Robust Learning Stability with Operational Monetary Policy Rules*

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

Seppo Honkapohja
Bank of Finland,
University of Cambridge

January 1, 2008

Abstract

We consider "robust stability" of a rational expectations equilibrium, which we define as stability under discounted (constant gain) least-squares learning, for a range of gain parameters. We find that for operational forms of policy rules, i.e. rules that do not depend on contemporaneous values of endogenous aggregate variables, many interest-rate rules do not exhibit robust stability. We consider a variety of interest-rate rules, including instrument rules, optimal reaction functions under discretion or commitment, and rules that approximate optimal policy under commitment. For some reaction functions we allow for an interest-rate stabilization motive in the policy objective. The expectations-based rules proposed in Evans and Honkapohja (2003, 2006) deliver robust learning stability. In contrast, many proposed alternatives become unstable under learning even at small values of the gain parameter.

*Paper prepared for the 11th Annual Conference of the Central Bank of Chile, "Monetary Policy Under Uncertainty and Learning," November 15-16, 2007. We thank our discussant Claes Berg, Carl Walsh and the conference participants for valuable comments. Support from National Science Foundation Grant No. SES-0617859 and ESRC grant RES-000-23-1152 is gratefully acknowledged. Any views expressed are those of the authors and do not necessarily reflect the views of the Bank of Finland.

Key words: Commitment, interest-rate setting, adaptive learning, stability, determinacy.

JEL classification: E52, E31, D84.

1 Introduction

Recently, the conduct of monetary policy in terms of interest rate rules has been examined from the viewpoint of imperfect knowledge and learning by economic agents. In this literature stability of rational expectations equilibrium (REE) is taken as a key desideratum for good monetary policy design.[1] Most of this literature postulates that agents use least squares or related learning algorithms to carry out real-time estimations of the parameters of their forecast functions as new data becomes available. Moreover, it is usually assumed that the learning algorithms have a decreasing gain; in the most common case the gain is the inverse of the sample size so that all data points have equal weights. Use of such a decreasing-gain algorithm makes it possible for learning to converge exactly at the REE in environments without structural change. Convergence requires that REE satisfies a stability condition, known as E-stability.

Decreasing-gain algorithms do not, however, perform well when occasional unobservable structural changes take place. So-called constant-gain algorithms are a natural alternative for estimating parameters in a way that is alert to possible structural changes. If agents use a constant-gain algorithm, then parameter estimates of the forecast functions do not fully converge to the REE values. Instead, they remain random, even asymptotically. However, for small values of the gain parameter the estimates remain for most of the time in a small neighborhood of the REE, provided that the REE is E-stable.[2] Recently, constant-gain algorithms have been employed in empirical work, e.g. see Milani (2005), Milani (2007a), Orphanides and Williams (2005b), Orphanides and Williams (2005a) and Branch and Evans (2006).

[1] For surveys see Evans and Honkapohja (2003a), Bullard (2006) and Evans and Honkapohja (2007).

[2] See Chapters 3 and 7 of Evans and Honkapohja (2001) for the basic theoretical results on constant-gain learning. See also Evans, Honkapohja, and Williams (2006) for references on recent papers on constant-gain learning.

It should be emphasized that the connection between convergence of constant-gain learning and E-stability noted above is a limiting result for gain parameters sufficiently small. For finite values of the gain parameter, the stability condition for constant-gain learning is more stringent than E-stability. In this paper we examine the stability implications of various interest rate rules when agents use constant-gain learning rules with plausible positive values of the gain. We will say that an interest rate rules yields *robust learning stability* of the economy if stability under constant-gain learning obtains for all values of the gain parameter in the range suggested by the empirical literature.[3]

In this study we focus on interest rate rules that are operational in the sense discussed by McCallum (1999). He argues that monetary policy cannot be conditioned on current values of endogenous aggregate variables. Thus, the rules we consider assume that policy responds to expectations of contemporaneous (or future) values of inflation and output but not on their actual values in the current period.

We consider robust learning stability for a variety of operational interest rate rules that have been suggested in the recent literature. These include Taylor rules and optimal reaction functions under discretion and commitment when central bank policy aims for interest-rate stabilization in addition to the usual motives for flexible inflation targeting. The reaction function may be expectations-based in the spirit of Evans and Honkapohja (2003b) and Evans and Honkapohja (2006), or of the Taylor-type form suggested by Duffy and Xiao (2007). We also analyze two interest rate rules that approximate optimal policy under commitment and were suggested by Svensson and Woodford (2005) and McCallum and Nelson (2004). Our results show that expectations-based rules deliver robust learning stability, whereas the proposed alternatives often become unstable under learning even at quite small values of the constant gain parameter.

[3] There are numerous concepts of robustness that are relevant to policymaking reflecting, e.g., uncertainty about the structure of the economy, and a desire by both private agents and policymakers to guard against the risk of large losses. We do not mean to downplay the importance of such factors, but we here abstract from them in order to focus on the importance of setting policy in such a way as to ensure stability in the face of constant-gain learning.

2 Constant Gain Steady-State Learning

2.1 Theoretical Results

In this paper we employ multivariate linear models. In this simplest case in which the shocks are white noise and there are no lagged endogenous variables, the REE takes the form of a stochastic steady state. We now briefly review the basics of steady state learning in linear models.[4]

The steady state can be computed by postulating that agents' beliefs, called the "perceived law of motion" (PLM), take the form

$$y_t = a + e_t,$$

for a vector y_t, where $e_t \sim iid(0, \sigma^2)$. Using the model, one then computes the "actual law of motion" (ALM), which describes the temporary equilibrium in the current period, given the PLM. We write the ALM using a linear operator T as

$$y_t = \alpha + Ta + e_t,$$

where the matrix T depends on the structural parameters of the model. Examples of the T map will be given below. An REE is a fixed point $\bar{a}$ of the T map, i.e.

$$\bar{a} = \alpha + T\bar{a}$$

We assume that $I - T$ is non-singular, so that there is a unique solution $\bar{\alpha} = (I - T)^{-1}\alpha$. For convenience, and without loss of generality, we now assume that the model has been written in deviation from the mean form, so that $\alpha = 0$. Thus in our analysis the REE corresponds to $\bar{a} = 0$. Under learning agents attempt to learn the value of $\bar{a}$, and hence in deviation from mean form we are examining whether agents' estimates of the mean converge to $a = 0$.

