. The second part of the paper provides an overview of the recent and ongoing developments in the literature. We first summarize research on learnability of rational expectations equilibria when the basic New-Keynesian model is

---

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).

extended to incorporate further features of the economy. We then 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.

1. THE MODEL

We conduct our discussion using the New-Keynesian model that has become the workhorse in the analysis of monetary policy, and we directly employ 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 firms' constraints on the frequency of price changes, as originally suggested by Calvo (1983).

The behavior of the private sector is summarized by 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, and i_t is the nominal interest rate, expressed as the deviation from the steady state real interest rate. The determination of i_t is discussed below. Private sector expectations of the output gap and inflation in the next period are denoted $E_t^* x_{t+1}$ and $E_t^* \pi_{t+1}$, respectively. Since our focus is on learning behavior, these expectations need not be rational (E_t without * denotes rational expectations). 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), which 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 in the next period.

Rules for forecasting the next period's values of the state variables are the other ingredient in the description of individual behavior. We assume 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 to follow

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

where

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

$0 < |\mu| < 1$, $0 < |\rho| < 1$, and $\tilde{g}_t \sim$ i.i.d. $(0, \sigma_g^2)$, $\tilde{u}_t \sim$ i.i.d. $(0, \sigma_u^2)$ are independent white noise. In addition, g_t represents shocks to government purchases or potential output (or both), and 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 the central bank's interest rate setting.[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, that is,

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

where $\overline{\pi}$ 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 percent).

---

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. Preston (2005, 2006) proposes a formulation in which long horizons matter in individual behavior.

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 that 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 (2007), McCallum (2003), and Evans, Guse, and Honkapohja (2008). We discuss some aspects of the interaction below.

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 take the inflation target to be $\bar{\pi} = 0$, so that this class of rules takes the form

$$i_t = \chi_\pi \pi_t + \chi_x x_t, \tag{4}$$

where χ_π, $\chi_x > 0$. 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, that is,

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

where $\bar{\pi}$ is the inflation target. This type of optimal policy is often called flexible inflation targeting in the current literature (see, for example, Svensson, 1999, 2003). The policymaker is assumed to have the same discount factor, β, as the private sector, while α is the relative weight placed by the policymaker on the output target. The case of $\alpha = 0$ represents strict inflation targeting. The loss function (5) can alternatively be viewed as a quadratic approximation to the welfare function of a representative agent.[8]

The literature on optimal policy under rational expectations distinguishes between optimal discretionary policy, in which the policymaker is unable to commit to policies for future periods, and optimal policy in which such commitment is possible. Under commitment, the policymaker can do better because of the effect on private expectations, but commitment policy exhibits time inconsistency, in the sense that policymakers 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.

---

8. See Rotemberg and Woodford, 1999; Woodford, 2003. In this formulation, α is a function of various deep structural parameters in the fully microfounded version of the model.

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 described in equations (1) and (2), the optimality condition (6), and rational expectations. However, some 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 and u_t and 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 to neutralize the effects of aggregate demand shocks, g_t, and to strike the optimal balance between output and inflation effects for inflation shocks, u_t. The dependence of i_t on x_{t-1} is optimally chosen to take advantage of the effects on expectations of commitment to a rule.[9]

Expectations-based optimal rules are advocated in Evans and Honkapohja (2003c, 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. When this is done appropriately, the rational expectations equilibrium will be stable under learning and optimal policy can thus be successfully implemented. The essence of these rules is that they do not assume rational expectations 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 rational expectations equilibrium under learning. (If expectations are rational, these rules deliver the optimal equilibrium.)

The optimal expectations-based rule under commitment is

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

---

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 $\zeta = 1 + \beta + \lambda^2/\alpha$, and $\bar{b}_\pi = (\alpha/\lambda)(1 - \bar{b}_x)$, $\bar{c}_x = -[\lambda + \beta\bar{b}_\pi + (1 - \beta\rho)(\alpha/\lambda)]^{-1}$, $\bar{c}_\pi = -(\alpha/\lambda)\bar{c}_x$.

The coefficients of equation (8) are

$$\delta_L = \frac{-\alpha}{\varphi(\alpha + \lambda^2)}, \ \delta_\pi = 1 + \frac{\lambda\beta}{\varphi(\alpha + \lambda^2)},$$

$$\delta_x = \varphi^{-1}, \ \delta_g = \varphi^{-1} \ \text{and} \ \delta_u = \frac{\lambda}{\varphi(\alpha + \lambda^2)}. \tag{9}$$

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

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

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

where $\theta > 0$. This rule combines the fundamentals-based rule of equation (7) with the correction for the first-order condition.[11] Rule (10) delivers the optimal equilibrium under rational expectations. McCallum and Nelson (2004) suggest another hybrid rule, which takes the form

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

where $\theta > 0$.

2. DETERMINACY AND STABILITY UNDER LEARNING

Given an interest rate rule, we can obtain the reduced form of the model and study its properties under rational expectations. Two basic properties of interest are determinacy of the rational expectations

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).

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

solution and stability under learning of the rational expectations equilibrium.

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

$$\mathbf{y}_t = \begin{pmatrix} x_t \\ \pi_t \end{pmatrix} \text{ and } \mathbf{v}_t = \begin{pmatrix} g_t \\ u_t \end{pmatrix},$$

the reduced form can be written as

$$\mathbf{y}_t = \mathbf{M}E_t^* \mathbf{y}_{t+1} + \mathbf{N}\mathbf{y}_{t-1} + \mathbf{P}\mathbf{v}_t, \tag{12}$$

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

$$\mathbf{y}_t = \mathbf{M}E_t^* \mathbf{y}_{t+1} + \mathbf{P}\mathbf{v}_t. \tag{13}$$

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

The first issue of concern is whether under rational expectations the system possesses a unique stationary equilibrium, in which case the model is said to be determinate. The model is said to be indeterminate if it has multiple stationary solutions. These multiple solutions include sunspot solutions, in which the rational expectations equilibrium depends on extraneous random variables that influence the economy solely through agents' expectations.[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 agents' beliefs. These beliefs concern the stochastic process followed by the endogenous (and exogenous) variables that need to be forecasted. The parameters

---

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).

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 rational expectations solution of interest be stable—that is, will it be reached asymptotically by the learning process? If not, the rational expectations equilibrium is unlikely to be attained. This is the focus of the papers by Bullard and Mitra (2002, 2007), Evans and Honkapohja (2003c, 2006), and many others.

2.1 Digression on Methodology

Consider first the simpler reduced-form equation (13) under rational expectations. For determinacy to hold, both eigenvalues of the 2×2 matrix $\mathbf{M}$ must lie inside the unit circle. In the determinate case, the unique stationary solution will be of the minimal state variable (MSV) form:

$$\mathbf{y}_t = \bar{\mathbf{c}}\mathbf{v}_t,$$

where $\bar{\mathbf{c}}$ 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.

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

$$\mathbf{y}_t = \mathbf{a} + \mathbf{c}\mathbf{v}_t, \tag{14}$$

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

---

13. Private agents and the policymaker are here assumed to observe the shocks $\mathbf{v}_t$. If $\mathbf{v}_t$ is not observable then the PLM would be adjusted to reflect relevant available information.

