

Learning and Macroeconomics*

George W. Evans

University of Oregon and University of Saint Andrews

email: gevans@uoregon.edu

Seppo Honkapohja

Bank of Finland and University of Cambridge

email: seppo.honkapohja@bof.fi

July 11, 2008

Abstract

Expectations play a central role in modern macroeconomic theories. The econometric learning approach models economic agents as forming expectations by estimating and updating forecasting models in real time. The learning approach provides a stability test for rational expectations and a selection criterion in models with multiple equilibria. In addition, learning provides new dynamics if older data is discounted, models are misspecified or agents choose between competing models. This paper describes the E-stability principle and the stochastic approximation tools used to assess equilibria under learning. Applications of learning to a number of areas are reviewed, including the design of monetary and fiscal policy, business cycles, self-fulfilling prophecies, hyperinflation, liquidity traps, and asset prices.

Key words: E-stability, least-squares, stochastic approximation, persistent learning dynamics, business cycles, monetary policy, asset prices, sunspots.

JEL classification: E32, D83, D84, C62

*Corresponding author: George W. Evans, Department of Economics, 1285 University of Oregon, Eugene, OR 97403-1285, USA. email: gevans@uoregon.edu.

1 Introduction

Expectations play a central role in modern macroeconomics. Economic agents are assumed to be dynamic optimizers whose current economic decisions are the first stage of a dynamic plan. Thus households must be concerned with expected future incomes, employment, inflation, and taxes, as well as the expected trajectory of the stock market and the housing market. Firms must forecast the level of future product demand, wage costs, productivity levels, and foreign exchange rates. Monetary and fiscal policy-makers must forecast inflation and aggregate economic activity and consider both the direct impact of their policies and the indirect effect of policy rules on private-sector expectations.

Macroeconomic models can be summarized as a reduced-form multivariate dynamic system

$$y_t = F(y_{t-1}, y_{t+1}^e, w_t), \quad (1)$$

where y_t is a vector of endogenous variables and w_t is a vector of stochastic exogenous variables. Typically, w_t is assumed to follow a stationary stochastic process such as a finite-dimensional vector autoregression. Crucially, y_t depends not only on the state of the system, captured by the exogenous variables and lagged endogenous variables, w_t and y_{t-1} , but also on expectations of future endogenous variables, y_{t+1}^e . The precise information set available to economic agents for forming expectations will depend on the specific model, and in some cases y_t will depend also on “forecasts” of contemporaneous variables.

Since the work of Muth (1961), Lucas (1972), and Sargent (1973), the benchmark model of expectation formation in macroeconomics has been rational expectations (RE). This posits, for both private agents and policy-makers, that expectations are equal to the true statistical conditional expectations of the unknown random variables. RE is clearly a very strong assumption, since it implicitly assumes knowledge of the correct form of the model, knowledge of all parameters, and knowledge that other agents are rational, as well as the knowledge that other agents know that other agents are rational, etc.

The “learning theory” approach in macroeconomics argues that although RE is the natural benchmark, it is implausibly strong. We need a more realistic model of rationality, which may, however, be consistent with agents eventually learning to have RE. A natural criterion for a model of rational-

ity is the “cognitive consistency principle,” that economic agents should be assumed to be about as smart as (good) economists. This still leaves open various possibilities, since we could choose to model households and firms like economic theorists or, alternatively, model them like econometricians.¹ The adaptive or econometric learning approach, which will here be our principal focus, takes the latter viewpoint, arguing that economists, when they forecast future economic aggregates, usually do so using time-series econometric techniques. This seems particularly natural since neither private agents nor economists at central banks know the true model. Instead economists formulate and estimate models. These models are re-estimated and possibly reformulated as new data become available. Economists themselves engage in processes of learning about the economy.²

The econometric learning approach to expectation formation leads to several distinct roles for learning in macroeconomics. Closest to the RE view, econometric learning can be viewed as a stability test for RE equilibria (REE): under what circumstances will least squares (LS) or closely related econometric learning schemes converge asymptotically to RE? The stability analysis can also be used as a selection device when there are multiple REE. This is of particular interest when the REE include “sunspot equilibria” or cycles that can be viewed as self-fulfilling prophecies. Can economic agents using econometric forecasting rules, updated over time in accordance with LS, converge over time to non-fundamental solutions like sunspot equilibria? These questions can be examined using the expectational stability (E-stability) tool. According to the E-stability principle the local stability of an REE under LS-type learning rules can be determined from using a differential equation that is often straightforward to compute.

The econometric learning approach also generates additional insights for macroeconomic theory and economic policy. For example, the reality that econometricians sometimes use misspecified models suggests that we should consider agents using misspecified econometric forecasting models. There is then the possibility of convergence to “restricted perceptions equilibria” in which the agents are doing the best they can, given their misspecified models. Another consideration is that, even if agents converge asymptotically to

¹The “eductive” approach models agents like economic theorists with common knowledge of the economic structure and of the rationality of other agents. See Guesnerie (2005) for the theory and applications.

²For recent survey articles and overviews see Evans and Honkapohja (1999), Marimon (1997) and Sargent (2008).

an REE, the economy will deviate from this equilibrium during the learning transition. Furthermore, if agents use “constant-gain” (“discounted”) LS, which weights more recent data more heavily, then convergence will be to a stochastic process near the REE, rather than to the REE itself. In some cases this can have major implications in applications or for economic policy. Additional learning dynamics arise when one allows for heterogeneous expectations, due to agents using either different learning rules or differing forecasting models.

In this paper we first survey the main tools of macroeconomic learning theory, in Section 2, and then consider a range of applications in Section 3. The applications examined in detail include monetary policy, business cycles, and asset prices. The section on learning and monetary policy describes the implications for optimal policy if agents use constant-gain learning, recent empirical work on inflation dynamics, and results on the stability of alternative interest-rate rules in New Keynesian models, including both Taylor-type rules and rules aiming to implement optimal policy. In examining business cycles under learning, we first consider stability of the REE under LS learning, in the standard RBC model, and then examine stability of sunspot equilibria in RBC-type models with distortions. We then turn to stability of sunspot equilibria in New Keynesian and in cash-in-advance models. Next, we summarize the implications of learning in models that analyze exceptional phenomena, namely hyperinflation and liquidity traps. A final section reviews some applications to asset pricing, specifically to stock-price returns and to exchange rates.

2 Theory and Techniques

2.1 Least-Squares Learning and E-stability

We develop the basic ideas of econometric learning using a simple linear model

$$p_t = \mu + \alpha E_{t-1}^* p_t + \delta' w_{t-1} + \eta_t. \quad (2)$$

Here p_t is a scalar endogenous variable, w_{t-1} is a vector of exogenous observable variables and η_t is an unobservable random shock. The key assumption is that the expectations of economic agents, $E_{t-1}^* p_t$, are not necessarily rational since the agents do not know all the structural parameters. Expectations are instead formed as forecasts from an estimated model and observations

w_{t-1} . The parameters of the forecasting model are estimated using past data and updated over time. For simplicity, all agents are assumed to have the same expectations. (We discuss heterogenous expectations below.)

As a benchmark we note that for model (2) the unique REE is

$$p_t = \bar{a} + \bar{b}'w_{t-1} + \eta_t, \bar{a} = (1 - \alpha)^{-1}\mu, \bar{b} = (1 - \alpha)^{-1}\delta,$$

as is easily verified by applying the method of undetermined coefficients with the functional form $p_t = a + b'w_{t-1} + \eta_t$. Two well-known economic examples lead to reduced-form model (2).

Example 1: (Lucas aggregate supply model). A simple version of the ‘‘Lucas islands’’ model, presented in Lucas (1973), consists of the aggregate supply function

$$q_t = \bar{q} + \pi(p_t - p_t^e) + \zeta_t,$$

where $\pi > 0$, and the aggregate demand function

$$m_t + v_t = p_t + q_t,$$

where v_t is a velocity shock. We assume that velocity depends in part on exogenous observables w_{t-1} so that

$$v_t = \mu + \gamma'w_{t-1} + \xi_t,$$

and that money supply follows the policy rule

$$m_t = \bar{m} + u_t + \rho'w_{t-1}.$$

Thus, m_t responds to past shocks to velocity. Here u_t, ξ_t and ζ_t are white noise shocks. The reduced form of the model is of the form (2) with $0 < \alpha = \pi(1 + \pi)^{-1} < 1$ and $\eta_t = (1 + \pi)^{-1}(u_t + \xi_t - \zeta_t)$.