Steady-state learning under decreasing gain is given by the recursive algorithm

$$a_t = a_{t-1} + \gamma_t(y_t - a_{t-1}), \tag{1}$$

where the gain γ_t is a sequence of small decreasing numbers such as $\gamma_t = 1/t$. Assuming that $y_t = Ta_{t-1} + e_t$, i.e. that expectations are formed using the estimate a_{t-1} based on data through time $t-1$, the convergence condition

[4] See Chapters 8 and 10 of Evans and Honkapohja (2001) for a detailed discussion of adaptive learning in linear models.

of algorithm (1) is given by the conditions for local asymptotic stability of $\bar{a}$ under an "associated differential equation"

$$\frac{da}{d\tau} = Ta - a,$$

known as the E-stability differential equation. Here τ denotes notional or virtual time. It is easily seen that the E-stability condition holds if and only if all eigenvalues of the matrix T have real parts less than one.[5]

Under constant-gain learning, the estimate a_t of a is updated according to

$$a_t = a_{t-1} + \gamma(y_t - a_{t-1}), \tag{2}$$

where $0 < \gamma \leq 1$ is the constant gain parameter. The only difference to (1) is constancy of the gain sequence. We now have

$$\begin{aligned} a_t &= a_{t-1} + \gamma\left(Ta_{t-1} + e_t - a_{t-1}\right), \text{ or} \\ a_t &= \left(\gamma T + (1-\gamma)I\right)a_{t-1} + \gamma e_t. \end{aligned}$$

This converges to a stationary stochastic process around the REE value $\bar{a} = 0$ (in deviation from mean form) provided all roots of the matrix $\gamma T + (1-\gamma)I$ lie inside the unit circle.

It is evident that stability under constant-gain learning depends on the value of γ, and we have the following result.

Proposition 1 *For a given $0 < \gamma \leq 1$, the stability condition is that the eigenvalues of T lie inside a circle of radius $1/\gamma$ and origin at $(1 - 1/\gamma, 0)$. This condition is therefore stricter for larger values of γ.*

Proof. The stability condition is that the roots of $\gamma(T + \gamma^{-1}(1-\gamma)I)$ lie inside the unit circle centered at the origin. Equivalently, the roots of $(T + \gamma^{-1}(1-\gamma)I)$ must lie inside a circle of radius $1/\gamma$ centered at the origin. Since the roots of $T + \gamma^{-1}(1-\gamma)I$ are the same as the roots of T plus $\gamma^{-1}(1-\gamma)$, this is equivalent to the condition given. ■

Note that the right edge of the circle is at $(1, 0)$ in the complex plane and that as $\gamma \to 0$ we obtain the standard (decreasing gain) E-stability condition that the real parts of all roots of T are less than than one. Looking at the other extreme $\gamma = 1$ gives the following:

[5] Throughout, we rule out boundary cases in which the real part of some eigenvalue of the T-map is one.

Corollary 2 *We have stability for all* $0 < \gamma \leq 1$ *if and only if all eigenvalues of* T *lie inside the unit circle.*

We remark that stability for all constant gains $0 < \gamma \leq 1$ is equivalent to a condition known as iterative E-stability, sometimes called "IE-stability." Iterative E-stability is said to hold when $T^j \to 0$ as $j \to \infty$.[6]

Note that when the stability condition holds, the parameter a_t converges to a stationary stochastic process that we can fully describe. This in turn induces a stationary stochastic process for $y_t = Ta_{t-1} + e_t$.

2.2 Application to Taylor Rules

Consider the standard forward-looking New Keynesian (NK) model,

$$\begin{aligned} x_t &= -\varphi(i_t - \pi^e_{t+1}) + x^e_{t+1} + g_t & (3) \\ \pi_t &= \lambda x_t + \beta \pi^e_{t+1} + u_t. & (4) \end{aligned}$$

For convenience we here assume that $(g_t, u_t)'$ are *iid*, so that the preceding technical results can be applied. Later we will consider cases with AR(1) shocks. We use x^e_{t+1} and π^e_{t+1} to denote expectations of π_{t+1} and x_{t+1}. Below we will be precise about the information sets available to agents when forming expectations and throughout the paper we will be exploring the implications of alternative assumptions.

Bullard and Mitra (2002) consider Taylor rules of various forms, including the "contemporaneous data" rule

$$i_t = \chi_\pi \pi_t + \chi_x x_t, \tag{5}$$

and the "contemporaneous expectations" rule

$$i_t = \chi_\pi \pi^e_t + \chi_x x^e_t. \tag{6}$$

In this section, when analyzing the contemporaneous expectations rule we follow Bullard and Mitra (2002) in assuming that all expectations are based on information at time $t-1$, i.e. $\pi^e_t = \hat{E}_{t-1}\pi_t$, $x^e_t = \hat{E}_{t-1}x_t$, $\pi^e_{t+1} = \hat{E}_{t-1}\pi_{t+1}$ and $x^e_{t+1} = \hat{E}_{t-1}x_{t+1}$. Since we have *iid* shocks, forecasts are based purely on the estimated intercept.

[6] In many models, iterative E-stability is known to be a necessary condition for the stability of "eductive" learning, e.g. see Evans and Guesnerie (1993).

Bullard and Mitra (2002) show that the determinacy and E-stability conditions are the same and are identical for both interest rate rules, and given by

$$\lambda(\chi_\pi - 1) + (1-\beta)\chi_x > 0. \tag{7}$$

They considered this finding important because of the argument by McCallum (1999) that it is not plausible that interest-rate rules can be conditioned on contemporaneous observations of endogenous aggregate variables like inflation and output, whereas they could plausibly be conditioned on central bank forecasts or "nowcasts" $\hat{E}_{t-1}\pi_t$, $\hat{E}_{t-1}x_t$.

We reconsider this issue from the vantage point of constant-gain learning. For the interest-rate rule (6) the model takes the form

$$y_t = M_0 y_t^e + M_1 y_{t+1}^e + P v_t, \tag{8}$$

where $y_t' = (x_t, \pi_t)$ and $v_t = (g_t, u_t)$, and where

$$M_0 = \begin{pmatrix} -\chi_x\varphi & -\chi_\pi\varphi \\ -\chi_x\varphi\lambda & -\chi_\pi\varphi\lambda \end{pmatrix} \text{ and } M_1 = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta+\varphi\lambda \end{pmatrix}, \tag{9}$$

and $P = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}$.