With this PLM and parameter estimates ($\mathbf{a}$, $\mathbf{c}$), agents would form expectations as

$$E_t^* \mathbf{y}_{t+1} = \mathbf{a} + \mathbf{cF}\mathbf{v}_t,$$

where $\mathbf{F}$ is either known or also estimated. Inserting these expectations into eqaution (13) and solving for $\mathbf{y}_t$, we get the implied actual law of motion (ALM), that is, the law that $\mathbf{y}_t$ would follow for a fixed PLM ($\mathbf{a}$, $\mathbf{c}$).[14] This is given by

$$\mathbf{y}_t = \mathbf{Ma} + (\mathbf{P} + \mathbf{McF})\mathbf{v}_t.$$

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

$$\mathbf{T}(\mathbf{a}, \mathbf{c}) = (\mathbf{Ma}, \mathbf{P} + \mathbf{McF})\mathbf{v}_t,$$

and the rational expectations solution ($\mathbf{0}$, $\bar{\mathbf{c}}$) 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 ($\mathbf{a}_t$, $\mathbf{c}_t$) of the PLM parameters computed on the basis of data through $t - 1$. Next, exogenous shocks $\mathbf{v}_t$ are realized, and private agents form expectations $E_t^* \mathbf{y}_{t+1} = \mathbf{a}_t + \mathbf{c}_t\mathbf{F}\mathbf{v}_t$ (assuming for convenience that $\mathbf{F}$ is known). Following, for example, policy rule (4), the central bank sets the interest rate i_t, and $\mathbf{y}_t$ is generated according to equations (1) and (2) together with the interest rate rule. This temporary equilibrium is summarized by equation (13). At the beginning of $t + 1$ agents add the new data point to their information set to update their parameter estimates to ($\mathbf{a}_{t+1}$, $\mathbf{c}_{t+1}$) using least squares, for example, and the process continues. The question of interest is whether ($\mathbf{a}_t$, $\mathbf{c}_t$) $\rightarrow$ ($\mathbf{0}$, $\bar{\mathbf{c}}$) 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

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

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 Evans and Honkapohja (2001, chap. 10).

where τ denotes notional time. If the rational expectations equilibrium $(\mathbf{0}, \bar{\mathbf{c}})$ is locally asymptotically stable under this differential equation, then the equilibrium 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 also refer to these stability conditions as the conditions for stability under adaptive learning, the conditions for stability under learning, or the conditions for learnability of the equilibrium.

For the reduced-form equation (13), it can be shown that the two E-stability conditions are that the eigenvalues of $\mathbf{M}$ have real parts less than one and that all products of eigenvalues of $\mathbf{M}$ times eigenvalues of $\mathbf{F}$ 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 equation (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

$$\mathbf{y}_t = \mathbf{a} + \mathbf{b}\mathbf{y}_{t-1} + \mathbf{c}\mathbf{v}_t, \tag{15}$$

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

To examine stability under learning, we treat equation (15) as the agents' PLM. Under real-time learning, agents estimate the coefficients $\mathbf{a}$, $\mathbf{b}$, $\mathbf{c}$ of equation (15). This is a vector autoregression (VAR) with exogenous variables $\mathbf{v}_t$. The estimates $(\mathbf{a}_t, \mathbf{b}_t, \mathbf{c}_t)$ are updated at each point in time by recursive least squares. Once again it can be shown that the E-stability principle gives the conditions for local convergence of real-time learning.

---

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 nonpredetermined variables with the number of roots of the forward-looking matrix that lie inside the unit circle.

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

$$E_t^* \mathbf{y}_{t+1} = \mathbf{a} + \mathbf{b}(\mathbf{a} + \mathbf{b}\mathbf{y}_{t-1} + \mathbf{c}\mathbf{v}_t) + \mathbf{c}\mathbf{F}\mathbf{v}_t, \tag{16}$$

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

$$\mathbf{T}(\mathbf{a}, \mathbf{b}, \mathbf{c}) = [\mathbf{M}(\mathbf{I} + \mathbf{b})\mathbf{a}, \ \mathbf{M}\mathbf{b}^2 + \mathbf{N}, \ \mathbf{M}(\mathbf{b}\mathbf{c} + \mathbf{c}\mathbf{F}) + \mathbf{P}], \tag{17}$$

E-stability is again determined by the differential equation

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

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

2.2 Results for Monetary Policy

We now describe the determinacy and stability results for the interest rate rules described in section 1.

2.2.1 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 χ_π, $\chi_x \geq 0$, they assume that the serial correlation parameters in $\mathbf{F}$ are nonnegative. The results are particularly straightforward and natural for policy rule (4).[18] Bullard and Mitra (2002) show that the rational expectations equilibrium is determinate and stable under learning if and only if (using our notation)

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

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

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

The situation is more complicated if lagged or forward-looking Taylor rules are used, and full analytical results are not available. For the lagged variable case, they find that for $\chi_\pi > 1$ and a sufficiently small $\chi_x > 0$, the policy leads to a rational expectations equilibrium 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

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

where χ_π, $\chi_x > 0$ and 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 policymakers. They find that for $\chi_\pi > 1$ and a sufficiently small $\chi_x > 0$, the policy leads to a rational expectations equilibrium that is determinate and stable under learning. Now for $\chi_\pi > 1$ and a large χ_x, the system is indeterminate, yet the MSV solution is stable under learning. E-stable sunspot equilibria are also possible, however, as 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.[19] 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 policymaker must appropriately select the parameters of the policy rule, χ_π, χ_x, when an instrument rule describes policy. Stability under learning provides a constraint for this choice.

---

19. Bullard and Mitra (2007) extend their analysis to include interest rate inertia, while Kurozumi (2006) considers modifications to the determinacy and E-stability results when the model structure is varied. Mitra (2003) examines performance of the related case of nominal income targeting.

2.2.2 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 this rule yields regions in which the optimal rational expectations equilibrium is determinate.[20] The basic intuition for this result can be seen from the following reduced-form equation:
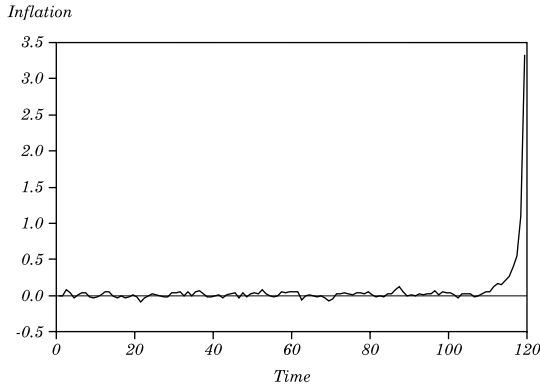
$$\begin{pmatrix} x_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta + \lambda\varphi \end{pmatrix} \begin{pmatrix} E_t^* x_{t+1} \\ E_t^* \pi_{t+1} \end{pmatrix} + \begin{pmatrix} -\varphi\psi_x & 0 \\ -\lambda\varphi\psi_x & 0 \end{pmatrix} \begin{pmatrix} x_{t-1} \\ \pi_{t-1} \end{pmatrix} + \begin{pmatrix} -\varphi\psi_u \\ 1 - \lambda\varphi\psi_u \end{pmatrix} u_t. \qquad (20)$$

Since typically $\beta + \lambda\varphi > 1$, 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 the rational expectations equilibrium. 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 rational expectations equilibrium occurs under rule (7).[21] The instability of the fundamentals-based rules, which are designed to obtain optimal policy, serves as a strong warning to policymakers not to automatically assume that rational expectations will be attained. It is necessary to examine explicitly the robustness of contemplated policy rules to private agent learning.

In Evans and Honkapohja (2003c, 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 partly conditioned on these expectations. In Evans and Honkapohja (2006), we show that under rule (8), the economy is determinate and the optimal rational expectations equilibrium 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,

---

20. The learning stability results are sensitive to the detailed information assumptions. With the PLM equation (15), if agents can make forecasts conditional also on $\mathbf{y}_t$, then there are regions of both stability and instability under the fundamentals-based rule, depending on the structural parameters.