Example 2: (Muth market model) Demand and supply functions are

$$\begin{aligned} d_t &= m_I - m_p p_t + v_{1t}, \\ s_t &= r_I + r_p E_{t-1}^* p_t + r_w' w_{t-1} + v_{2t}. \end{aligned}$$

Assume that w_t is white noise with $Ew_t = 0$, $Ew_t w_t' = \Omega$. With market clearing $d_t = s_t$, we obtain (2) as the reduced form with $\eta_t = (v_{1t} - v_{2t})/m_p$, $\mu = (m_I - r_I)/m_p$, $\delta = -m_p^{-1}r_w$ and $\alpha = -r_p/m_p$. Note that $\alpha < 0$ for $r_p, m_p > 0$.

2.1.1 Econometric Learning

We now develop the formal details of econometric learning. There two key building blocks to learning. First, agents' beliefs are described by means of a forecasting model. Agents are assumed to use a **perceived law of motion (PLM)**

$$p_t = a + b'w_{t-1} + \eta_t,$$

where true values of a and b are not known. Note that this PLM has the same functional form as the unique REE. This is a natural benchmark, but in some cases it is important to allow for possible misspecification of the PLM with respect to the REE of interest. Agents may either over- or under-parameterize their PLM relative to the REE. We will discuss situations of misspecification later.

Second, we need to describe how agents obtain estimates for the parameters in the PLM. It is postulated that agents use the most popular estimation method, least squares. Thus, agents estimate a and b by recursive least squares (RLS) from past data $\{p_i, w_i\}_{i=0}^{t-1}$ and they forecast using the estimated model:

$$E_{t-1}^* p_t = a_{t-1} + b'_{t-1} w_{t-1}.$$

Here a_{t-1} and b'_{t-1} denote the estimated parameter values from using data up to period $t - 1$.

Given the forecasts, the economy attains a temporary equilibrium in period t . Alternatively, defining $\phi'_t = (a_t, b'_t)$ and $z'_i = (1, w'_i)$, the **actual law of motion (ALM)**

$$\begin{aligned} p_t &= (\mu + \alpha a_{t-1}) + (\delta + \alpha b'_{t-1})' w_{t-1} + \eta_t \\ &\equiv T(\phi_{t-1}) z_{t-1} + \eta_t \end{aligned}$$

describes the temporary equilibrium relations between the variables.

Formally, RLS estimation is given by equations

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} (p_t - \phi'_{t-1} z_{t-1}) \quad (3)$$

$$R_t = R_{t-1} + t^{-1} (z_{t-1} z'_{t-1} - R_{t-1}). \quad (4)$$

Making the shift $S_{t-1} = R_t$ and defining $E z_t z'_t = M$, RLS formally becomes a **stochastic recursive algorithm (SRA)**, as was first shown by Marcet and Sargent (1989). There are general methods for analyzing the dynamics of SRAs, which we outline in Section 2.2. In particular, conditions for

convergence of SRA are given by local stability conditions of an associated ordinary differential equation (ODE). For the RLS algorithm the ODE takes the form

$$d\phi/d\tau = S^{-1}M(T(\phi) - \phi), \quad (5)$$

$$dS/d\tau = M - S. \quad (6)$$

Finally, we remark that some papers in the literature employ the stochastic gradient (also known as “least mean squares”) algorithm in place of LS. In the current setting, the gradient algorithm and its associated ODE take the form

$$\begin{aligned} \phi_t &= \phi_{t-1} + t^{-1}z_{t-1}(p_t - \phi'_{t-1}z_{t-1}), \\ d\phi/d\tau &= M(T(\phi) - \phi). \end{aligned}$$

Generalized stochastic gradient (GSG) algorithms are well-motivated when agents allow for parameter drift or model uncertainty. See Evans, Honkapohja, and Williams (2008) for a discussion of GSG algorithms and references to the literature.

2.1.2 Expectational Stability (E-stability)

Inspecting the differential equations (5)-(6), it is seen that $\lim_{\tau \rightarrow \infty} S = M$ in the second equation. Thus, the local stability of the fixed point for the whole ODE is determined by local stability under the “small” ODE

$$d\phi/d\tau = T(\phi) - \phi. \quad (7)$$

Note that the unique REE for model (2) is the fixed point of the system (7). We say that a fixed point $\bar{\phi} = T(\bar{\phi})$ is **expectationally stable (E-stable)** if it is locally stable under the small ODE (7), as defined in Evans (1989) and Evans and Honkapohja (1992). In economic terms, the small ODE is just partial adjustment in virtual time τ . The relationship between RLS learning and E-stability is highlighted by the following result:

Proposition. *The economy converges to the REE under RLS learning if and only if the REE is E-stable. The latter occurs iff $\alpha < 1$.*

For further details of the outlined steps and a proof of the result for the model (2), see Chapter 2 of Evans and Honkapohja (2001).

The result that E-stability of REE gives the conditions for (local) convergence of RLS and related learning schemes is quite general and it holds for a wide variety of models, as discussed in Evans and Honkapohja (2001).

2.2 Stochastic Approximation Techniques

Demonstrations of convergence of RLS learning, and additional approximation results, are available using stochastic approximation techniques.

2.2.1 Decreasing-Gain Algorithms

A general form of SRA is given by

$$\theta_t = \theta_{t-1} + \gamma_t Q(t, \theta_{t-1}, X_t), \quad (8)$$

where θ_t is a vector of parameter estimates, X_t is the state vector, and γ_t is a deterministic sequence of “gains.” The function Q expresses the way in which the estimate θ_{t-1} is revised in line with the last period’s observations. In our simple model (2) with RLS learning, θ_{t-1} will include all components of ϕ_{t-1} and R_t , X_t will include the effects of w_{t-1} and η_t , and $\gamma_t = t^{-1}$. Although, in our example, X_t follows an exogenous process, this is not at all essential. In particular, X_t can be permitted to follow a vector autoregression (VAR) with parameters that may depend on θ_{t-1} .

The stochastic approximation approach associates an ODE with the SRA,

$$\frac{d\theta}{d\tau} = h(\theta(\tau)), \quad (9)$$

where $h(\theta)$ is obtained as

$$h(\theta) = \lim_{t \rightarrow \infty} EQ(t, \theta, \bar{X}_t(\theta)), \quad (10)$$

provided this limit exists. Here $\bar{X}_t(\theta)$ is the stochastic process for X_t obtained by holding θ_{t-1} at the fixed value $\theta_{t-1} = \theta$ (thus $\bar{X}_t(\theta) = X_t$ if X_t does not depend on θ_{t-1}), and E denotes the expectation of $Q(t, \theta, \bar{X}_t(\theta))$, for θ fixed, taken over the invariant distribution of the stochastic process \bar{X}_t .

For the RLS algorithm (3)-(4), with notation $R_t = S_{t-1}$, the associated ODE is

$$\begin{aligned} h_\phi(\phi) &= \lim_{t \rightarrow \infty} E[S^{-1} z_{t-1} z'_{t-1} (T(\phi) - \phi)] = S^{-1} M (T(\phi) - \phi) \\ h_S(S) &= \lim_{t \rightarrow \infty} E \frac{t}{t+1} (z_t z'_t - S) = M - S, \end{aligned}$$

which verifies the claim (5)-(6).

The stochastic approximation results show that the behavior of the SRA is well approximated by the behavior of the associated ODE for large t . In particular, possible limit points of the SRA correspond to locally stable equilibria of the ODE:

Under suitable assumptions, if $\bar{\theta}$ is a locally stable equilibrium point of the ODE then $\bar{\theta}$ is a possible point of convergence of the SRA. If $\bar{\theta}$ is not a locally stable equilibrium point of the ODE then $\bar{\theta}$ is not a possible point of convergence of the SRA, i.e. $\theta_t \rightarrow \bar{\theta}$ with probability 0.

The precise theorems are complex in detail. First, there are various ways to formalize the positive convergence result (when $\bar{\theta}$ is a locally stable equilibrium point). In certain cases, when there is a unique solution and under the SRA the ODE is globally stable, $\theta_t \rightarrow \bar{\theta}$ with probability 1 from any starting point. When there are multiple equilibria, such a strong result will not be possible. If one artificially constrains θ_t to an appropriate neighborhood of a locally stable equilibrium $\bar{\theta}$ (using a so-called “projection facility”), one can still obtain convergence with probability 1. Other versions of local stability results are also available.

Second, a careful statement is required of the technical assumptions under which the convergence conditions obtain. There are three broad classes of assumptions: (i) regularity assumptions on Q , (ii) conditions on the rate at which $\gamma_t \rightarrow 0$, (iii) assumptions on the properties of the stochastic process followed by X_t .

For details of assumptions and precise statements, see Part II of Evans and Honkapohja (2001).

2.2.2 Constant-Gain LS Mean Dynamics and Escape Dynamics

Under constant gain we replace γ_t in (8) by a constant $0 < \gamma < 1$. For example in RLS equations for ϕ_t and R_t given above, t^{-1} would be replaced by a small constant γ . We rewrite the SRA as

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

where for convenience we have dropped any explicit time dependence in Q and where we have indexed θ_t by γ so that we can consider the implications for the stochastic process for small gains $\gamma > 0$. The SRA and suitable

regularity conditions are assumed to hold for all θ_{t-1} within an open set D , and we are given an initial condition $\theta_0 = a \in D$. Several types of results are available in this setting. See, in particular, Ch. 7 of Evans and Honkapohja (2001), Cho, Williams, and Sargent (2002), and Williams (2004b).