Since our shocks are *iid* the PLM is simply $y_t = a + e_t$, and the corresponding ALM is $y_t = (M_0+M_1)a+e_t$, where $e_t = Pv_t$. The usual E-stability condition is that the eigenvalues of $M_0 + M_1$ have real parts less than one, which leads to the condition (7). Applying Corollary 2, for convergence of constant-gain learning for all gains $0 < \gamma \leq 1$ we need that both eigenvalues of $M_0 + M_1$ lie inside the unit circle.

We investigate stability of constant-gain learning numerically, using the Woodford calibration of $\varphi^{-1} = 0.157$, $\lambda = 0.024$, $\beta = 0.99$. Setting $\chi_\pi = 1.5$, eigenvalues with real parts less than -1 arise for $\chi_x > 0.31$ and eigenvalues with real parts less than -9 arise for $\chi_x > 1.57$. This implies that when $\chi_\pi = 1.5$ and $\chi_x > 1.57$ the equilibrium is unstable under learning for constant gains $\gamma \geq 0.10$. This is perhaps not a significant practical concern since Taylor's recommended parameters are $\chi_\pi = 1.5$ and (based on the quarterly calibration of Woodford) $\chi_x = (0.5)/4 = 0.125$. However, it does show a previously unrecognized danger that arises under constant-gain learning if the Taylor rule has too strong a response to $\hat{E}_{t-1}x_t$, and this finding foreshadows instability problems that arise in more sophisticated rules discussed below.

Finally, we remark that the potential for instability under constant-gain learning arises specifically because of the necessity to use forecasts $\hat{E}_{t-1}y_t$. For the current data Taylor rule (5) it can be shown that the condition (7) guarantees stability under learning for all constant gains $0 < \gamma \leq 1$.[7]

3 Optimal Discretionary Monetary Policy

We now consider optimal policy under constant-gain learning, starting in this Section with optimal discretionary policy. We focus on homogeneous learning by private agents and the policy-maker. We initially restrict attention to the case of *iid* exogenous shocks, so that steady-state learning is appropriate. However, we also analyze the more general case, where the observable shocks follow AR(1) processes.

Consider the loss function

$$E_0 \sum_{t=0}^{\infty} \left[(\pi_t - \pi^*)^2 + \alpha_x (x_t - x^*)^2 + \alpha_i (i_t - i^*)^2 \right], \tag{10}$$

where π^*, x^* and i^* represent target values. For simplicity, we set $\pi^* = x^* = 0$. The weights $\alpha_x, \alpha_i > 0$ represent relative weights given by policy-makers to squared deviations of x_t and i_t from their targets, compared to squared deviations of π_t from its target.

The first-order condition (FOC) for discretionary optimal policy is

$$\lambda \pi_t + \alpha_x x_t - \alpha_i \varphi^{-1} (i_t - i^*) = 0. \tag{11}$$

We first consider a Taylor-type rule proposed by Duffy and Xiao (2007) and then discuss the expectations-based rule recommended by Evans and Honkapohja (2003b).

3.1 Taylor-type Optimal Rules

Duffy and Xiao (2007) propose using the equation (11) directly to obtain a Taylor-type rule that implements optimal discretionary policy. Solving the FOC for i_t yields the rule

$$i_t = \frac{\varphi \lambda}{\alpha_i} \pi_t + \frac{\varphi \alpha_x}{\alpha_i} x_t,$$

[7] The model now takes the form $y_t = M_1 \hat{E}_t y_{t+1} + P v_t$ and the required condition is the same as the determinacy condition.

where at this point we drop the term i^* since for brevity we are suppressing all intercepts. As discussed by Duffy and Xiao (2007), this is formally a contemporaneous-data Taylor rule. They show that for calibrated values of structural parameters and policy weights this leads to a determinate and E-stable equilibrium.

Because it is problematical that the Central Bank can observe contemporaneous output and inflation,[8] we instead examine the rule

$$i_t = \frac{\varphi\lambda}{\alpha_i}\hat{E}_{t-1}\pi_t + \frac{\varphi\alpha_x}{\alpha_i}\hat{E}_{t-1}x_t, \tag{12}$$

where the information set for the "nowcasts" $\pi_t^e = \hat{E}_{t-1}\pi_t$, $x_t^e = \hat{E}_{t-1}x_t$ is past endogenous variables and exogenous variables. This again leads to a model of the form (8) with coefficients (9), where $\chi_\pi = \varphi\lambda/\alpha_i$ and $\chi_x = \varphi\alpha_x/\alpha_i$. We assume that private agents and Central Banks estimate the same PLM, Since we are here assuming steady-state learning we also have $\pi_t^e = \hat{E}_{t-1}\pi_{t+1}$ and $x_t^e = \hat{E}_{t-1}x_{t+1}$.

We first note that for α_i sufficiently large the model under this Taylor-type rule will suffer from indeterminacy. This follows from the Bullard-Mitra result that the determinacy condition is (7), from which the critical value of α_i can be deduced. The condition for determinacy is

$$\alpha_i < \bar{\alpha}_i \equiv \varphi\lambda + (1-\beta)\lambda^{-1}\varphi\alpha_x. \tag{13}$$

If the central bank's desire to stabilize the interest rate is too strong, i.e. condition (13) is not met, then the central bank fails to adjust the interest rate sufficiently to ensure that the generalized Taylor principle (7) is satisfied. To assess this point numerically, we use the calibrated parameter values of Table 6.1 of Woodford (2003), with $\alpha_x = 0.048$, $\varphi = 1/0.157$, $\lambda = 0.024$, $\beta = 0.99$ and get approximately $\bar{\alpha}_i = 0.28$. Woodford's calibrated values of α_i are 0.077 or 0.233, where the latter value is from Woodford (1999). Thus the condition for determinacy does hold for these calibrations.

We next consider stability under learning. For the PLM $y_t = a + e_t$ we again get the ALM $y_t = (M_0 + M_1)a + e_t$, and

$$T \equiv M_0 + M_1 = \begin{pmatrix} 1-\alpha_i^{-1}\alpha_x\varphi^2 & \varphi - \alpha_i^{-1}\lambda\varphi^2 \\ \lambda - \alpha_i^{-1}\lambda\varphi^2\alpha_x & \beta + \lambda\varphi - \alpha_i^{-1}\lambda^2\varphi^2 \end{pmatrix}.$$

[8] An alternative would be to assume that agents and the policymaker sees the contemporaneous value of the exogenous shocks but not the contemporaneous values of x_t and π_t. This would not alter our results.