21. 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 a Fundamentals-Based Rule
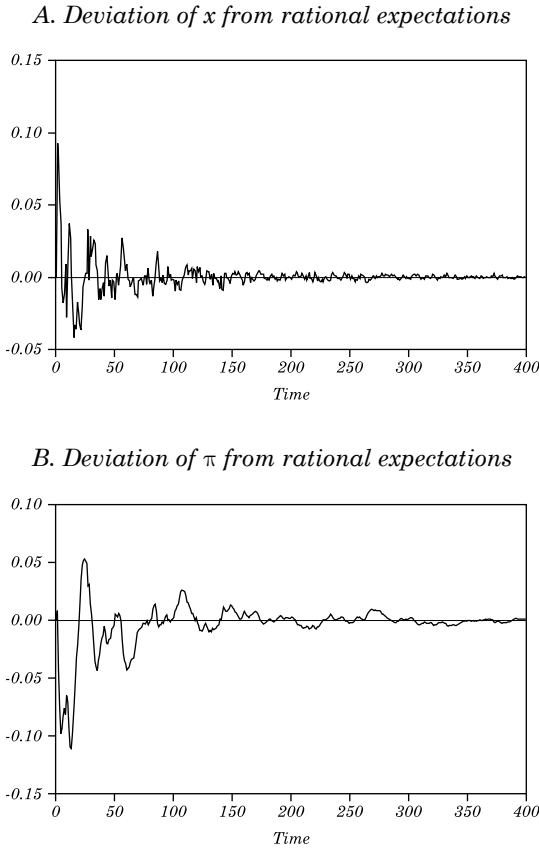

Source: Authors' calculations.

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

In equation (21), the feedback from inflation expectations to actual inflation is stabilizing since the coefficient $\alpha\beta/(\alpha+\lambda^2)$ is less than one and the influence of x_{t-1} is also weak. Deviations from rational expectations are thus offset by policy in such a way that under learning private agents are guided over time to form expectations consistent with the optimal equilibrium. Our expectations-based rule obeys a form of the Taylor principle, since $\delta_\pi > 1$. Figure 2 illustrates convergence of learning under rule (8).

Our optimal policy rule is conditioned on both private expectations and observable exogenous shocks, as well as lagged output. Moreover, when computing the optimal expectations-based rule, the central bank must 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 agents followed the long-horizon decision rules considered by Preston (2005, 2006).

Variations of fundamentals-based rules can perform well in some cases, at least for a relevant region of structural parameter values.

Figure 2. Stability with an Expectations-Based Rule

A. Deviation of x from rational expectations

B. Deviation of π from rational expectations

Source: Authors' calculations.

For the hybrid rule suggested by Svensson and Woodford (2005), numerical analysis shows that, in calibrated models, 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 rational expectations equilibrium. Another favorable case emerges if policy objective (5) is extended to include a motive for interest rate stabilization. Duffy and Xiao (2007b) show that in this case an optimal Taylor-type rule can deliver determinacy and E-stability for a region of parameter values

that includes the usual calibrations used in the literature. We comment further below on stability with constant-gain learning for operational versions of these rules.

Finally, some researchers have proposed monetary policy formulations other than interest rate rules. For example, policy could be formulated as a money supply rule, such as 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, but it does not perform well in terms of the policy objective function. Dennis and Ravenna (2008) examine stability of the economy under optimal discretionary policy, formulated as a targeting rule, for different forms of private agent learning.

2.3 Some Practical Concerns

Many of the i_t rules discussed above may not be operational, as discussed in McCallum (1999). For example, McCallum and Nelson (2004) note that it may be unrealistic to assume that policymakers can condition policy on current x_t and π_t. Similarly, policymakers may not have access to accurate observations on private expectations. We consider these points in 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).

2.3.1 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. If measurement errors are large, however, 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 equation (8).

Suppose now that agents and the central bank begin with different initial estimates, possibly have different learning rules, and use data

sets with different initial dates. When the private agents and the central bank are separately estimating and forecasting 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). Honkapohja and Mitra (2005b) analyze this issue for the case of optimal discretionary policy and expectations-based interest rate rules. Evans and Honkapohja (2003a) show that using VAR proxies can also achieve convergence to the optimal rational expectations equilibrium with commitment. Finally, Muto (2008) considers the consequences of learning from the published central bank's forecast.

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, such as when private agents estimate parameters using stochastic gradient techniques while the central bank uses least squares.

2.3.2 Unavailability of current data

A difficulty with the standard Taylor rule (equation 4), as well as the hybrid rules of Svensson and Woodford (2005) and McCallum and Nelson (2004), is that they presuppose that the policymaker can observe both the 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 unobservability can be avoided by the use of "nowcasts" $E_t^* \mathbf{y}_t$ in place of the actual data $\mathbf{y}_t$. 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 the Svensson-Woodford rule under standard values of the parameters. The situation is more complex for the McCallum-Nelson rule. McCallum and Nelson (2004) suggest using forward expectations in place of actual data. Doing so, however, means that 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.

An additional issue with stability under learning arises when current data are unobservable to the policymaker. If private agents are using constant-gain learning (see section 4.2 for details), the stability conditions are more demanding. As discussed in Evans and Honkapohja (2008), the 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.

2.3.3 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 (2003a, 2003c) extend the basic analysis to a situation in which the central bank estimates the structural parameters φ and λ in equations (1) and (2) and in each period uses the current estimates in its optimal interest rate rule.[22] 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 the central bank can learn their values asymptotically. 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.

Orphanides and Williams (2007) also stressed the importance of structural uncertainty. 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 along all of these dimensions.

---

22. It is natural to assume that the central bank knows the discount factor, β, and the policy weight, α.

3. FURTHER DEVELOPMENTS

A great deal of recent work extends the results on monetary policy and learning. One of the more significant issues, from an applied point of view, is the incorporation of constant-gain or perpetual learning, in which private agents update estimates using least squares, but discount past data. Consequently, agents' expectations never fully converge to the rational expectations equilibrium, but they are (asymptotically) in a neighborhood of the equilibrium, provided the equilibrium is stable. 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 (Orphanides and Williams, 2005a, 2007; Molnar and Santoro, 2006; Gaspar, Smets, and Vestin, 2005, 2006). 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 (2008), 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. Another extension of the basic model considers determinacy and E-stability of rational expectations equilibrium when long-term interest rates are introduced to the model (see McGough, Rudebusch, and Williams, 2005; Tesfaselassie, Schaling, and Eijffinger, 2008).

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 (2008b) extend the analysis to the model with search in labor markets, while Wieland (2008) analyzes the role of endogenous indexation for inflation targeting. Kurozumi and van Zandweghe (2008a), Duffy and Xiao (2007a), and Pfajfar and Santoro (2007a) examine 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, that is, on whether adjustment costs are included and 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 require the interest rate rule to have a positive response to the output gap.

Learning plays a key role in a number of detailed policy issues. 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. Honkapohja and Mitra (2005a) explore how this affects stability under learning. 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.[23] Benhabib, Schmitt-Grohé, and Uribe (2001) investigate this issue under perfect foresight. Evans and Honkapohja (2005) analyze this set-up under learning for the case of flexible prices, while Evans, Guse, and Honkapohja (2008) focus on a sticky-price version. The latter paper is discussed further below. Sticky-information models that incorporate learning have also been developed (Branch and others, 2007, 2008).

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? Honkapohja and Mitra (2004), Carlstrom and Fuerst (2004), and Evans and McGough (2005a) examine this issue in the New-Keynesian setting, where conditions for stable sunspots are obtained in linearized models, while Eusepi (2007) 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 seigniorage.