One can, first of all, obtain the mean dynamics of θ_t^γ for the limiting case of small gains, i.e. for $\gamma > 0$ sufficiently small. Defining $h(\theta)$ as in (10), we consider the solution to the associated ODE (9). Let $\tilde{\theta}(\tau, a)$ denote the solution to (9) for initial condition $\theta_0 = a \in D$. We often refer to τ as “notional time.” Consider a fixed notional time $\mathcal{T} > 0$ and a fixed compact set $\bar{D} \subset D$. Assume that $\tilde{\theta}(\tau, a) \in \bar{D}$ for all $0 \leq \tau \leq \mathcal{T}$. The main result is that the solution $\tilde{\theta}(\tau, a)$ approximates the mean dynamics of θ_t over $0 \leq \tau \leq \mathcal{T}$. Define

$$\theta^\gamma(\tau) = \theta_t^\gamma \text{ for } t\gamma \leq \tau \leq (t+1)\gamma.$$

Thus $\theta^\gamma(\tau)$ is a continuous-time interpolation of the realization θ_t^γ of the SRA. It can be shown that as $\gamma \rightarrow 0$ the normalized random variables $U^\gamma(\tau) = \gamma^{-1/2}(\theta^\gamma(\tau) - \tilde{\theta}(\tau, a))$ over $0 \leq \tau \leq \mathcal{T}$ converge weakly to the solution $U(\tau)$ of the stochastic differential equation

$$dU(\tau) = D_\theta h(\tilde{\theta}(\tau, a))U(\tau)d\tau + \mathcal{R}^{1/2}(\tilde{\theta}(\tau, a))dW(\tau),$$

with initial condition $U(0) = 0$, where $W(\tau)$ is a standard vector Wiener process and $\mathcal{R}(\theta)$ can be computed from $Q(\theta, \bar{X}_t(\theta))$. The solution is described in Evans and Honkapohja (2001), Chapter 7.4. In particular, the solution satisfies $E(U(\tau)) = 0$. It follows that for $0 \leq \tau \leq \mathcal{T}$ the mean dynamics of the SRA for $\gamma > 0$ small can be approximated using $E(\theta^\gamma(\tau)) \approx \tilde{\theta}(\tau, a)$. Specifically, the mean dynamics of the SRA satisfy

$$E\theta_t^\gamma \approx \tilde{\theta}(\gamma t, a)$$

for $0 \leq t \leq \mathcal{T}/\gamma$. The solution to the stochastic differential equation in $dU(\tau)$ can also be used to approximate $Var(\theta_t^\gamma)$ for small γ , and this approximation often takes a simple form as γt becomes large.

The other class of results that are available for constant-gain algorithms are “escape dynamics” based on large deviation theory. Although as $\gamma \rightarrow 0$ the solution to the above stochastic differential equation provides a good approximation to the distribution of θ_t^γ , with positive probability there will be large deviations from the mean dynamics, and over long stretches of time these unusual escape dynamics may be of considerable practical interest, as

argued in Sargent (1999), Cho, Williams, and Sargent (2002) and Williams (2004b). Using large deviation tools it is possible to compute the direction of the paths that are most likely to leave a specified neighborhood of the mean dynamics and thus provide useful information on these escape dynamics.

2.3 The Planning Horizon

In the Lucas/Muth model and in overlapping generations models with two-period lifetimes, agents in the current period make forecasts for the values of aggregate variables in the next period. However, many modern macroeconomic models are set in a representative-agent framework with infinitely-lived agents who solve infinite-horizon dynamic optimization problems. Typically, under RE the reduced-form equations for these models can be stated in the form (1). This reduction relies on the use of the Euler equations to describe the first-order conditions.

Under learning there are alternative approaches in infinite-horizon settings. In Evans and Honkapohja (2001), Chapter 10, the learning framework was kept close to the RE reduced-form set-up, a procedure that can be justified if agents make decisions based directly on their Euler equations. This approach has been used, for example, in Bullard and Mitra (2002) and Evans and Honkapohja (2003c). An alternative approach, followed by Preston (2005), assumes that households use estimated models to forecast aggregate quantities infinitely far into the future to solve for their current decisions. We now illustrate the two approaches using a simple endowment economy.³

A representative consumer makes consumption-saving decisions using the intertemporal utility function

$$E_t^* \sum_{s=t}^{\infty} \beta^{s-t} U(C_s). \quad (11)$$

Each period the household has a random endowment of Y_t and there is a market in safe one-period loans with gross rate of return R_t , assumed known at t . Initial wealth for each agent is zero. Output Y_s follows an exogenous process given by

$$\log Y_s = \mu + \rho \log Y_{s-1} + v_s,$$

³The passage is based largely on Honkapohja, Mitra, and Evans (2002), and we also draw on Evans, Honkapohja, and Mitra (2007).

where $|\rho| < 1$ and v_s is white noise. Expectations are not necessarily rational, which is indicated by $\hat{\cdot}$ in the expectations operator. The household has an intertemporal budget constraint

$$C_t + \sum_{s=t+1}^{\infty} \mathcal{R}_{t+1,s} C_s = Y_t + \sum_{s=t+1}^{\infty} \mathcal{R}_{t+1,s} Y_s, \quad (12)$$

where $\mathcal{R}_{t+1,s} = (R_{t+1} \dots R_s)^{-1}$.

Maximizing (11) subject to (12) yields the Euler equation as a necessary first-order condition (FOC),

$$U'(C_t) = \beta R_t E_t^* U'(C_{t+1}). \quad (13)$$

In equilibrium $C_t = Y_t$, as output is assumed to be perishable. In the ‘‘Euler-equation learning’’ approach, (13) is treated as a behavioral equation, determining for each agent their temporary equilibrium demand for C_t as a function of R_t and the forecast $E_t^* U'(C_{t+1})$ of their $t + 1$ marginal utility of consumption. Imposing the market clearing condition $C_t = Y_t$, and using the representative agent setting, (13) determines the temporary equilibrium interest rate according to

$$R_t^{-1} = \beta (E_t^* U'(C_{t+1})) / U'(Y_t).$$

Log-linearizing the Y_t process and (13) around the non-stochastic steady state yields

$$y_t = \rho y_{t-1} + v_t, \quad (14)$$

where $y_t = \log(Y_t/\bar{Y})$, and the consumer’s demand schedule

$$c_t = E_t^* c_{t+1} - \sigma r_t. \quad (15)$$

Here $c_t = \log(C_t/\bar{C})$, r_t is the net return, based on the approximation $r_t \approx \log(R_t/\bar{R})$ and $\sigma = -\frac{U'(\bar{C})}{U''(\bar{C})\bar{C}}$ is the coefficient of intertemporal substitution. (Bars over the variables denote the non-stochastic steady state). In the temporary equilibrium $c_t = y_t$ and $r_t = \sigma^{-1}(E_t^* c_{t+1} - y_t)$.

REE of the linearized model is given by

$$r_t = -(1 - \rho)\sigma^{-1}y_t$$

and for rational forecasts we have

$$E_t c_{t+1} = \rho y_t. \quad (16)$$

To formulate “Euler equation” (EE) learning, based on (15), suppose that agents have a PLM forecast function nesting the REE:

$$E_t^* c_{t+1} = m + ny_t, \quad (17)$$

with coefficient estimates (m_t, n_t) obtained using a regression of c_s on y_{s-1} using data $s = 1, \dots, t - 1$. As usual, (m_t, n_t) are updated over time.

The question is whether $(m_t, n_t) \rightarrow (0, \rho)$ over time. This can easily be verified using E-stability arguments. Given the PLM (17), in the temporary equilibrium $r_t = -\sigma^{-1}[y_t(1 - n) - m]$. However the ALM forecasts are $E_t c_{t+1} = \rho y_t$, so that the T-map is $T(m, n) = (0, \rho)$. Clearly the E-stability differential equation

$$\frac{d(m, n)}{d\tau} = T(m, n) - (m, n),$$

is stable, and hence there is convergence of LS learning to RE in this model.

Under EE learning, agents choose their consumption demand using the Euler equation between today’s and tomorrow’s consumption. To implement this FOC requires a forecast of agent’s own C_{t+1} . This forecast is made, assuming that the agent’s future consumption is related (as it is in the REE) to the key state variable, y_t . (In an RBC model the state would include capital and technology). Thinking one step ahead, in this way, appears to us to be a plausible and natural form of bounded rationality. Furthermore, although this formulation does not explicitly impose the intertemporal budget constraint, it can be verified that along the learning path both (12) and the transversality condition are also satisfied.