It can be shown that

$$\det(T) = \beta(1 - \alpha_i^{-1}\alpha_x\varphi^2).$$

For stability under all values $0 < \gamma < 1$ we need

$$\left|\beta(1 - \alpha_i^{-1}\alpha_x\varphi^2)\right| < 1,$$

and it is clear that for given β, α_x, φ this condition will not be satisfied for $\alpha_i > 0$ sufficiently small. Hence

Proposition 3 *Let* $\hat{\alpha}_i = \beta(1+\beta)^{-1}\alpha_x\varphi^2$*. For* $0 < \alpha_i \leq \hat{\alpha}_i$ *there exists* $0 < \hat{\gamma}(\beta, \varphi, \alpha_i, \alpha_x) < 1$ *such that the optimal discretionary Taylor-type rule (12) renders the REE unstable under learning for* $\hat{\gamma} < \gamma \leq 1$*.*

Thus, in addition to the indeterminacy problem for "large values" of α_i, the Taylor-type optimal rule suffers from a more serious problem of instability under constant-gain learning for "small values" of α_i. The source of this difficulty is the interaction of strong policy responses seen in equation (12) and a large gain parameter. This combination leads to cyclical overshooting of inflation and output gap. This is particularly evident as α_i tends to zero since in this case, e.g., a positive change in inflation expectations $\hat{E}_{t-1}\pi_t$ leads to large increase in i_t, which in turn leads to large negative changes in x_t and π_t via equations (3) and (4). The severity of this problem depends on the value of $\hat{\gamma}$ in Proposition 3. Ideally, stability would hold for all $0 < \gamma \leq 1$, but if $\hat{\gamma}$ is high the problem might not be a major concern.

We investigate the magnitude of $\hat{\gamma}$ numerically by computing the eigenvalues of $\gamma T + (1-\gamma)I$. As an example, for the Woodford calibration $\beta = 0.99, \varphi = 1/0.157, \lambda = 0.024$, we find that with $\alpha_x = 0.048$ and $\alpha_i = 0.077$, the critical value $\hat{\gamma} \approx 0.04$. Since estimates in the macro literature suggest gains in the range 0.02 to 0.06, this indicates that optimal Taylor-type rules may not be stable under learning.[9] The source of the problem is that with low α_i the implied weights on $\hat{E}_{t-1}\pi_t$ and especially $\hat{E}_{t-1}x_t$ are very high. Under constant-gain learning this can lead to instability unless the gain parameter is very low. We will demonstrate later that this problem is avoided by using a suitable expectations-based optimal rule.

[9] Milani (2007b) considers a setting in which agents switch between decreasing gain and constant gain estimators, depending on recent average mean-square errors. In the constant-gain regime the estimated gains are even higher, around 0.07 to 0.08.

We next consider the case in which the exogenous shocks are $AR(1)$ processes. In this setting various information assumptions have been used in the literature. Perhaps the most common assumption is that agents see current and lagged exogenous variables and lagged, but not current, endogenous variables. Expectations under this assumption are denoted as $\hat{E}_t\pi_t$, $\hat{E}_t x_t$, $\hat{E}_t\pi_{t+1}$ and $\hat{E}_t x_{t+1}$. An alternative would replace these by $\hat{E}_{t-1}\pi_t$, $\hat{E}_{t-1}x_t$, $\hat{E}_{t-1}\pi_{t+1}$ and $\hat{E}_{t-1}x_{t+1}$, indicating that agents only see lagged information.[10] Whether agents see current or only lagged exogenous shocks is not particularly crucial and does not affect our main results. Consequently, we follow the most common assumption that expectations are specified as $\hat{E}_t\pi_t$, $\hat{E}_t x_t$, $\hat{E}_t\pi_{t+1}$ and $\hat{E}_t x_{t+1}$.[11] In contrast, as we have already seen, whether agents and policy-makers are able to see current endogenous variables is an important issue for stability under learning. This is why we use the term operationality to indicate an interest rate rule that does not depend on current endogenous variables.

We now assume that the exogenous shocks g_t and u_t follow $AR(1)$ processes, i.e.

$$g_t = \mu g_{t-1} + \tilde{g}_t \text{ and } u_t = \rho u_{t-1} + \tilde{u}_t$$

where $0 < |\mu|, |\rho| < 1$ and $\tilde{g}_t \sim iid(0, \sigma_g^2)$, $\tilde{u}_t \sim iid(0, \sigma_u^2)$ are independent white noise processes. We write this in vector form as

$$v_t = F v_{t-1} + \tilde{v}_t.$$

Under the current assumptions, the PLM of the agents is

$$y_t = a + c v_t,$$

and the forecasts are now $\hat{E}_t y_t = a + c v_t$ and $\hat{E}_t y_{t+1} = a + cF v_t$. Using the general model (8), the ALM is

$$y_t = (M_0 + M_1)a + (M_0 c + M_1 cF + P)v_t,$$

[10] A third alternative, which is occasionally used in the literature, allows agents to see the contemporaneous values of endogenous variables. However, this assumption runs against the requirement of operationality that we want to emphasize here.

[11] The standard assumption under RE is that agents have contemporaneous information. Our information assumption takes account of the operationality critique, but nonetheless allows for the possibility of convergence under learning to the REE.