Constant-gain learning raises the issue of the appropriate choice of gain parameter (see Evans and Honkapohja, 1993, 2001, chap. 14; Marcet and Nicolini, 2003). Evans and Ramey (2006) consider this problem 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 rational expectations, can carry over to learning dynamics in which agents have misspecified models.

---

23. Bullard and Cho (2005) study the possibility of liquidity traps under learning using a linearized New-Keynesian model.

A number of papers model monetary policy with near-rational expectations. Woodford (2005) develops a min-max concept of policy robustness in which policymakers protect against agents' expectations being distorted away from rational expectations within some class of near rational expectations. Bullard, Evans, and Honkapohja (2008) 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 that is increasingly receiving attention. Theoretical work on monetary policy that allows for learning heterogeneity across private agents, or between policymakers and private agents, includes Evans, Honkapohja, and Marimon (2001), Giannitsarou (2003), and Honkapohja and Mitra (2005b, 2006). Guse (2008) and Berardi (2008) introduce misspecified expectations to the New-Keynesian model, while Tetlow and von zur Muehlen (2008) introduce robustness considerations to the analysis of stability. A related approach 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 (see Brock and Hommes, 1997; Branch and Evans, 2006b). For an application to monetary inflation models and monetary policy, see Branch and Evans (2007) and Brazier and others (2008).

Empirical applications of learning to macroeconomics and monetary policy include Bullard and Eusepi (2005) and Orphanides and Williams (2005b), who look at estimated models that focus on the explanation of the large increase in inflation rates in the 1970s. Milani (2005, 2007) incorporates learning as a way to explain persistence in New-Keynesian models, using U.S. data. The first attempts to incorporate learning to applied dynamic stochastic general equilibrium (DSGE) models have recently been undertaken by Slobodyan and Wouters (2007) and Murray (2007). Several papers use least-squares learning models or dynamic predictors to explain expectations data, including Branch (2004), Branch and Evans (2006a), Orphanides and Williams (2005c), 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 (we discuss this paper in detail below). In addition, Cogley and Sargent (2005), Sargent, Williams, and Zha (2006), Primiceri (2006), Ellison and Yates (2007), and Carboni and Ellison (2007, 2008)

emphasize the importance of policymaker 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 analyses, but the main focus in each case goes well beyond the stability of rational expectations equilibrium under learning.

4. PERPETUAL LEARNING AND PERSISTENCE

The preceding sections were concerned with the stability of the rational expectations equilibrium under least squares (LS) learning. That is, we used LS learning to assess whether a rational expectations equilibrium is attainable if we model agents as econometricians. Orphanides and Williams (2005a) 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 equilibrium is stable under learning.

Orphanides and Williams work with a simple two-equation macroeconomic model. The first equation is a new classical expectations-augmented Phillips curve with inertia:

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

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

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

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

Here, x_t is set by monetary policy at t, and u_{t+1} is white noise. Through monetary policy it is assumed that one period ahead, policymakers are able 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 suppresses issues of monetary control.

4.1 Optimal Policy under Rational Expectations

At time t the only state variable is π_t. Policymakers 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,

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

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

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

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

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

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

Substituting into equation (22) yields

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

Substituting in 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 L over θ to get θ^P, the optimal choice of θ under rational expectations. Orphanides and Williams (2005a) show that

$$\theta^P = \theta^P\left(\omega, \frac{1 - \phi}{\alpha}\right),$$

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, which is sometimes called the Taylor curve.

For this choice of feedback parameter, in the rational expectations equilibrium inflation follows the process

$$\pi_t = c_0^P + c_1^P \pi_{t-1} + noise_t$$

and

$$E_t \pi_{t+1} = c_0^P + c_1^P \pi_t,$$

where $c_0^P = \alpha\theta^P /(1 - \phi)$ and $c_1^P = 1 - [\alpha\theta^P /(1 - \phi)]$. Here $noise_t$ is white noise. The superscript P refers to perfect knowledge, which Orphanides and Williams use as a synonym for rational expectations.

The problem is thus quite straightforward under rational expectations. How "aggressive" policy should be with respect to deviations of inflation from target depends naturally on the structural parameters ϕ and α and on the policymaker preferences as described by ω.

4.2 Least-Squares Learning

We now make the crucial step of backing away from rational expectations. Instead of assuming that agents are endowed a priori with rational expectations, 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, estimating its parameters, and using the estimated model to forecast. Specifically, suppose private agents believe that inflation follows a first-order autoregressive, or AR(1), process, as it does in a rational expectations equilibrium, but that they do not know c_0^P, c_1^P. Instead they estimate the parameters of

$$\pi_t = c_0 + c_1 \pi_t + v_t$$

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

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

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 policymakers do not explicitly take account of private agent learning and follow the feedback rule with $\theta = \theta^P$. Then, with infinite memory (that is, no discounting of observations), one can show that

$$c_{0,t}, c_{1,t} \rightarrow c_0^P, c_1^P$$

with probability 1. Asymptotically, we get the optimal rational expectations equilibrium.

Orphanides and Williams (2005a) 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 (that is, recursive least squares, or RLS), the coefficient 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 $t \rightarrow \infty$. Orphanides and Williams 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, that is, by setting $\kappa_t = \kappa$ (for example, 0.05). This is equivalent to using weighted least squares with weights declining geometrically in time as we move backward 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 approach that is sometimes used in practice is to implement a rolling data-window of finite length.[24]

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

4.3 Implications of Constant-Gain Least Squares

With constant-gain procedures, estimates no longer fully converge to the rational expectations equilibrium. The estimators $c_{0,t}$, $c_{1,t}$ converge instead to a stochastic process. Orphanides and Williams (2005a) therefore 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 the equilibrium values most of the time with a high probability, and output and inflation will be near their equilibrium paths. Nonetheless, small plausible values like $\kappa = 0.05$ can lead to very different outcomes in the calibrations Orphanides and Williams 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 $\omega = 0.01, 0.50,$ and 1.00, respectively, under rational expectations.

Their main findings are threefold. First, the standard deviations of $c_{0,t}$ and $c_{1,t}$ are large even though forecast performance remains good. Second, the persistence of inflation increases substantially, compared with the rational expectations equilibrium, as measured by the AR(1) coefficient for π_t. Finally, the policy frontier shifts out very substantially and sometimes in a nonmonotonic way.

4.4 Policy Implications

Under perpetual learning by private agents, if policymakers keep to the same class of rules,

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

then they should choose a different θ than under rational expectations. Here the notation θ^S indicates that we restrict policymakers to choosing 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, that is, the policy that assumes rational expectations (perfect knowledge) on the part of agents, can be strictly inefficient when the agents are, in fact, 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, policy should generally be more hawkish—that is, under perpetual learning the monetary authorities should pick a larger θ^S than if agents had rational expectations.

Third, following a sequence of unanticipated inflation shocks, inflation doves (that is, policymakers with a low θ, reflecting a low ω) can do very poorly, as these shocks can lead expectations to temporarily but persistently deviate substantially from rational expectations. 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 indicates how the performance of policy depends on expectations formation and what the policymaker assumes about it. The middle curve is the efficient policy under learning, while "naive" refers to the case in which policy presumes rational expectations while agents are in fact learning with gain $\kappa = 0.05$.