An alternative approach postulates that consumption demand each period is based on forecasts over an infinite horizon. We call this approach, presented for the New Keynesian model in Preston (2005), infinite-horizon (IH) learning, and we describe it for the current context. Log-linearizing the intertemporal budget constraint (12) yields

$$c_t + \sum_{s=t+1}^{\infty} \beta^{s-t} E_t^* c_s = y_t + \sum_{s=t+1}^{\infty} \beta^{s-t} E_t^* y_s, \quad (18)$$

where we have used $\bar{C} = \bar{Y}$. Iterating the linearized Euler equation (15) backwards for $s \geq t + 1$ gives $E_t^* c_s = c_t + \sigma \sum_{j=t}^{s-1} E_t^* r_j$. Substituting this into

(18) and solving for c_t leads to the behavioral equation

$$c_t = (1 - \beta)y_t - \sigma\beta r_t + \sum_{s=t+1}^{\infty} \beta^{s-t} [(1 - \beta)E_t^*y_s - \sigma\beta E_t^*r_s]. \quad (19)$$

Here we have assumed that both y_t and r_t are known at t .

Suppose that agents do not know the RE relationship between y_t and r_t , but have the PLM

$$r_t = d + fy_t,$$

where at time t the coefficients are estimated to be d_t, f_t . To determine whether there is convergence to the REE, we again turn to E-stability. Using $E_t^*r_s = d + fy_s$ and $E_t^*y_s = \rho^{s-t}y_t$ in (19), and imposing market clearing $c_t = y_t$ we obtain

$$y_t = -\sigma\beta r_t + \frac{1 - \beta - \sigma\beta^2\rho f}{1 - \beta\rho}y_t - \frac{\sigma\beta^2d}{1 - \beta}.$$

Solving for r_t , the T-mapping is

$$d \rightarrow -\frac{\beta d}{1 - \beta} \text{ and } f \rightarrow -\sigma^{-1} \frac{1 - \rho + \beta\sigma\rho f}{1 - \beta\rho}.$$

The fixed point of T is the REE $d = 0$ and $f = -(1 - \rho)\sigma^{-1}$, and the E-stability ODE is clearly stable.

Although for this particular model, learning stability holds for both EE and IH learning, in more general models it is possible for stability to depend on the planning horizon of the agents. For an analysis of these issues in a general framework, see Evans and McGough (2008).

2.4 Multiple Equilibria and Learning

Thus far, we have considered linear or linearized models. One attractive feature of linear models is that for them the class of possible REE can be described explicitly. For some models, such as model (2), a unique stationary solution exists. The economy or model is then said to be **determinate**. If there are multiple non-explosive solutions, the model is said to be **indeterminate**.

Consider the simple forward-looking linear model

$$x_t = \mu + \alpha E_t^*x_{t+1} + v_t, \quad (20)$$

where x_t is a scalar endogenous variable and the random shock v_t is *iid* with mean zero. The model (20) is determinate if $|\alpha| < 1$ and indeterminate if $|\alpha| > 1$. The fundamental REE is $x_t = (1 - \alpha)^{-1}\mu + v_t$. Writing $x_{t+1} = E_t x_{t+1} + \eta_{t+1}$, where under RE $E_t \eta_{t+1} = 0$, substituting into (20), and rearranging we get

$$x_{t+1} = -\alpha^{-1}\mu + \alpha^{-1}x_t - \alpha^{-1}v_t + \eta_{t+1}.$$

Thus for any stationary η_{t+1} this defines a stationary stochastic process for x_{t+1} when $|\alpha| > 1$.

Next, consider the nonlinear forward-looking model

$$x_t = E_t^* F(x_{t+1}), \tag{21}$$

where for simplicity the random shock has been omitted (for extensions to models with intrinsic random shocks see, e.g., Part IV of Evans and Honkapohja (2001)). Samuelson's overlapping generations model is a well-known example of the nonlinear model.⁴ An important difference between linear and nonlinear models is that it is usually not possible to describe all REE for the nonlinear model. Potential equilibria include steady states, deterministic cycles and sunspot solutions.

A steady state \bar{x} is defined by the equation $\bar{x} = F(\bar{x})$ while a deterministic k -period cycle $(\hat{x}_1, \dots, \hat{x}_K)$ is defined by the equations $\hat{x}_{k-1} = F(\hat{x}_k)$ for $k = 2, \dots, K$ and $\hat{x}_K = F(\hat{x}_1)$. A stationary sunspot equilibrium (SSE) is a stochastic REE in which agents' expectations depend on an extraneous random variable that has no fundamental significance for the economy. A widely-discussed case of SSEs takes the form of a finite Markov chain. Consider a finite Markov chain s_t that can take values $(1, \dots, K)$ with transition probabilities π_{ij} , where $\pi_{ij} = \Pr[s_{t+1} = j | s_t = i]$. A k -tuple (x_1^*, \dots, x_K^*) is an SSE with transition probabilities π_{ij} if

$$x_t = x_i^* \text{ if } s_t = i \text{ where } x_i^* = \sum_{s=1}^K \pi_{is} F(x_s^*) \text{ for all } i = 1, \dots, K.$$

Figures 1 and 2 illustrate these different types of REE in the nonlinear model (21).

⁴See e.g. Chapter 9 of Ljungqvist and Sargent (2003) for an overview of the overlapping generations model.

FIGURES 1 AND 2 ABOUT HERE

Global and local determinacy and indeterminacy for the nonlinear model are discussed, e.g., in Chiappori and Guesnerie (1991). We use model (21) to discuss learning as a selection criterion when there are multiple REE.

Evans and Honkapohja (1995) develop E-stability conditions for steady states and cycles and show that the relationship between convergence of LS learning and E-stability continues to hold for models of type (21).⁵ For sunspot equilibria the first result about convergence of learning to SSE was obtained by Woodford (1990) in the context of a specific model. SSEs near deterministic equilibria are of special interest in many contexts. Local stability results for models of type (21) were developed by Evans and Honkapohja (1994) and Evans and Honkapohja (2003b). They showed that

- 1) E-stable SSEs exist near a pair of distinct steady states iff both steady states are E-stable,
- 2) E-stable SSEs exist near an E-stable deterministic cycle, and
- 3) E-stable SSEs exist near a single steady state iff $F'(\bar{x}) < -1$.

These results help to select among REE when multiple equilibria exist. By narrowing consideration to “reasonable” equilibria, learning is a useful selection criterion, even if it does not necessarily select a unique solution.

A variety of selection results also exist for linear models, e.g., see Evans and Honkapohja (1992) and Part III of Evans and Honkapohja (2001). For the basic forward-looking model (20) it is easy to check that the E-stability condition is $\alpha < 1$. It follows that determinacy is sufficient but not necessary for E-stability of the fundamental solution. This result has been significantly generalized by McCallum (2007). However, it should be noted that the result depends on the timing of information, see Bullard and Eusepi (2008).

2.5 Further Issues

We close this section with a brief general discussion of heterogeneous expectations and dynamic predictor selection.

⁵Convergence of learning to steady states and cycles was initially considered for finite-memory learning rules by several researchers, see Grandmont (1998) for a review. These rules do not converge fully to REE in stochastic models.

2.5.1 Heterogeneous Expectations

The preceding discussion has assumed homogeneous expectations for analytical convenience. In practice, heterogeneous expectations can be a major concern. In some models the presence of heterogeneous expectations does not have major effects on stability conditions, as first suggested by Evans and Honkapohja (1996) and substantially generalized by Giannitsarou (2003). However, Honkapohja and Mitra (2006) showed that interaction of structural and expectational heterogeneity can make the conditions for convergence of learning significantly more stringent than those obtained under homogeneous expectations.

Consider a forward-looking model with S classes of agents:

$$\begin{aligned} y_t &= \alpha + \sum_{i=1}^S A_i E_t^* y_{t+1} + Bw_t, \\ w_t &= Fw_{t-1} + v_t. \end{aligned}$$

where y_t is a scalar endogenous variable and w_t follows a stationary $AR(1)$. F is taken to be known (if not, it could be estimated). $M_w = \lim_{t \rightarrow \infty} Ew_t^2 > 0$. The fundamental or minimal state variable (MSV) solution is

$$y_t = a + bw_t,$$

where a and b can be solved (usually uniquely) from equations

$$a = \alpha + A^M a \text{ and } b = A^M bF + B,$$

where $A^M = \sum_{i=1}^S A_i$.