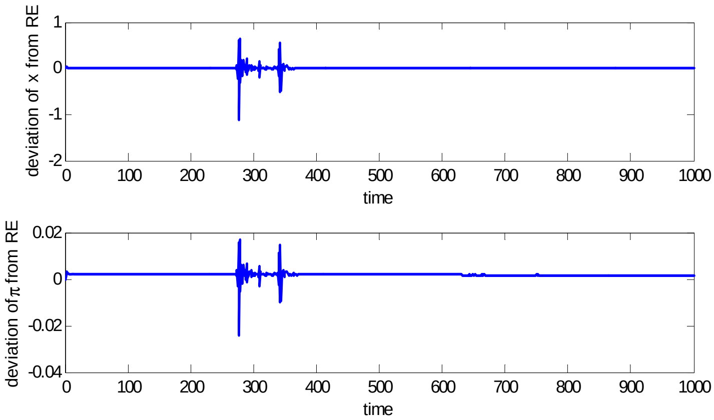

Figure 1: Stability of optimal Taylor-type rule with $\gamma = 0.02$.

and the E-stability conditions are that all eigenvalues of the matrices M_0+M_1 and $I \otimes M_0 + F' \otimes M_1$ have real parts less than one. Here $\otimes$ denotes the Kronecker product of two matrices.[12]

To examine stability under constant-gain learning, we simulate the model under constant-gain recursive least squares (RLS) estimation of the PLM parameters a and c.[13] Under constant-gain least squares agents discount old data geometrically at the rate $1-\gamma$. Let a_t, c_t denote the estimates based on data through $t-1$. For the recursive formulation of (constant-gain) least squares see the Appendix. Given these estimates, expectations are formed as $y_t^e = \hat{E}_t y_t = a_t + c_t v_t$ and $y_{t+1}^e = \hat{E}_t y_{t+1} = a_t + c_t F v_t$ and the temporary equilibrium is then given by (8) with these expectations.

We use the previous values for the structural parameters and also set $\mu = \rho = 0.8$. Simulations of the system indicate instability under constant-gain RLS learning for gain parameters at or in excess of 0.024. Thus, with regressors that include exogenous AR(1) observables instability arises at even lower gain values than in the case of steady state learning. Figures 1 and 2 illustrate the evolution of parameters over time under constant-gain RLS learning with the Taylor-type rule (12) in stable and unstable cases.

[12] In the case of lagged information the PLM is specified as $y_t = a + cv_{t-1} + \eta_t$ and the ALM is then $y_t = (M_0 + M_1)a + (M_0 c + M_1 cF + PF)v_{t-1} + P\tilde{v}_t$.

[13] The RLS formulae are given in the Appendix.

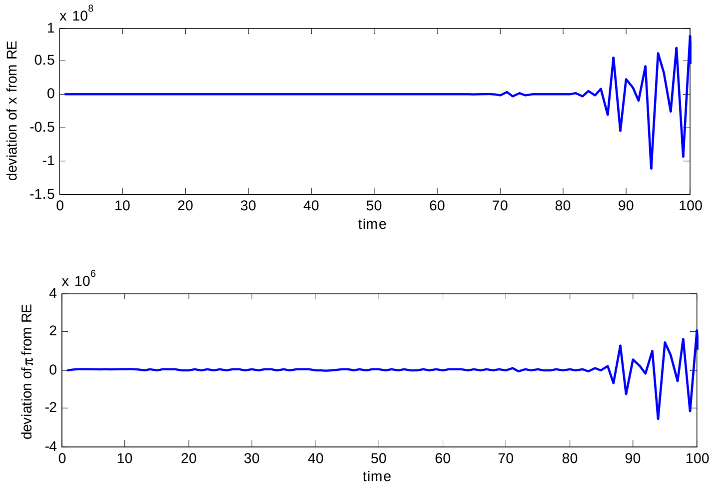

Figure 2: Instability of optimal Taylor-type rule with $\gamma = 0.04$.

3.2 Expectations-Based Optimal Rules

Assume now that at time t the exogenous shocks g_t, u_t and private-sector expectations $\hat{E}_t\pi_{t+1}$ and $\hat{E}_t x_{t+1}$ are observed by the Central Bank. The expectations-based (EB) rule is constructed so that it exactly implements (11), the FOC under discretion, even outside an REE for given expectations, as suggested by Evans and Honkapohja (2003b). To obtain the rule, we combine (3), (4) and (11), and solve for i_t in terms of the exogenous shocks and the expectations.

The resulting EB-rule is

$$\begin{aligned} i_t &= \frac{(\alpha_x+\lambda^2)\varphi}{\alpha_i+(\alpha_x+\lambda^2)\varphi^2}\hat{E}_t x_{t+1} + \frac{\beta\lambda\varphi+(\alpha_x+\lambda^2)\varphi^2}{\alpha_i+(\alpha_x+\lambda^2)\varphi^2}\hat{E}_t\pi_{t+1} \\ &\quad + \frac{(\alpha_x+\lambda^2)\varphi}{\alpha_i+(\alpha_x+\lambda^2)\varphi^2}g_t + \frac{\lambda\varphi}{\alpha_i+(\alpha_x+\lambda^2)\varphi^2}u_t. \end{aligned}$$

This leads to a reduced form

$$y_t = M\hat{E}_t y_{t+1} + P v_t. \tag{14}$$

Determinacy of the REE corresponding to optimal discretionary monetary policy requires that M has both eigenvalues inside the unit circle.[14] We again have the condition $\alpha_i < \bar{\alpha}_i$, where $\bar{\alpha}_i$ is given by (13).

For stability under learning, first consider the case where the exogenous shocks v_t are *iid* and agents use steady state learning under constant gain. For this reduced form the PLM $y_t = a + e_t$ gives the ALM $y_t = Ma + e_t$ (where $e_t = Pv_t$), as discussed in Section 2.1. Thus $T = M$ and there is a very close connection between determinacy and stability under learning. We have:

Proposition 4 *Assume $\alpha_i < \bar{\alpha}_i$ and the shocks are iid. Then the EB-rule, which implements the FOC, yields a reduced form that is stable under steady-state learning for all constant-gain rules* $0 < \gamma \leq 1$.

Provided $\alpha_i < \bar{\alpha}_i$, so that determinate optimal policy is possible, the EB-optimal rule will successfully implement the optimal REE: under decreasing gain learning there will be convergence to the REE, and under small constant gain it will converge to a stochastic process near the optimal REE. Furthermore, for all constant gains $0 < \gamma \leq 1$ there will be convergence to a stationary process centered at the optimal REE.

Second, we examine numerically the case of AR(1) shocks with (constant-gain) RLS learning. For the Woodford calibration $\beta = 0.99, \varphi = 1/0.157, \lambda = 0.024$, $\alpha_x = 0.048$ and $\alpha_i = 0.077$ (and $\rho = \mu = 0.8$) we find that learning converges for gain values at or below $\gamma = 0.925$. In other words, the expectations-based optimal discretionary rule is quite robustly stable under learning. We also make a technical remark that when the agents have to run genuine regressions, as in the current case, then the IE-stability condition does not imply convergence of constant-gain learning for all $0 < \gamma \leq 1$. However, we see that stability does hold even for γ quite close to one.