Figure 3. The Policymaker's Loss
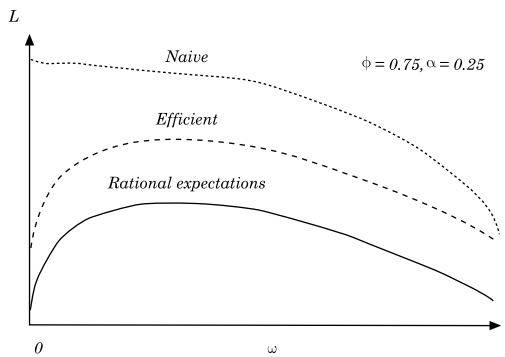
Source: Authors' drawing, adapted from Orphanides and Wiliams (2005a), figure 7.

Perpetual learning thus turns out to have major implications for policy, even when the deviation from the rational expectations equilibrium might not be thought to be too large. The main policy implication is that with perpetual learning, there should be a policy bias toward hawkishness. The intuition for this result that a more hawkish policy (high θ) helps to keep inflation expectations, π^e_{t+1},

in line, or closer to rational expectations values. This qualitative result also emerges in the more general setting in Orphanides and Williams (2007).

5. ESTIMATED MODELS WITH LEARNING

The Orphanides and Williams (2005a) results suggest another implication of learning that goes beyond policy, namely, that learning itself can be a source of persistence in macroeconomic dynamics. The starting point for this line of thought, as pursued by Milani (2005, 2007), 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 rational expectations assumption. The source of the backward-looking component used in these hybrid models, however, is 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 a 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 rational expectations empirically finds values of γ that are close to one.

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

$$E_t^*\pi_{t+1} = \phi_{0,t} + \phi_{1,t}E_t^*\pi_t$$

$$= \phi_{0,t} + \phi_{1,t}(\phi_{0,t} + \phi_{1,t}\pi_{t-1}),$$

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 the output gap, x_t. To estimate the model for the United States, Milani computes inflation from the GDP deflator and the output gap as detrended GDP, while real marginal costs are proxied by the deviation of the labor income share from 1960:1 to 2003:4. Agents' initial parameter estimates are obtained by using presample data from 1951–59.

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. Milani then 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: $\phi_{1,t}$ was initially low in the 1950s and 1960s, before rising (up to 0.958) and then declining somewhat to values above 0.8; $\phi_{0,t}$ was also initially low before rising sharply and then gradually declining 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's Bayesian information criterion. The estimate of γ is not significantly different from zero, and it constrasts sharply with the high levels of γ found under the rational expectations assumption. It thus 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 the degree of habit persistence is also low in the IS curve. This contrasts with the usual extension of the New-Keynesian model under rational expectations 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.

6. RECURRENT HYPERINFLATIONS

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

6.1 The Basic Hyperinflation Model

The starting point is the theoretical model sometimes known as the seigniorage model of inflation (see Evans and Honkapohja, 2001, chap. 11). The Cagan model is based on the linear money demand equation, which can be obtained from an overlapping generations (OG) endowment economy with log utility. Specifically,

$$\frac{M_t^d}{P_t} = \phi - \phi\gamma\frac{P_{t+1}^e}{P_t},$$

if $1 - \gamma(P_{t+1}^e / P_t) > 0$ and 0 otherwise. 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 (that is, $P_{t+1}^e / P_t = P_{t+1} / P_t$) there are two steady states, $\beta_L < \beta_H$, provided $d \geq 0$ is not too large. 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 a steady state, that is, $P_{t+1} / P_t = \beta + \eta_t$, 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, that is, $(P_{t+1}/P_t)^e = \beta_t$, where

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

This is simply a recursive algorithm for the average inflation rate, which is equivalent to a least-squares regression on a constant.[25] The stability of this learning rule is governed by the E-stability differential equation

$$\frac{d\beta}{d\tau} = T(\beta;d) - \beta,$$

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.

Figure 4. Steady-State Learning in the Hyperinflation Model
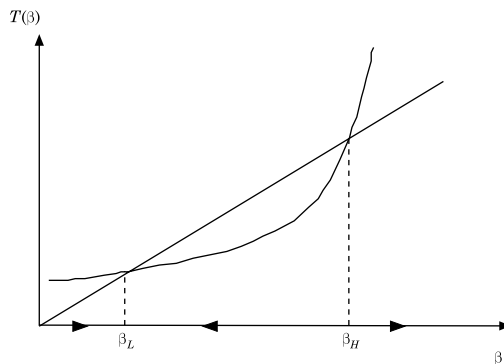
25. One can consider more general classes of PLM. Adam, Evans, and Honkapohja (2006) study the circumstances in which autoregressive PLMs can converge to hyperinflation paths.

An increase in d shifts $T(\beta)$ up, so the comparative statics of β_L are natural while 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 rational expectations version of this model as a theoretical explanation for hyperinflations.

6.2 Empirical Background

Marcet and Nicolini (2003) list four stylized facts about hyperinflation episodes in the 1980s in a number of South American countries (as well as some episodes in other places and at other times): (1) hyperinflation episodes are recurrent; (2) exchange rate rules stop hyperinflations, although new hyperinflations eventually occur; (3) during a hyperinflation, seigniorage and inflation are not highly correlated; and (4) average inflation and seigniorage are strongly positively correlated across countries, with hyperinflations only occurring in countries where seigniorage is high, on average. Stabilization plans to deal with hyperinflation have been based either on heterodox policy (exchange rate rules) or orthodox policy (permanently reducing the deficit). Policies that combine both elements appear to have been successful in stopping hyperinflations permanently.

6.3 The Marcet-Nicolini Model

Marcet and Nicolini (2003) use an open economy version of the overlapping-generations hyperinflation model. This is a flexible price model with purchasing power parity (PPP), so that

$$P_t^f e_t = P_t,$$

where P_t^f is the foreign price of goods, which is assumed to be 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 i.i.d.

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 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, that is, $(M_t - M_{t-1})/P_t = d_t + [(R_t - R_{t-1})e_t]/P_t$. In equilibrium, any increase in reserves must be matched by a trade surplus, that is, $(R_t - R_{t-1})e_t = TB_t \cdot P_t$, where TB_t is total endowment minus total private consumption minus d_t.

The key question is the form of the exchange rate rule. When an exchange rate rule 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 the exchange rate rule, 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 its exchange rate regime. We assume there is a maximum inflation rate tolerated, β_U. The exchange rate regime is imposed only in periods when inflation would otherwise exceed this bound (or if no positive P_t would otherwise clear the market).

6.4 Learning

Marcet and Nicolini (2003) argue that under rational expectations, the model cannot properly explain the stylized facts of hyperinflation outlined above. An adaptive learning formulation will be more successful. They 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 a given β_0. Here $1/\alpha_t = \kappa_t$ is what we have called the gain, $\alpha_t = \alpha_{t-1} + 1$ corresponds to decreasing gain learning, and $\alpha_t = \bar{\alpha} > 1$ is a constant-gain algorithm (α_t can also be thought of as the effective sample size). Marcet and Nicolini consider a version in which agents switch between decreasing and constant gain according to recent performance. Specifically,

$$\alpha_t = \alpha_{t-1} + 1 \quad \text{if} \quad \left| \left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) \middle/ \beta_{t-1} \right|$$

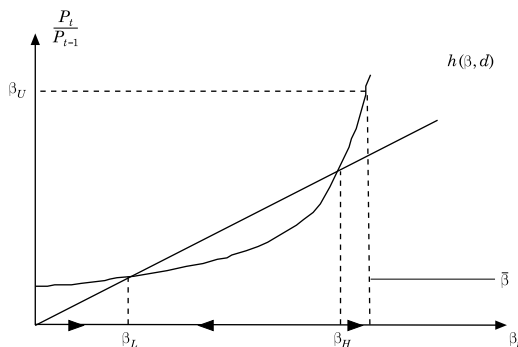falls below some bound v, and otherwise $\alpha_t = \bar{\alpha}$.