Define the state variables $z'_t = (1, w_t)$ and the matrix of parameters $\varphi'_{i,t} = (a_{i,t}, b_{i,t})$, $i = 1, \dots, S$. Agents are assumed to have PLMs

$$y_t = a_{i,t} + b_{i,t}w_t = \varphi'_{i,t}z_t, \quad i = 1, \dots, S.$$

The resulting ALM is

$$\begin{aligned} y_t &= \left[\alpha + \sum_{i=1}^S A_i a_{i,t}, \left(\sum_{i=1}^S A_i b_{i,t} \right) F + B \right] \begin{bmatrix} 1 \\ w_t \end{bmatrix} \\ &= T(\varphi'_{1,t}, \dots, \varphi'_{S,t})z_t. \end{aligned}$$

We consider **mixed RLS/SG learning** when initial conditions of the different types of agents are different and different agents may use different learning rules. The learning rules may have different gain parameters. Specifically, it is assumed that types $i = 1, \dots, S_0$ agents use RLS and types $j = S_0 + 1, \dots, S$ agents use SG learning rules. For agents $i = 1, \dots, S_0$ the algorithm is given by

$$\begin{aligned}\varphi_{i,t} &= \varphi_{i,t-1} + \gamma_{i,t} R_{i,t}^{-1} z_{t-1} (y_{t-1} - \varphi'_{i,t-1} z_{t-1})', \\ R_{i,t} &= R_{i,t-1} + \gamma_{i,t} (z_{t-1} z'_{t-1} - R_{i,t-1}),\end{aligned}$$

while for agents $i = S_0 + 1, \dots, S$ it is given by

$$\varphi_{i,t} = \varphi_{i,t-1} + \gamma_{i,t} z_{t-1} (y_{t-1} - \varphi'_{i,t-1} z_{t-1})'.$$

The gain sequences $\gamma_{i,t}$ are assumed to satisfy $\lim_{t \rightarrow \infty} E(\gamma_{i,t}/\gamma_t) = \delta_i > 0$, where the decreasing and positive sequence γ_t satisfies

- (i) $\hat{\gamma}_{i,t} \leq K_i \gamma_t$ for some constant $K_i > 0$,
- (ii) $\sum_{t=1}^{\infty} \gamma_t = \infty$, $\sum_{t=1}^{\infty} \gamma_t^2 < \infty$ and $\limsup (1/\gamma_{t+1} - 1/\gamma_t) < \infty$.

Also

- (iii) $\delta_i \neq \delta_j$, i.e. mean gains of the agents can differ asymptotically.

Generalizations to random gains are possible, see Honkapohja and Mitra (2006).

It can be shown that in the case of mixed RLS/SG learning, stability is determined by

$$\begin{aligned}d\varphi_i/d\tau &= \delta_i (T(\varphi'_1, \dots, \varphi'_S)' - \varphi_i), i = 1, \dots, S_0 \\ d\varphi_i/d\tau &= \delta_i M_z (T(\varphi'_1, \dots, \varphi'_S)' - \varphi_i), i = S_0 + 1, \dots, S,\end{aligned}$$

and the **generalized E-stability condition** is stricter than usual E-stability or SG-stability. The new feature is that, in general, speeds of learning, as indicated by parameters δ_i , affect convergence. However, for the univariate case $n = k = 1$ we have the following result. Assume that the aggregate economy is E-stable and the parameters A_i have the same sign. If the different agents use either RLS or SG learning rules, the economy converges to the MSV REE for all $\{\delta_i\}_{i=1}^S$ and M_z .

2.5.2 Dynamic Predictor Selection

Another natural way to introduce heterogeneity is assume that different agents have different types of forecasting models, with the model choices

at each point in time determined endogenously. Brock and Hommes (1997) postulate that agents have a finite set of predictors or expectation functions for predicting price. Each predictor has a fitness measure associated with it, based on past performance, as well as a cost of using that predictor. The proportion of agents who select a predictor depends on its fitness, i.e., an estimate of the profits net of costs.

Brock and Hommes (1997) study the resulting “adaptively rational expectations dynamics” for the standard “cobweb” model with two predictors: rational and naive forecasts. The model is nonstochastic, so that RE is equivalent to perfect foresight. Demand is assumed to be linear $D(p_t) = A - Bp_t$. Firms have a quadratic cost function $c(q) = q^2/2b$ and thus the supply curve $S(p_t^e) = bp_t^e$. There are two predictors available, the perfect foresight predictor $p_t^e = p_t$, which costs $C \geq 0$, and the naive predictor $p_t^e = p_{t-1}$, which is free. Letting n_{1t} and n_{2t} denote the proportion of agents using the perfect foresight and naive predictors, respectively, market equilibrium at t is given by

$$D(p_t) = n_{1,t-1}S(p_t) + n_{2,t-1}S(p_{t-1}).$$

The main fitness studied measure examined is net realized profit in the last period. Since profits are $\pi_t = p_t S(p_t^e) - c(S(p_t^e))$ we have realized time t profits given by

$$\pi_{1t} = \frac{b}{2}p_t^2 - C \text{ and } \pi_{2t} = \frac{b}{2}p_{t-1}(2p_t - p_{t-1})$$

The proportion of agents using the j th predictor is given by the “multinomial logit” ratios

$$n_{1,t} = \exp(\beta\pi_{1,t})/(\exp(\beta\pi_{1,t}) + \exp(\beta\pi_{2,t})) \text{ and } n_{2,t} = 1 - n_{1,t}.$$

The parameter β measures the intensity with which agents choose predictors with higher “fitness.” For $\beta = +\infty$ all agents choose the predictor with highest previous period net profit.

These equations fully define the “adaptively rational equilibrium dynamics.” This system has a unique steady state. Brock and Hommes (1997) focus on the case $b/B > 1$, in which the model is locally unstable under naive expectations. If $C > 0$, the dynamics depend crucially on β . Above a critical value β_1 the steady state is an unstable saddlepoint. Stable two-cycle and higher-order cycles, the coexistence of low periodic attractors, and chaotic

attractors appear as β increases. Brock and Hommes call the resulting complicated dynamical phenomena a “rational route to randomness”.

The economic mechanisms generating the complex dynamics are straightforward. When agents use the cheapest predictor (here static expectations) the steady state is unstable, whereas the costly sophisticated predictor is stabilizing. Near the steady state, it pays to use the cheap predictor, but this pushes the economy away from the steady state. For a high enough intensity of choice β this tension leads to local instability and complex global dynamics.

The dynamic selector framework is extended by Branch and Evans (2006) and Branch and Evans (2007) to incorporate stochastic features and econometric learning. Based on the assumption that agents employ parsimonious models, Branch and Evans study the implications of agents choosing between equally costly misspecified models. Consider a stochastic cobweb model, driven by two exogenous shocks, with agents choosing between two models, each of which omits one of the variables. For this set-up a misspecification equilibrium (ME) has the following elements: (i) the coefficients of each forecasting model are given by the true linear projections, (ii) fitness of a forecast rule is measured by the mean equilibrium profits of an agent using that rule, and (iii) the proportions of agents using the different rules are in accordance with the multinomial logit ratios described above. Branch and Evans (2006) show that, under weak assumptions, a ME exists. In the two-predictor case with exogenous variables the ME is unique. However, in contrast to the REE, there are cases of “intrinsic heterogeneity,” in which both predictors are used, even with $\beta = +\infty$. Branch and Evans (2006) also examine the stability of this ME under real-time learning and dynamic predictor selection.

Branch and Evans (2007) use the same framework to examine a Lucas-type monetary model in which there is positive expectational feedback. The new phenomenon is that this model can have two ME, in which agents coordinate on either of the two models. The stochastic process for inflation and output can then exhibit regime-switching or parameter drift, in line with much macroeconomic evidence.

3 Economic Applications

3.1 Monetary Policy

Analysis of monetary policy rules from the learning viewpoint has recently become a very popular research topic.⁶ We now discuss some key aspects of this rapidly growing literature.

3.1.1 New Classical Model with Constant Gain

Orphanides and Williams (2005b) (OW) use a simple two-equation macro model to show that constant-gain learning by private agents has major implications for economic policy. Their model is based on a New Classical expectations-augmented Phillips curve with inertia:

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

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 $t + 1$, and e_{t+1} is a white noise inflation shock. $(1 - \phi)\pi_t$ represents intrinsic inflation persistence. We assume $0 < \phi < 1$.

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

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

x_t is set by monetary policy at t and u_{t+1} is white noise. Through monetary policy it is assumed that policy-makers are able one period ahead to control aggregate output up to the unpredictable random disturbance u_{t+1} .

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

$$x_t = -\theta(\pi_t - \pi^*). \quad (23)$$

The policy-makers aim to choose θ optimally given their loss function

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

with $0 \leq \omega \leq 1$ parameterizing the relative weight on inflation vs. output stabilization. Under RE, the optimal choice takes the form $\theta^P = \theta^P(\omega, (1 -$

⁶For surveys see Evans and Honkapohja (2003a), Bullard (2006), and Evans and Honkapohja (2007).