4 Optimal Policy with Commitment

For brevity, in the remainder of the paper we assume that $\alpha_i = 0$, i.e. the Central Bank does not have an interest rate stabilization objective.[15] Given the model (3)-(4) and the loss function (10) with $\alpha_i = 0$, it is well-known

[14] Equivalently we need $|\mathrm{tr}(M)| < 1 + \det(M)$ and $|\det(M)| < 1$.

[15] See Duffy and Xiao (2007) for the extension to the case where the Central Bank also has an interest-rate stabilization motive.

that optimal monetary policy under commitment (in a timeless perspective) is characterized by the condition[16]

$$\lambda \pi_t = -\alpha_x (x_t - x_{t-1}), \tag{15}$$

which is often called the optimal targeting rule. It can be shown that the optimal rational expectations equilibrium of interest has the form

$$\begin{aligned} x_t &= b_x x_{t-1} + c_x u_t \\ \pi_t &= b_\pi x_{t-1} + c_\pi u_t, \end{aligned}$$

where we choose the unique $0 < b_x < 1$ that solves the equation $\beta b_x^2 - (1+\beta+\lambda^2/\alpha_x)b_x + 1 = 0$ and $b_\pi = (\alpha_x/\lambda)(1-b_x)$, $c_x = -[\lambda + \beta b_\pi + (1-\beta\rho)(\alpha_x/\lambda)]^{-1}$ and $c_\pi = -(\alpha_x/\lambda)c_x$.

Different optimal reaction functions that implement the optimal targeting rule (15) have been proposed in the literature. Under rational expectations one obtains the fundamentals-based reaction function

$$i_t = \psi_x x_{t-1} + \psi_g g_t + \psi_u u_t, \tag{16}$$

where

$$\begin{aligned} \psi_x &= b_x[\varphi^{-1}(b_x - 1) + b_\pi] \\ \psi_g &= \varphi^{-1} \\ \psi_u &= [b_\pi + \varphi^{-1}(b_x + \rho - 1)]c_x + c_\pi \rho. \end{aligned}$$

Evans and Honkapohja (2006) show that the reaction function (16) often leads to indeterminacy and always leads to expectational instability. They propose instead the expectations-based reaction function

$$i_t = \delta_L x_{t-1} + \delta_\pi \hat{E}_t \pi_{t+1} + \delta_x \hat{E}_t x_{t+1} + \delta_g g_t + \delta_u u_t, \tag{17}$$

where the coefficients are[17]

$$\delta_L = \frac{-\alpha_x}{\varphi(\alpha_x + \lambda^2)}, \; \delta_\pi = 1 + \frac{\lambda\beta}{\varphi(\alpha_x + \lambda^2)}, \; \delta_x = \delta_g = \varphi^{-1}, \; \delta_u = \frac{\lambda}{\varphi(\alpha_x + \lambda^2)}.$$

[16] See e.g. Clarida, Gali, and Gertler (1999) and Woodford (1999). For the exposition, we follow Evans and Honkapohja (2006).

[17] In the discretionary case with $\alpha_i = 0$ the same coefficients would obtain, except that $\delta_L = 0$.

Under the interest-rate reaction rule (17) the reduced form model is of the form

$$y_t = M_1 \hat{E}_t y_{t+1} + N y_{t-1} + P v_t,$$

with $y_t' = (x_t, \pi_t)$ and $v_t' = (g_t, u_t)$. The corresponding REE takes the form $y_t = \bar{b} y_{t-1} + \bar{c} v_t$. Evans and Honkapohja (2006) show that the optimal expectations-based reaction function (17) delivers a determinate and E-stable optimal REE for all values of the parameters. It is therefore clearly preferred to the fundamentals-based rule (16).

In connection with constant gain learning we have the following partial result:[18]

Proposition 5 *The EB-rule under commitment (17) yields a reduced form for which the eigenvalues of the $T-$map are inside the unit circle for all values of the structural parameters.*

This result is partial in the sense that the eigenvalues condition is no longer sufficient for stability of constant-gain learning for all $0 < \gamma \leq 1$. This is because in the model the regressors include exogenous and lagged endogenous variables.

We now examine numerically the performance of constant-gain RLS learning under the expectations-based optimal rule with commitment. Using Woodford's parameter values (but with $\alpha_i = 0$), we find that constant-gain RLS learning converges for values of the gain parameter below $\hat{\gamma} \approx 0.25$. The inclusion of a lagged variable among the regressors appears to have a significant effect on stability of learning for large gains. However, the rule is still robust for all plausible values of the gain parameter.

As noted above, the Duffy and Xiao (2007) formulation under commitment breaks down when $\alpha_i = 0$ (as it does in the discretionary case). One might investigate numerically the performance of the Duffy-Xiao rule under constant-gain RLS for calibrated values of α_i. Based on the results in the discretionary case, we are not optimistic about robust learning stability of the Duffy-Xiao rule with commitment.

[18] See the Appendix for a proof.

5 Alternative Rules for Optimal Policy under Commitment

5.1 Svensson-Woodford Rule

Given that the fundamentals-based optimal rules (without interest rate stabilization) lead to problems of indeterminacy and learning instability, Svensson and Woodford (2005) suggest a modification to such a rule. In this rule the fundamentals-based rule (16) is complemented with a term that is based on the commitment optimality condition. We again assume that contemporaneous data are not available to the policy-maker, so that current values of inflation π_t and the output gap x_t are replaced by their nowcasts $\hat{E}_t\pi_t$ and $\hat{E}_t x_t$. This results in the interest rate rule

$$i_t = \psi_x x_{t-1} + \psi_g g_t + \psi_u u_t + \theta[\hat{E}_t\pi_t + \frac{\alpha_x}{\lambda}(\hat{E}_t x_t - x_{t-1})], \quad (18)$$

where $\theta > 0$.