The qualitative features of the model are approximated by the system

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

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

Figure 5 describes the dynamics of system (24).

**Figure 5. Inflation as a Function of Expected Inflation**

Source: Authors' drawing, adapted from Marcet and Nicolini (2003), figure 3.

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. 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 the exchange rate rule. Then the process begins again. The model with learning has the following three features. First, there may be eventual convergence to rational expectations. This can occur if the random shocks/learning dynamics do not push β_t into the unstable region for a long time. Decreasing gain may then lead to asymptotic convergence to β_L. Second, a higher $E(d_t)$ raises both average inflation and the frequency of hyperinflations. A combination of orthodox and heterodox policies make sense as a way to end hyperinflations. Third, 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.

7. LIQUIDITY TRAPS AND DEFLATIONARY SPIRALS

Deflation and liquidity traps have been a concern in recent times. Evans, Guse, and Honkapohja (2008) consider 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 to ensure determinacy and stability under learning. As emphasized by Benhabib, Schmitt-Grohé, and Uribe (2001), however, if one considers the interest rate rule globally, rather than 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 that, for any continuous rule, 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.[26] 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.

---

26. Taylor rules usually also include a dependence on aggregate output, which we omit for simplicity.

Figure 6. Multiple Steady States with a Global Taylor Rule
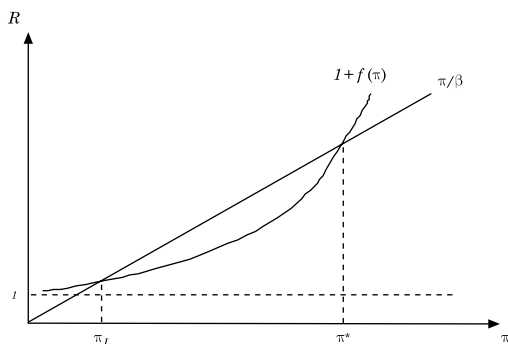
Source: Authors' drawing.

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 $\pi - 1$ 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. In addition, π_L may correspond to either a very low positive inflation rate or to a negative net inflation rate, that is, deflation. The zero lower bound corresponds to $R = 1$. Benhabib, Schmitt-Grohé, and Uribe (2001) show that under rational expectations, there is a continuum of liquidity trap paths that converge on π_L. The pure rational expectations 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 show that deflationary paths are possible, but that the real risk, under learning, involves paths in which inflation slips below π_L and then continue to fall further. For this flexible-price model, we show that this can 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 (2008) 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

monopolistic competition and 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-Grohé, 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

$$\frac{\alpha\gamma}{\nu}(\pi_t - 1)\pi_t = \beta\frac{\alpha\gamma}{\nu}(\pi_{t+1}^e - 1)\pi_{t+1}^e + (c_t + g_t)^{(1+\varepsilon)/\alpha}$$
$$-\alpha\left(1 - \frac{1}{\nu}\right)(c_t + g_t)c_t^{-\sigma_1}$$

and

$$c_t = c_{t+1}^e\left(\frac{\pi_{t+1}^e}{\beta R_t}\right)^{\sigma_1}.$$

The first equation is the New-Keynesian Phillips curve, relating π_t positively to π_{t+1}^e 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.

There are two stochastic steady states at π_L and π_H. If the random shocks are i.i.d., then steady-state learning is appropriate for both c^e and π^e, that is,

$$\pi_{t+1}^e = \pi_t^e + \phi_t(\pi_{t-1} - \pi_t^e)$$

and

$$c_{t+1}^e = c_t^e + \phi_t(c_{t-1} - c_t^e),$$

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 and π^e will follow paths leading to deflation and stagnation. This is illustrated in figure 7, based on E-stability dynamics.

Figure 7. The Dynamics of π^e and c^e under Normal Policy
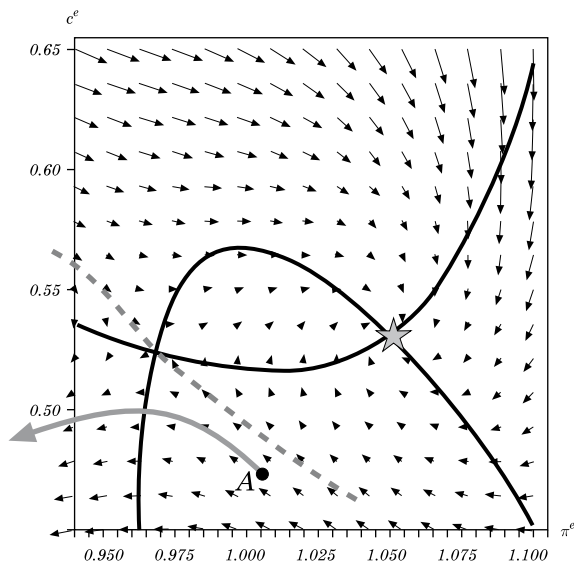
Source: Based on Evans, Guse, and Honkapohja (2008), figure 4.

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 a fixed R this would lead through the IS curve to a lower c and thus through the Phillips curve, to a lower π. A sufficient reduction in R would prevent the reductions in c and π, but this is not possible since we are close to the zero lower bound, and the global Taylor rule here dictates only small reductions in R. The falls in realized c and π then leads under learning to reductions in c^e and π^e, and this sets the deflationary spiral in motion.

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.

Evans, Guse, and Honkapohja (2008) show that policy can be altered to avoid the deflationary spiral. 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, $\tilde{\pi} = 1$. 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 (2008) 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^*.[27] Perhaps surprisingly, they also show that it is essential to use an inflation threshold, since using an output threshold to trigger aggressive polices will not always avoid deflationary spirals.

8. Conclusions

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

Convergence to rational expectations is possible in many economic models, with an appropriate econometric perceived law of motion. However, the stability of rational expectations equilibrium 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

27. For non-Ricardian economies, Bénassy (2007) develops an alternative interest rate rule that leads to global uniqueness.

well in this respect, but implementing optimal policy appears to require an appropriate response to private sector expectations about the future.

More generally, policymakers 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 also shows 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(12): 2725–48.

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. 2007. *Inflation Report*. London.

Basdevant, O. 2005. "Learning Process and Rational Expectations: An Analysis Using a Small Macroeconomic Model for New Zealand." *Economic Modelling* 22(6): 1074–89.

Bénassy, J-P. 2007. *Money, Interest, and Policy: Dynamic General Equilibrium in a non-Ricardian World.* MIT Press.

Benhabib, J., S. Schmitt-Grohé and M. Uribe. 2001. "The Perils of Taylor Rules." *Journal of Economic Theory* 96(1–2): 40–69.

Berardi, M. 2008. "Monetary Policy with Heterogeneous and Misspecified Expectations." *Journal of Money, Credit, and Banking* (forthcoming).

Berardi, M. and J. Duffy. 2007. "The Value of Central Bank Transparency When Agents Are Learning." *European Journal of Political Economy* 23(1): 9–29.

Bernanke, B. 2007. "Inflation Expectations and Inflation Forecasting." Speech given at the Monetary Economics Workshop. National Bureau of Economic Research, Summer Institute, Cambridge, Massachusetts, 10 July.

Bernanke, B. and M. Woodford. 1997. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit, and Banking* 29(4): 653–84.

Branch, W.A. 2004. "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations." *Economic Journal* 114(497): 592–621.

Branch, W.A., J. Carlson, G.W. Evans and B. McGough. 2007. "Adaptive Learning, Endogenous Inattention, and Changes in Monetary Policy." Working paper 2006-6, revised. University of Oregon.