$\phi)/\alpha)$, where θ^P is increasing in its arguments, and inflation follows an AR(1) process. Under LS learning, private agents estimate the PLM

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

by a LS-type regression, and at time t forecast $\pi_{t+1}^e = c_{0,t} + c_{1,t}\pi_t$. The REE can be shown to be E-stable, so under decreasing gain, LS learning would converge to the REE.

With constant-gain LS (which OW call “perpetual learning”), estimates $c_{0,t}, c_{1,t}$ no longer fully converge to the REE, but instead to a stochastic process. If the gain parameter $\kappa > 0$ is very small, then estimators will be close to the REE values for most of the time with high probability, and output and inflation will be near their REE paths. Nonetheless, small plausible values like $\kappa = 0.05$ can lead to very different outcomes in the calibrations OW consider. Using simulations OW find that (i) the standard deviations of $c_{0,t}$ and $c_{1,t}$ are large even though forecast performance remains good, (ii) there is a substantial increase in the persistence of inflation, compared to the REE, as measured by the AR(1) coefficient for π_t , and (iii) the policy trade-off between standard deviations σ_π and σ_y shifts out substantially and sometimes in a non-monotonic way.

Under perpetual learning by private agents, if policy-makers keep to the same class of rules then they should choose a different θ than under RE. One key finding is that the “naive” policy choice $\theta = \theta^P$, can be strictly inefficient when agents are learning. In general, policy should be more hawkish, i.e. under learning the monetary authorities should pick $\theta > \theta^P$.⁷ Finally, following a sequence of unanticipated inflation shocks, inflation “doves” can do very poorly, with expectations deviating substantially from RE. The intuition for these results is that a more hawkish policy helps to keep inflation expectations π_{t+1}^e “in line,” i.e. closer to RE values.

3.1.2 The Rise and Fall of Inflation

Several recent papers have argued that the learning approach plays a central role in the historical explanation of the rise and fall of US inflation over the 1960-1990 period. Sargent (1999) and Cho, Williams, and Sargent (2002) emphasize the role of policy-maker learning. They argue that if monetary policymakers attempt to implement optimal policy while estimating

⁷Similar results emerge in the more general setting in Orphanides and Williams (2007).

and updating the coefficients of a misspecified Phillips curve, there will be both periods of inefficiently high inflation and occasional escapes to low inflation. Sargent, Williams, and Zha (2006) estimate a version of this model. They find that shocks in the 1970s led the monetary authority to perceive a trade-off between inflation and unemployment, leading to high inflation, and subsequent changed beliefs about this trade-off account for the conquest of US inflation during the Volker period.

Primiceri (2006) makes a related argument, emphasizing both (i) policymaker learning about both the Phillips curve parameters and the aggregate demand relationship, and (ii) uncertainty about the unobserved natural rate of unemployment, U_t^n . The great inflation of 1970s initially resulted from a combination of underestimates of both U_t^n and the persistence of inflation. This also led policymakers to underestimate the impact of unemployment on inflation until estimates of the perceived trade-off between inflation and unemployment changed during the Volker period.

Other empirical accounts of the period that emphasize learning include Bullard and Eusepi (2005), which examines the implications of policymaker learning about the growth rate of potential output, Orphanides and Williams (2005a), which underscores both private-agent learning and policymaker misestimates of the natural rate of unemployment, and Cogley and Sargent (2005), which develops a historical account of inflation policy emphasizing Bayesian model averaging and learning by policymakers uncertain about the true economic model.

3.1.3 New Keynesian Models and Policy Rules

The New Keynesian (NK) model is currently the most widely-used vehicle for studying monetary policy. The NK model is a dynamic stochastic general equilibrium model with a representative consumer, and price rigidity modelled using monopolistic competition with constraints on price setting.⁸

We directly employ the log-linearized version of the NK model. The aggregate demand and supply curves summarize private-sector behavior. The simplest version of the NK model takes the form

$$x_t = -\varphi(i_t - E_t^* \pi_{t+1}) + E_t^* x_{t+1} + g_t, \quad (24)$$

$$\pi_t = \lambda x_t + \beta E_t^* \pi_{t+1} + u_t. \quad (25)$$

⁸See Clarida, Gali, and Gertler (1999) for a survey article and the books Woodford (2003), Walsh (2003), and Galí (2008).

where x_t denotes the output gap, π_t is the rate of inflation, i_t is the nominal rate of interest, and g_t and u_t are exogenous AR(1) stationary shocks. The notation $E_t^*(\cdot)$ allows for non-rational expectations. The aggregate demand or IS curve is obtained by log-linearizing the consumer's Euler equation and employing the goods market-clearing condition, so that the equation is expressed in terms of the output gap. The aggregate supply or AS (or NK Phillips) curve is derived as a linearization of the firms' optimality condition under the price setting constraint.

The model is completed by specifying an interest-rate rule for monetary policy, e.g., of the contemporaneous or forward-looking Taylor-type form proposed by Taylor (1993):

$$\begin{aligned} i_t &= \chi_\pi \pi_t + \chi_x x_t, \text{ or} \\ i_t &= \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}. \end{aligned}$$

The form of the policy rules affects the determinacy and learnability properties of the NK model. Multiplicity of equilibria or expectational instability of equilibrium under learning means that there can be undesirable fluctuations in the economy. To avoid this possibility, good policy should focus on interest-rate rules that deliver stability under learning and determinacy.

For Taylor rules Bullard and Mitra (2002) showed:

(1) The standard Taylor rule $i_t = \chi_\pi \pi_t + \chi_x x_t$ yields both E-stability and determinacy iff the inequality

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

holds.

(2) The forward-looking rule $i_t = \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}$ delivers E-stability and determinacy of equilibrium when $\chi_\pi > 1$ and $\chi_x \geq 0$ is sufficiently small.

Taylor rules do not usually describe optimal policy in the NK model. Optimal monetary policy under learning has been considered by Evans and Honkapohja (2003c) under discretion and Evans and Honkapohja (2006) under commitment (using the "timeless perspective" described in Woodford (2003)). The FOC for the timeless-perspective optimum (often called the targeting rule) is

$$\lambda \pi_t = \alpha(x_t - x_{t-1}), \tag{26}$$

where α is the relative weight on output variance in the loss function.

There are different ways for attempting to implement optimal monetary policy. One natural formulation is to solve (25)-(26) for the optimal REE for

x_t, π_t and to insert this solution into (24) to obtain a fundamentals-based rule of the form $i_t = \psi_x x_{t-1} + \psi_g g_t + \psi_u u_t$. (There are also “hybrid rules” in which i_t responds to deviations from the targeting rule.) Another implementation solves (24)-(25)-(26), given expectations, for an “expectations-based” rule of the form $i_t = \delta_L x_{t-1} + \delta_\pi E_t^* \pi_{t+1} + \delta_x E_t^* x_{t+1} + \delta_g g_t + \delta_u u_t$, where RE has not been imposed on private-sector expectations. We have:

Proposition. *Optimal rules based only on fundamentals lead to E-instability and (often) to indeterminacy. Optimal expectations-based rules deliver both E-stability and determinacy.*

This proposition is based on the formulation (24)-(25), which presumes that agents aim to satisfy their subjective Euler conditions, as discussed in Section 2.3. Preston (2005) and Preston (2006) analyze modifications to the preceding analysis when it is assumed that agents instead have infinite planning horizons.

Our discussion in this section is just the beginning of what is already a large and growing literature. Further aspects of policy design in the NK model include lack of observability of expectations and other variables, imperfect knowledge of structural parameters for optimal policy rules, implications of constant-gain learning, and extensions of the NK model to open economies and supply-side channels of monetary policy transmission. See Evans and Honkapohja (2007) for a discussion of these and other topics, with references.

3.2 Business Cycles

3.2.1 The Basic RBC Model

Real Business Cycle (RBC) models have been widely discussed since the 1980s, see e.g. Cooley (1995). The basic RBC model describes an infinite-horizon, representative-agent economy with flexible prices and perfect competition. In such economies the competitive equilibrium is Pareto efficient, so that computing the equilibrium can be done by solving the corresponding planning problem.

The social planner maximizes

$$\sum_{t=0}^{\infty} E_0 \mathfrak{B}^t (\log(C_t) - L_t)$$

subject to the constraints

$$\begin{aligned} C_t + K_{t+1} &\leq S_t K_t^\alpha (\gamma^t L_t)^{1-\alpha} + (1-d)K_t, \\ S_t &= S_{t-1}^\rho V_t, \\ K_0 &= \bar{K}_0, S_0 = \bar{S}_0. \end{aligned}$$

C_t and K_t denote consumption and capital, respectively. S_t is a productivity shock and V_t is an *iid* innovation with mean one. The log of S_t thus follows an *AR*(1) process and ρ captures the persistence of the technology shocks.