The full model is now given (3), (4) and (18). By substituting (18) into (3) this model can be reduced to a bivariate model of the form

$$y_t = M_0\hat{E}_t y_t + M_1\hat{E}_t y_{t+1} + N y_{t-1} + P v_t, \quad (19)$$

where the information set in the forecasts and nowcasts includes current values of the exogenous shocks but not of the endogenous variables. It is also assumed for convenience that $v_t = F v_{t-1} + \tilde{v}_t$ is a known, stationary process. The coefficient matrices are

$$\begin{aligned} M_0 &= \begin{pmatrix} -\varphi\alpha_x\theta\lambda^{-1} & -\varphi\theta \\ -\varphi\alpha_x\theta & -\varphi\theta\lambda \end{pmatrix}, M_1 = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta+\lambda\varphi \end{pmatrix} \\ N &= \begin{pmatrix} -\varphi\psi_x + \varphi\alpha_x\theta\lambda^{-1} & 0 \\ -\lambda\varphi\psi_x + \varphi\alpha_x\theta & 0 \end{pmatrix}, P = \begin{pmatrix} 0 & -\varphi\psi_u \\ 0 & 1-\lambda\varphi\psi_u \end{pmatrix}. \end{aligned}$$

The PLM has the form

$$y_t = a + b y_{t-1} + c v_t$$

and the $T-$mapping is

$$T(a,b,c) = ((M_0 + M_1(I+b))a, M_1 b^2 + M_0 b + N, M_0 c + M_1(bc + cF) + P).$$

The usual E-stability conditions are stated in terms of the eigenvalues of the derivative matrices

$$\begin{aligned} DT_a &= M_0 + M_1(I + \bar{b}) \\ DT_b &= \bar{b}' \otimes M_1 + I \otimes M_1\bar{b} + I \otimes M_0 \\ DT_c &= F' \otimes M_1 + I \otimes M_1\bar{b} + I \otimes M_0, \end{aligned}$$

where $\otimes$ is the Kronecker product and $\bar{b}$ is the RE value of b.

We compute numerically the E-stability eigenvalues for the Woodford calibration with $\alpha_x = 0.048$ and $\theta = 1$.[19] For this case the eigenvalues of DT_a are -9.570 and 0.99, while the eigenvalues of DT_b are -10.605, -9.672, 0.878 and -0.0118. However, $\theta = 1$ is very close to the lower bound on θ needed for E-stability (since one root of DT_a is almost one), and the eigenvalues are sensitive to the value of θ. For example, for $\theta = 1.5$ the eigenvalues of DT_a are -15.975 and 0.949, while the eigenvalues of DT_b are -17.059, -16.082, 0.842 and -0.0110. It is seen that large negative eigenvalues appear.

The calculation of the E-stability eigenvalues suggests that the interest rate rule (18) can be subject to instability if learning is based on constant gain. We now examine numerically the performance of the rule (18) under different values of the constant gain using the Woodford calibrated values of the model parameters and $\theta = 1.5$. Numerical simulations show that under the interest rate rule (18) constant-gain RLS learning becomes unstable for values of γ at 0.019 or higher.

We also examine numerically the sensitivity of the stability upper bound on γ for different values of α_x, i.e. the degree of flexibility of inflation targeting. Table 1 gives the approximate highest value $\hat{\gamma}$ of the gain for which stability under constant-gain learning obtains.

α_x	0.01	0.02	0.03	0.04	0.05	0.06	0.08	0.1
$\hat{\gamma}$	0.185	0.06	0.035	0.02	0.018	0.014	0.009	0.007

Table 1: Critical values of γ for stability, Svensson-Woodford rule

It is seen from Table 1 that robust learning stability of the Svensson-Woodford hybrid rule is very sensitive to the degree of flexibility in inflation targeting. Robust stability obtains only when the Central Bank is an “inflation hawk”.

[19] We remark that the eigenvalues of the same model but with contemporaneous data available would not deliver large negative eigenvalues in the E-stability calculation for this parameterization.

5.2 McCallum-Nelson Rule

McCallum and Nelson (2004) propose a different rule that approximates optimal interest-rate policy in the timeless-perspective sense. They suggest that the interest rate be raised above inflation whenever the timeless-perspective optimality condition is above zero.

Their rule performs well if y_t is observable, but as McCallum and Nelson (2004) themselves point out, such a rule would be subject to the operationality problem that we have encountered several times: it presupposes that contemporaneous data on inflation and the output gap is available. One way to overcome this problem is to replace unknown contemporaneous data by nowcasts of the variables. In this case the interest-rate rule becomes

$$i_t = \hat{E}_t\pi_t + \theta[\hat{E}_t\pi_t + \frac{\alpha_x}{\lambda}(\hat{E}_t x_t - x_{t-1})]. \tag{20}$$

Under RE this rule approximates optimal policy under (timeless perspective) commitment, provided $\theta > 0$ is large.

The model is then given by equations (3), (4) and (20). The model can be reduced to a bivariate model of the form (19), where the coefficient matrices are

$$\begin{aligned} M_0 &= \begin{pmatrix} -\varphi\alpha_x\lambda^{-1} & -\varphi(1+\theta) \\ -\varphi\alpha_x & -\lambda\varphi(1+\theta) \end{pmatrix}, \ M_1 = \begin{pmatrix} 1 & \varphi \\ \lambda & \beta+\lambda\varphi \end{pmatrix} \\ N &= \begin{pmatrix} \varphi\alpha_x\lambda^{-1} & 0 \\ \varphi\alpha_x & 0 \end{pmatrix}, \ P = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}. \end{aligned}$$

Using the same parameter values as above in the case of the Svensson-Woodford hybrid rule, with $\alpha_x = 0.048$, we obtain that for $\theta = 1$ the eigenvalues of DT_a are -9.719 and 0.869, while the eigenvalues of DT_b are -10.780, -9.833, 0.750 and -0.213. For $\theta = 1.5$ the eigenvalues of DT_a are -9.997 and 0.841, while the eigenvalues of DT_b are -11.087, -10.138, 0.701 and -0.213.

The results are very sensitive to α_x. For $\alpha_x = 0.1$, we obtain that for $\theta = 1$ the eigenvalues of DT_a are -22.954 and 0.912, while the eigenvalues of DT_b are -24.042, -23.033, 0.835 and -0.143.

It can be seen that the problem of large negative eigenvalues appears with this rule, so that the potential of instability under constant-gain learning exists. Using the Woodford calibration (including $\alpha_x = 0.048$) and choosing $\theta = 1.5$, we find that constant-gain RLS learning becomes unstable for values of the gain at or above 0.029.

We again examine numerically the sensitivity of the stability upper bound on γ for different values of α_x, i.e. the degree of flexibility of inflation targeting. Table 2 gives the approximate highest value $\hat{\gamma}$ of the gain for which stability under constant-gain learning obtains.