————. 2008. "Monetary Policy, Endogenous Inattention, and the Volatility Trade-Off." *Economic Journal* (forthcoming).

Branch, W.A. and G.W. Evans. 2006a. "A Simple Recursive Forecasting Model." *Economic Letters* 91(2): 158–66.

————. 2006b. "Intrinsic Heterogeneity in Expectation Formation." *Journal of Economic Theory* 127(1): 264–95.

————. 2007. "Model Uncertainty and Endogenous Volatility." *Review of Economic Dynamics* 10(2): 207–37.

Brazier, A., R. Harrison, M. King and T. Yates. 2008. "The Danger of Inflating Expectations of Macroeconomic Stability: Heuristic Switching in an Overlapping Generations Monetary Model." *International Journal of Central Banking* 4(2): 219–54.

Brock, W.A., and C.H. Hommes. 1997. "A Rational Route to Randomness." *Econometrica* 65(5): 1059–95.

Bullard, J. 2006. "The Learnability Criterion and Monetary Policy." *Federal Reserve Bank of St. Louis Review* 88 (May): 203–17.

Bullard, J. and I-K. Cho. 2005. "Escapist Policy Rules." *Journal of Economic Dynamics and Control* 29(11): 1841–66.

Bullard, J. and S. Eusepi. 2005. "Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?" *Review of Economic Dynamics* 8(2): 324–59.

Bullard, J., G.W. Evans and S. Honkapohja. 2008. "Monetary Policy, Judgment, and Near-Rational Exuberance." *American Economic Review* 98(3): 1163–77.

Bullard, J. and K. Mitra. 2002. "Learning About Monetary Policy Rules." *Journal of Monetary Economics* 49(6): 1105–29.

————. 2007. "Determinacy, Learnability, and Monetary Policy Inertia." *Journal of Money, Credit, and Banking* 39(5): 1177–212.

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.

Bullard, J. and A. Singh. 2006. "Worldwide Macroeconomic Stability and Monetary Policy Rules." Working paper 2006-040b. Federal Reserve Bank of St Louis.

Calvo, G.A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12(3): 383–98.

Carboni, G. and M. Ellison. 2007. "Learning and the Great Inflation." Discussion paper 6250. London: Centre for Economic Policy Research.

————. 2008. "The Great Inflation and the Greenbook." Frankfurt: European Central Bank.

Carlstrom, C.T. and T.S. Fuerst. 2004. "Learning and the Central Bank." *Journal of Monetary Economics* 51(2): 327–38.

Clarida, R., J. Galí and M. Gertler. 1998. "Monetary Policy Rules in Practice: Some International Evidence." *European Economic Review* 42(6): 1033–67.

———. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37(4): 1661–707.

Cogley, T. and T.J. Sargent. 2005. "The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics* 8(2): 528–63.

Dennis, R. and F. Ravenna. 2008. "Learning and Optimal Monetary Policy." *Journal of Economic Dynamics and Control* 32(6): 1964–94.

Duffy, J. and W. Xiao. 2007a. "Investment and Monetary Policy: Learning and Determinacy of Equilibria." Working paper 324. University of Pittsburgh, Department of Economics.

———. 2007b. "The Value of Interest Rate Stabilization Policies When Agents are Learning." *Journal of Money, Credit, and Banking* 39(8): 2041–56.

Ellison, M. and T. Yates. 2007. "Escaping Volatile Inflation." *Journal of Money, Credit, and Banking* 39(4): 981–93.

European Central Bank. 2007. *Monthly Bulletin* (June and December). Frankfurt.

Eusepi, S. 2007. "Learnability and Monetary Policy: A Global Perspective." *Journal of Monetary Economics* 54(4): 1115–31.

Eusepi, S. and B. Preston. 2007. "Central Bank Communication and Expectations Stabilization." Working paper 13259. Cambridge, Mass.: National Bureau of Economic Research.

Evans, G.W., E. Guse and S. Honkapohja. 2008. "Liquidity Traps, Learning, and Stagnation." *European Economic Review* (forthcoming).

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 *Learning and Rationality in Economics*, edited by A. Kirman and M. Salmon, chap. 4. Oxford: Basil Blackwell.

———. 1999. "Learning Dynamics." In *Handbook of Macroeconomics*, vol. 1, edited by J. Taylor and M. Woodford, chap. 7. Amsterdam: Elsevier.

———. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.

————. 2003a. "Adaptive Learning and Monetary Policy Design." *Journal of Money, Credit, and Banking* 35(6): 1045–72.

————. 2003b. "Expectational Stability of Stationary Sunspot Equilibria in a Forward-looking Model." *Journal of Economic Dynamics and Control* 28(1): 171–81.

————. 2003c. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies* 70(4): 807–24.

————. 2003d. "Friedman's Money Supply Rule vs. Optimal Interest Rate Policy." *Scottish Journal of Political Economy* 50(5) 550–66.

————. 2005. "Policy Interaction, Expectations, and the Liquidity Trap." *Review of Economic Dynamics* 82 303–23.

————. 2006. "Monetary Policy, Expectations, and Commitment." *Scandinavian Journal of Economics* 108(1): 15–38.

————. 2007. "Policy Interaction, Learning, and the Fiscal Theory of Prices." *Macroeconomic Dynamics* 11(5): 665–90.

————. 2008. "Robust Learning Stability with Operational Monetary Policy Rules." In *Monetary Policy Under Uncertainty and Learning* edited by K. Schmidt-Hebbel and C.E. Walsh, XX-YY. Central Bank of Chile.

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

————. 2007. "Stable Sunspot Equilibria in a Cash-in-Advance Economy." *The Berkeley Electronic Journal of Macroeconomics (Advances)* 7(1): 1165.

Evans, G.W. and B. McGough. 2005a. "Monetary Policy, Indeterminacy, and Learning." *Journal of Economic Dynamics and Control* 29(11): 1809–40.

————. 2005b. "Stable Sunspot Solutions in Models with Predetermined Variables." *Journal of Economic Dynamics and Control* 29(4): 601–25.

————. 2007. "Optimal Constrained Interest-Rate Rules." *Journal of Money, Credit, and Banking* 39(6): 1335–56.

Evans, G.W. and G. Ramey. 2006. "Adaptive Expectations, Underparameterization, and the Lucas Critique." *Journal of Monetary Economics* 53(2): 249–64.

Ferrero, G. 2007. "Monetary Policy, Learning, and the Speed of Convergence." *Journal of Economic Dynamics and Control* 31(9): 3006–41.

Gaspar, V., F. Smets and D. Vestin. 2005. "Optimal Monetary Policy Under Adaptive Learning." Bank of Portugal.

————. 2006. "Adaptive Learning, Persistence, and Optimal Monetary Policy." *Journal of the European Economic Association* 4(2–3): 376–85.

Giannitsarou, C. 2003. "Heterogeneous Learning." *Review of Economic Dynamics* 6(4): 885–906.

Guse, E. 2008. "Learning in a Misspecified Multivariate Self-Referential Linear Stochastic Model." *Journal of Economic Dynamics and Control* 32(5): 1517–42.

Honkapohja, S. and K. Mitra. 2003. "Learning with Bounded Memory in Stochastic Models." *Journal of Economic Dynamics and Control* 27(8): 1437–57.

————. 2004. "Are Non-Fundamental Equilibria Learnable in Models of Monetary Policy?" *Journal of Monetary Economics* 51(1): 1743–70.

————. 2005a. "Performance of Inflation Targeting Based on Constant Interest Rate Projections." *Journal of Economic Dynamics and Control* 29(11): 1867–92.

————. 2005b. "Performance of Monetary Policy with Internal Central Bank Forecasting." *Journal of Economic Dynamics and Control* 29(4): 627–58.