This planning problem does not have an explicit solution but dynamics of the economy can be described using a linearization around a non-stochastic steady state. (See Section 10.4 of Evans and Honkapohja (2001) for formal details.) Defining detrended variables $\tilde{K}_t = \frac{K_t}{\gamma^t}$, $\tilde{C}_t = \frac{C_t}{\gamma^t}$ etc., the first-order optimality conditions are transformed to equations with asymptotically stationary variables and a unique steady state. Log-linearizing around the steady state, defining the variables $k_t = \log(\tilde{K}_t/\bar{K})$, $c_t = \log(\tilde{C}_t/\bar{C})$, $s_t = \log S_t$ and $\vartheta_t = \log V_t$, and introducing vector notation $y_t' = (c_t, k_t, s_t)$, the model has the standard form

$$y_t = AE_t y_{t+1} + B y_{t-1} + C \vartheta_t. \quad (27)$$

It is well-known that the basic RBC model is determinate (saddle-point stable) and the unique solution has *VAR*(1) form

$$y_t = a + b y_{t-1} + c \vartheta_t, \quad (28)$$

with particular values $\bar{a}, \bar{b}, \bar{c}$. To check E-stability of this solution one treats (28) with general values for a, b, c as the PLM. The ALM can then be computed in a standard way. It is possible to develop general E-stability conditions for models of the form (27) and check their validity numerically. For the Farmer (1999) parameter values, the RE solution is E-stable. The analysis of RLS learning can be done in a standard way, provided the dimension of shocks is increased to avoid an exact linear relationship between contemporaneous consumption, capital and the technology shock.

3.2.2 Applications and Extensions of the RBC Model

Williams (2004a) explores further features of the RBC model dynamics under learning. Using simulations he shows that the dynamics under RE and learning are not very different unless agents need to estimate structural aspects

as well as the reduced form PLM parameters. Huang, Liu, and Zha (2008) focus on the role of misspecified beliefs and suggest that these can substantially amplify the fluctuations due to technology shocks in the standard RBC model.⁹

Other papers on learning and business cycle dynamics include Van Nieuwerburgh and Veldkamp (2006) and Eusepi and Preston (2008). The former formulates a model of Bayesian learning about productivity and suggests that the resulting model can explain the sharp downturns that are an empirical characteristic of business cycles. The latter paper introduces the notion of infinite-horizon decision rules (discussed in Section 2.3 above) to RBC models and argues that variants of the model under learning can resolve some of the empirical difficulties of RE models with business cycles.

Giannitsarou (2006) extends the basic RBC model to include government spending financed by capital and labour taxes. Her objective is to compare the transitional dynamics of the model under RE and under RLS learning when an unanticipated reduction in the capital tax displaces the steady state equilibrium. Under RE there is the usual saddle-path adjustment: consumption jumps instantaneously and the economy monotonically converges to the new steady state. In contrast, the nature of dynamics under learning depends on the nature of technology shocks near the time of tax change. With negative shocks the adjustment shows a delayed response in economic activity, though eventually the dynamics approximate the saddle-path dynamics. In contrast, under positive technology shocks the learning and RE adjustment paths are nearly identical. This is an important finding as tax reductions are often carried out in bad times.

3.2.3 Sunspot Fluctuations

When the RBC model is generalized to include externalities, monopolistic competition or other distortions, it is possible for the steady state to be indeterminate, i.e. to possess multiple solutions, including a dependence on sunspots, in a neighborhood of the steady state. Examples are the models of Farmer and Guo (1994), Benhabib and Farmer (1996) and Schmitt-Grohe and Uribe (1997). This line of research suggests SSEs as a possible model of the business cycle.

⁹Williams (2004a) also considers misspecification in an extended RBC model with complementarities and shows that there can large fluctuations taking the form of “escape dynamics.”

Are these SSEs stable under learning? This issue was initially studied in Chapter 10.5 of Evans and Honkapohja (2001), where it was found that the SSE in the Farmer-Guo calibrated model was not stable under learning. Stability of SSEs in these models was examined further in Evans and McGough (2005a) and Duffy and Xiao (2007). In general, stability of SSEs can depend both on their parametric representations and the precise information set available to agents. However, for this class of models both Evans and McGough (2005a) and Duffy and Xiao (2007) obtain predominately negative results. Duffy and Xiao further argue that empirically plausible adjustment dynamics rule out stable sunspots in this class of models. Evans and McGough do find some cases of stable sunspots in the case of “common factor” representations of SSEs, when the information sets of private agents include contemporaneous aggregate endogenous variables. Stable SSEs of this type arise only in small parameter regions and are sensitive to the information set, but they do arise for some plausible calibrations of the Schmitt-Grohe and Uribe (1997) model. On balance, the existing results provide a challenge to future researchers to design versions of RBC-type models that exhibit robustly stable SSEs.

Turning to other models, we have already seen that stable SSEs (i.e. stable under learning) have been shown to exist by Woodford (1990) in a monetary overlapping generations model. Stable SSEs have also been obtained by Howitt and McAfee (1992) in a model with search externalities and by Evans, Honkapohja, and Romer (1998) in an endogenous growth model. We briefly discuss some positive results for the standard NK model, introduced in Section 3.1.3, and for a cash-in-advance (CA) model.

It is well-known that in NK models indeterminacy and existence of SSEs arise for some policy parameters χ_π, χ_x . Stability of SSEs under learning is examined in Honkapohja and Mitra (2004) and Evans and McGough (2005b). Writing the model in bivariate form

$$y_t = ME_t^* y_{t+1} + Pv_t \text{ where } y_t' = (x_t, \pi_t) \text{ and } v_t' = (g_t, u_t)$$

indeterminacy arises when at least one eigenvalue of M is outside the unit circle. When the model is indeterminate, SSEs can be represented in different ways. VAR representations take the form

$$y_t = a + by_{t-1} + cv_t + dv_{t-1} + f\varepsilon_t$$

where $E_t \varepsilon_{t+1} = 0$, and noisy Markov SSEs can be represented as

$$y_t = as_t + bv_t,$$

where s_t a finite-state Markov process. The stability of SSEs under learning can in some cases depend on the representation, which can be interpreted as the econometric forecasting model.

In many cases with indeterminacy, SSEs in the NK model are not stable under learning. For example, if $i_t = \chi_\pi \pi_t + \chi_x x_t$ with $0 < \chi_\pi < 1$ there is indeterminacy and SSEs exist, but they are never stable under learning. However, there do exist cases of stable SSEs (in some representations) for the forward-looking rule $i_t = \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}$, with $\chi_x > 0$ sufficiently large and χ_π not too small. This was demonstrated for noisy finite-state Markov SSEs in Honkapohja and Mitra (2004). The result was generalized by Evans and McGough (2005b), who show that the noisy Markov SSEs are special cases of “Common Factor” representations in which the sunspot takes an AR(1) form. That is, s_t can be replaced by a sunspot $\zeta_t = \lambda \zeta_{t-1} + \varepsilon_t$ where ε_t is an exogenous martingale difference sequence and λ satisfies a “resonant frequency” condition. (For finite-state Markov processes the resonant frequency corresponds to specific transition probabilities).

Stable SSEs can also arise in representative-agent CA models when the government deficit is at least partially financed by seigniorage. Evans, Honkapohja, and Marimon (2007) consider a standard representative-agent model with cash goods, credit goods and variable labor supply. There is no capital, but agents can hold assets in the form of money or bonds, and there is a CA constraint for purchases of cash goods. Government spending g_t is assumed to be an exogenous *iid* process, and in the simplest version of the model g_t is entirely financed by seigniorage.

There are two regimes depending on the magnitude of the elasticity of intertemporal substitution (ITS). When ITS is high, there are two steady states, with differing inflation rates. This is a CA version of the hyperinflation model, which is discussed in the next section. When ITS is low there is a single steady state, and with sufficiently low ITS the steady state is indeterminate. Evans, Honkapohja, and Marimon (2007) show that in this case there are finite-state Markov sunspot equilibria that are stable under learning. These stable SSEs exhibit random variations over time in inflation and output in response to the extraneous sunspot variable. Evans, Honkapohja, and Marimon (2007) also examine the impact of changes in fiscal policy: a sufficient reduction in the mean of g_t , or a sufficient increase in taxes will in many cases eliminate the SSEs.

3.3 Hyperinflations and Liquidity Traps

3.3.1 Hyperinflations

Marcet and Nicolini (2003) aim to provide a unified theory to explain the empirical regularities of Latin American hyperinflations experienced by many countries in the 1980s using the seigniorage model of inflation. The model is based on the linear money demand equation

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

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

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

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

$$\begin{aligned} \frac{P_t}{P_{t-1}} &= \hat{T}(\beta_t, \beta_{t-1}; d) \equiv \frac{1 - \gamma\beta_{t-1}}{1 - \gamma\beta_t - d/\phi} \text{ where} \\ \beta_t &= P_{t+1}^e/P_t \text{ and } \beta_{t-1} = P_t^e/P_{t-1}. \end{aligned}$$

Under perfect foresight, there are two steady states, $\beta_L < \beta_H$, provided $d \geq 0$ is not too large. 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. The learning approach provides a different perspective.