α_x	0.01	0.02	0.03	0.04	0.05	0.06	0.08	0.1
$\hat{\gamma}$	0.395	0.107	0.058	0.037	0.026	0.021	0.015	0.01

Table 2: Critical values of γ for stability, McCallum-Nelson rule

Comparing Tables 1 and 2, it is seen that the stability performance of the McCallum-Nelson rule (20) is somewhat better than that of the hybrid rule (18) for the same parameter values. However, it is still the case that the McCallum-Nelson rule is not robust for many plausible values of the gain parameter.

We remark that McCallum and Nelson (2004) suggest that a preferable alternative to (20) is to use forward expectations in place of nowcasts, since this delivers superior results under rational expectations. In this case, the model has no lagged endogenous variables, i.e. $N = 0$ in (19). This case has been analyzed numerically in Evans and Honkapohja (2003a) and Evans and Honkapohja (2006). In this formulation, large negative eigenvalues no longer arise. However, we found that determinacy and E-stability require a small value of the parameter θ, while for small values of θ the welfare losses for optimal policy can be significant.

6 Conclusions

A lot of recent applied research on learning and monetary policy has emphasized discounted (constant-gain) least-squares learning by private agents. We have examined the stability performance of various operational interest-rate rules under constant-gain learning for different values of the gain parameter. Since estimates of the gain parameter tend to be in the range of 0.02 to 0.06 for quarterly macro data, ideally there should convergence of learning for gain parameters up to 0.1. Based on this criterion, we have found that many proposed interest-rate rules are not robustly stable under learning in this sense. An exception to this finding is the class of expectations-based optimal rules in which the interest rate feeds directly back on private expectations in an appropriate way.

A Appendix

Constant-gain RLS Algorithm

Suppose the economy is described in terms of a multivariate linear model, which includes possible dependence on lagged endogenous variables.

Under least-squares learning agents have the PLM

$$y_t = a + by_{t-1} + cv_t + e_t, \tag{21}$$

where a, b and c denote parameters to be estimated. Here y_t is a $p \times 1$ vector of endogenous variables. v_t is $k \times 1$ vector of observable exogenous variables, and e_t is a vector of white noise shocks. If the model does not have lagged endogenous variables, then the term by_{t-1} is omitted.

At time t agents compute their forecasts using (21) with the estimated values (a_t, b_t, c_t) based on data up to period $t-1$. Constant-gain RLS takes the form

$$\begin{aligned} \xi_t &= \xi_{t-1} + \gamma R_t^{-1} Z_{t-1}(y_{t-1} - \xi_{t-1}' Z_{t-1})', \\ R_t &= R_{t-1} + \gamma (Z_{t-1} Z_{t-1}' - R_{t-1}) \end{aligned}$$

where $\xi_t' = (a_t, b_t, c_t)$, $Z_t' = (1, y_{t-1}', v_t')$ and $1 > \gamma > 0$. The algorithm starts at $t = 1$ with a complement of initial conditions. We remark that the only difference from standard RLS is that the latter assumes a decreasing gain $\gamma_t = 1/t$.[20]

Proof of Proposition 5

We now sketch a proof of Proposition 5. We examine the formulas given in equations (A7)-(A9) on p. 36 of Evans and Honkapohja (2006). Two of the eigenvalues of DT_b are 0, while the remaining eigenvalues are those of the matrix

$$K_b = \begin{pmatrix} \frac{-\lambda\beta b_\pi}{\alpha_x+\lambda^2} & \frac{-\lambda\beta b_x}{\alpha_x+\lambda^2} \\ \frac{\alpha_x\beta b_\pi}{\alpha_x+\lambda^2} & \frac{\alpha_x\beta b_x}{\alpha_x+\lambda^2} \end{pmatrix}$$

[20] The formal analysis of recursive least squares (RLS) learning in linear multivariate models is developed e.g. in Evans and Honkapohja (1998) and Chapter 10 of Evans and Honkapohja (2001).

The eigenvalues of K_b are 0 and $-1 < \frac{\alpha_x \beta(2b_x - 1)}{\alpha_x + \lambda^2} < 1$. Likewise, two of the eigenvalues of DT_c are 0 while the other two eigenvalues of those of the matrix

$$K_c = \begin{pmatrix} \frac{-\lambda\beta b_\pi}{\alpha_x+\lambda^2} & \frac{-\lambda\beta\rho}{\alpha_x+\lambda^2} \\ \frac{\alpha_x\beta b_\pi}{\alpha_x+\lambda^2} & \frac{\alpha_x\beta\rho}{\alpha_x+\lambda^2} \end{pmatrix},$$

and the eigenvalues of K_c are 0 and $\frac{\alpha_x \beta(b_x - 1 + \rho)}{\alpha_x + \lambda^2}$, which is inside the unit circle unless ρ is negative and large in magnitude. Finally,

$$DT_a = \begin{pmatrix} \frac{-\lambda\beta b_\pi}{\alpha_x+\lambda^2} & \frac{-\lambda\beta}{\alpha_x+\lambda^2} \\ \frac{\alpha_x\beta b_\pi}{\alpha_x+\lambda^2} & \frac{\alpha_x\beta}{\alpha_x+\lambda^2} \end{pmatrix}$$

and its eigenvalues are 0 and $0 < \frac{\alpha_x \beta b_x}{\alpha_x + \lambda^2} < 1$.

References

Bernanke, B., and M. Woodford (eds.) (2005): *The Inflation-Targeting Debate*. University of Chicago Press, Chicago, USA.

Branch, W., and G. W. Evans (2006): "A Simple Recursive Forecasting Model," *Economic Letters*, 91, 158–166.

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

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

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

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

Evans, G. W., and R. Guesnerie (1993): "Rationalizability, Strong Rationality, and Expectational Stability," *Games and Economic Behaviour*, 5, 632–646.

Evans, G. W., and S. Honkapohja (1998): "Economic Dynamics with Learning: New Stability Results," *Review of Economic Studies*, 65, 23–44.

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

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

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

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

——— (2007): "Expectations, Learning and Monetary Policy: An Overview of Recent Research," mimeo.

Evans, G. W., S. Honkapohja, and N. Williams (2006): "Generalized Stochastic Gradient Learning," mimeo.

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

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

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

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

——— (2007b): "Learning and Time-varying Macroeconomic Volatility," Working paper, University of California, Irvine.

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

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

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

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

Woodford, M. (1999): "Optimal Monetary Policy Inertia," Working paper, NBER WP7261.

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