————. 2006. "Learning Stability in Economies with Heterogeneous Agents." *Review of Economic Dynamics* 9(2): 284–309.

Kurozumi, T. 2006. "Determinacy and Expectational Stability of Equilibrium in a Monetary Sticky-Price Model with Taylor Rule." *Journal of Monetary Economics* 53(4): 827–46.

Kurozumi, T. and W. Van Zandweghe. 2008a. "Investment, Interest Rate Policy, and Equilibrium Stability." *Journal of Economic Dynamics and Control* 32(5): 1489–516.

————. 2008b. "Labor Market Search and Interest Rate Policy." Tokyo: Bank of Japan.

Llosa, G. and V. Tuesta. 2007. "E-Stability of Monetary Policy When the Cost Channel Matters." Washington: Inter-American Development Bank.

————. 2008. "Determinacy and Learnability of Monetary Policy Rules in Small Open Economies." *Journal of Money, Credit and Banking* 40(5): 1033-63.

Lucas, R.E., Jr. 1976. "Econometric Policy Evaluation: A Critique." In *Carnegie-Rochester Conference Series on Public Policy* 1: 19–46.

Marcet, A. and J.P. Nicolini. 2003. "Recurrent Hyperinflations and Learning." *American Economic Review* 93(5): 1476–98.

Marimon, R. 1997. "Learning from Learning in Economics." In *Advances in Economics and Econometrics: Theory and Applications*, vol. 1, edited by D. Kreps and K. Wallis, chap. 9. Cambridge University Press.

McCallum, B.T. 1999. "Issues in the Design of Monetary Policy Rules." In *Handbook of Macroeconomics*, vol. 1, edited by J. Taylor and M. Woodford, chap. 23. Amsterdam: Elsevier.

————. 2003. "Is the Fiscal Theory of the Price Level Learnable?" *Scottish Journal of Political Economy* 50(5): 634–49.

————. 2007. "E-stability vis-à-vis Determinacy Results for a Broad Class of Linear Rational Expectations Models." *Journal of Economic Dynamics and Control* 31(4): 1376–91.

McCallum, B.T. and E. Nelson. 1999. "Performance of Operational Policy Rules in an Estimated Semi-Classical Model." In *Monetary Policy Rules*, edited by J. Taylor, pp. 15–45. University of Chicago Press.

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

McGough, B., G. Rudebusch and J. Williams. 2005. "Using a Long-Term Interest Rate as the Monetary Policy Instrument." *Journal of Monetary Economics* 52(5): 855–79.

Milani, F. 2005. "Adaptive Learning and Inflation Persistence." Working paper 050607. University of California at Irvine, Department of Economics.

————. 2007. "Expectations, Learning, and Macroeconomic Persistence." *Journal Of Monetary Economics* 54(7): 2065–82.

Mitra, K. 2003. "Desirability of Nominal GDP Targeting under Adaptive Learning." *Journal of Money, Credit, and Banking* 35(2): 197–220.

Molnar, K. and S. Santoro. 2006. "Optimal Monetary Policy When Agents Are Learning." Discussion paper 2006/01. Institute of Economics, Hungarian Academy of Sciences.

Murray, J. 2007. "Empirical Significance of Learning in a New Keynesian Model with Firm-Specific Capital." Working paper 2007-027. Indiana University, Center for Applied Economics and Policy Research.

Muto, I. 2008. "Monetary Policy and Learning from the Central Bank's Forecast." Discussion paper 2008-E-1. Tokyo: Bank of Japan.

Orphanides, A. and J.C. Williams. 2005a. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation-*

Targeting Debate, edited by B. Bernanke and M. Woodford, chap. 5. University of Chicago Press.

———. 2005b. "Inflation Scares and Forecast-Based Monetary Policy." *Review of Economic Dynamics* 8(2): 498–527.

———. 2005c. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations." *Journal of Economic Dynamics and Control* 29(11): 1927–50.

———. 2007. "Robust Monetary Policy with Imperfect Knowledge." *Journal of Monetary Economics* 54(5): 1406–35.

Pfajfar, D. 2007. "Formation of Rationally Heterogeneous Expectations." University of Cambridge.

Pfajfar, D. and E. Santoro. 2007a. "Credit Market Distortions, Asset Prices, and Monetary Policy." University of Cambridge.

———. 2007b. "Heterogeneity and Learning in Inflation Expectation Formation: An Empirical Assessment." University of Cambridge.

Preston, B. 2005. "Learning about Monetary Policy Rules When Long-Horizon Expectations Matter." *International Journal of Central Banking* 1(2): 81–126.

———. 2006. "Adaptive Learning, Forecast-based Instrument Rules, and Monetary Policy." *Journal of Monetary Economics* 53(3): 507–35.

Primiceri, G.E. 2006. "Why Inflation Rose and Fell: Policymakers' Beliefs and U.S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121(3): 867–901.

Rotemberg, J.J. 1982. "Sticky Prices in the United States." *Journal of Political Economy* 90(6): 1187–211.

Rotemberg, J.J. and M. Woodford. 1999. "Interest Rate Rules in an Estimated Sticky-Price Model." In *Monetary Policy Rules*, edited by J. Taylor, pp. 57–119. University of Chicago Press.

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

———. 1999. *The Conquest of American Inflation*. Princeton University Press.

Sargent, T.J. and N. Wallace. 1975. "'Rational Expectations,' the Optimal Monetary Instrument, and the Optimal Money Supply Rule." *Journal of Political Economy* 83(2): 241–54.

Sargent, T.J., N. Williams and T. Zha. 2006. "Shocks and Government Beliefs: The Rise and Fall of American Inflation." *American Economic Review* 96(4): 1193–224.

Slobodyan, S. and R. Wouters. 2007. "Learning in an Estimated DSGE Model." Prague: CERGE-EI.

Svensson, L.E.O. 1999. "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics* 43(3): 607–54.

————. 2003. "What Is Wrong with Taylor Rules? Using Judgement in Monetary Policy through Targeting Rules." *Journal of Economic Literature* 41(2): 426–77.

Svensson, L.E.O. and M. Woodford. 2005. "Implementing Optimal Policy through Inflation-Forecast Targeting." In *The Inflation-Targeting Debate*, edited by B. Bernanke and M. Woodford, pp. 19–83. University of Chicago Press.

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

Tesfaselassie, M., E. Schaling and S. Eijffinger. 2008. "Learning about the Term Structure and Optimal Rules for Inflation Targeting." Discussion paper 88. Tilburg University, Center for Economic Research.

Tetlow, R.J. and P. von zur Muehlen. 2008. "Robustifying Learnability." *Journal of Economic Dynamics and Control* (forthcoming).

Trichet, J-C. 2005. "Monetary Policy and Private Expectations." Zolotas Lecture at the Bank of Greece. Athens, 25 February.

Walsh, C. 2003. *Monetary Theory and Policy*, 2d ed. MIT Press.

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

Wieland, V. 2008. "Learning, Endogenous Indexation, and Disinflation in the New Keynesian Model." Working paper 6749. London: Centre for Economic Policy Research.

Woodford, M. 1999a. "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" In *New Challenges for Monetary Policy*, edited by the Federal Reserve Bank of Kansas City, pp. 277–316.

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

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

————. 2005. "Robustly Optimal Monetary Policy with Near-Rational Expectations." Working paper 11896. Cambridge, Mass.: National Bureau of Economic Research.

Zanna, L-F. 2006. "PPP Rules, Macroeconomic (In)stability, and Learning." International finance discussion paper 814. Washington: Board of Governors of the Federal Reserve System.