Consider now the situation under adaptive learning. Suppose the PLM is that the inflation process is perceived to be a steady state, i.e. $P_{t+1}/P_t = \beta + \eta_t$, where η_t is perceived white noise. For this PLM expectations are $\left(\frac{P_{t+1}}{P_t}\right)^e = \beta$, all t , and the corresponding ALM is

$$\frac{P_t}{P_{t-1}} = \hat{T}(\beta, \beta; d) \equiv T(\beta; d).$$

The map $T(\beta; d)$ corresponds in Figure 3 to the part of $h(\beta, d)$ that lies below the value β_U . Under steady-state learning, agents estimate β based on past average inflation, i.e.

$$\beta_t = \beta_{t-1} + t^{-1}(P_{t-1}/P_{t-2} - \beta_{t-1}). \quad (29)$$

This is a recursive algorithm for the average inflation rate, which is equivalent to LS regression on a constant.¹⁰

Stability under this learning rule is governed by the E-stability differential equation

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

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 can be seen from Figure 3.

Marcet and Nicolini (MN) extend the preceding model to an open economy setting. They assume price flexibility with purchasing power parity (PPP), so that $P_t^f e_t = P_t$, where P_t^f is the exogenous foreign price of goods. A cash-in-advance constraint for local currency generates the money demand as in the basic model. d_t is assumed to be *iid*. There are two exchange rate regimes. In the floating regime the economy behaves just like the closed economy model, with PPP determining the price of foreign currency. In the exchange rate rule (ERR) regime, the government buys or sells foreign exchange as needed to guarantee $\frac{P_t}{P_{t-1}} = \bar{\beta} \equiv \beta_L$. The government imposes ERR if the inflation rate would otherwise exceed $\beta_U (> \beta_H)$, a maximum acceptable level.

MN argue that under RE the model cannot properly explain the main stylized facts of hyperinflation, and that a learning formulation is more successful. They use a variation of learning rule (29) in which t is replaced by α_t , where $\alpha_t = \alpha_{t-1} + 1$ if $\left| \left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1} \right) / \beta_{t-1} \right|$ falls below some bound, and otherwise $\alpha_t = \bar{\alpha}$, i.e. a constant gain is used. The qualitative features of the model are approximated by the system $\frac{P_t}{P_{t-1}} = h(\beta_{t-1}, d_t)$ where

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

Figure 3 describes the dynamics of system.

FIGURE 3 ABOUT HERE

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. This

¹⁰One can consider more general classes of PLMs. Adam, Evans, and Honkapohja (2006) study the circumstances in which autoregressive PLMs can converge to hyperinflation paths.

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

The model with learning has useful policy implications. ERR is valuable as a way of ending hyperinflations if the economy enters the explosive regime. However, a higher $E(d_t)$ makes average inflation higher and the frequency of hyperinflations greater. This indicates the importance of the orthodox policy of reducing deficits as a way of minimizing the likelihood of hyperinflation paths.

3.3.2 Liquidity Traps and Deflationary Spirals

Deflation and liquidity traps have at times been a concern. As we have seen, in contemporaneous Taylor rules, interest-rates should respond to the inflation rate more than one-for-one in order to ensure determinacy and stability under learning near the target inflation rate. However, as emphasized by Benhabib, Schmitt-Grohe, and Uribe (2001), if one considers the interest-rate rule globally, the requirement that net nominal interest rates must be nonnegative implies that the rule must be nonlinear and also, for any continuous rule, the existence of a second steady state at a lower (possibly negative) inflation rate. This is illustrated in Figure 4, which shows the interest-rate policy $R = 1 + f(\pi)$ as a function of π (a dependence on aggregate output is omitted for simplicity). The straight line in the figure is the Fisher equation $R = \pi/\beta$, which is obtained from the usual Euler equation for consumption in a steady state.

FIGURE 4 ABOUT HERE

Here R stands for the interest rate factor (the net interest rate is $R-1$) and $\pi_t = P_t/P_{t-1}$ for the inflation factor ($\pi-1$ is the net inflation rate). In Figure 4, π^* denotes the intended steady state, at which the “Taylor principle” of a more than one-for-one response is satisfied, and π_L is the unintended steady state. π_L may correspond to either a very low positive or a negative net inflation rate, i.e. deflation. The zero lower bound corresponds to $R = 1$. Benhabib, Schmitt-Grohe, and Uribe (2001) show that under RE, there is a continuum of “liquidity trap” paths that converge on π_L . The pure RE

analysis thus suggests a serious risk of the economy following these “liquidity trap” paths.

What happens under learning? Evans and Honkapohja (2005) analyzed a flexible-price perfect competition model. We showed that deflationary paths are possible, but that the real risk, under learning, were paths in which inflation slipped below π_L and then continued to fall further. For this flexible-price model we showed that this could be avoided by a switch to an aggressive money supply rule at low inflation rates.

Evans, Guse, and Honkapohja (2008) reconsider the issues in an NK model with sticky prices due to adjustment costs and deviations of output from flexible-price levels. 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 equations are nonlinear, and the nonlinearity in its analysis under learning is retained. The key equations are

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

The first equation is the nonlinear NK Phillips curve and the second equation is the IS curve. 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 *iid* then “steady-state” learning is appropriate for both c^e and π^e , i.e.

$$\begin{aligned} \pi_{t+1}^e &= \pi_t^e + \phi_t (\pi_{t-1} - \pi_t^e) \\ c_{t+1}^e &= c_t^e + \phi_t (c_{t-1} - c_t^e), \end{aligned}$$

where ϕ_t is the gain sequence. The intended steady state π^* is locally stable under learning, while the unintended steady state π_L is unstable. The key observation is that π_L is a saddlepoint, which implies the existence of deflationary spirals under learning. In particular, after a sufficiently pessimistic expectational shock, c^e, π^e will follow paths leading to deflation and stagnation. This is illustrated in Figure 5, giving the E-stability dynamics.

FIGURE 5 ABOUT HERE

For the intuition, suppose that we are initially near the π_L steady state and consider a small drop in π^e . With fixed R this would lead through the IS curve to lower c and thus, through the Phillips curve, to lower π . Because only small reductions in R are possible given the global Taylor rule, the reduction in c and π cannot be offset. The falls in realized c and π lead, under learning, to reductions in c^e and π^e , and this sets in motion the deflationary spiral.

Thus, large adverse shocks to expectations or structural changes can set in motion unstable downward paths. Can policy be altered to avoid deflationary spiral? Evans, Guse, and Honkapohja (2008) show that it can. The recommended policy is to set a minimum inflation threshold $\tilde{\pi}$, where $\pi_L < \tilde{\pi} < \pi^*$. The authorities would follow normal monetary and fiscal policy provided this delivers $\pi_t > \tilde{\pi}$. However, if π_t threatens to fall below $\tilde{\pi}$, 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. It can be shown that these policies can indeed ensure $\pi_t \geq \tilde{\pi}$ always under learning and lead to global stability of the intended steady state at π^* . Perhaps surprisingly, it is essential to have an *inflation* threshold. Using instead an output threshold to trigger aggressive policies will not always avoid deflationary spirals.

3.4 Asset Prices

Asset pricing is another area of recent focus in the learning literature. The potential for adaptive learning to generate new phenomena for asset prices was already apparent in the early work of Timmermann (1993) and Timmermann (1996). Consider the standard risk-neutral asset-pricing framework

$$p_t = \beta E_t^*(p_{t+1} + d_{t+1}), \quad (30)$$

where p_t is the real price of equalities, d_{t+1} is the real dividend paid at the end of period $t+1$, and $0 < \beta \equiv (1+r)^{-1} < 1$ is the discount factor, assumed constant. Assume also that d_t is an exogenous stochastic process, e.g.

$$\ln(d_t) = \mu + \ln(d_{t-1}) + \varepsilon_t, \quad (31)$$

where $\varepsilon_t \stackrel{iid}{\sim} N(0, \sigma^2)$ and $\mu + \sigma^2/2 < \ln(1+r)$. Under RE the “fundamentals” solution is

$$p_t = \frac{1+g}{r-g} d_t, \text{ where } 1+g = \exp(\mu + \sigma^2/2). \quad (32)$$

Iterating (30) forward, one obtains the present value formula

$$p_t = \sum_{j=1}^{\infty} \beta^j E_t^* d_{t+j}, \quad (33)$$

and imposing RE yields (32).

There are a number of empirical puzzles in asset-pricing based on this model, including excess volatility of stock prices and predictability of stock returns. A potentially simple and appealing explanation is that traders do not have *a priori* knowledge of the parameters of the (p_t, d_t) process. The equilibrium price-process under learning is generated, as usual, assuming that the parameter estimates are updated over time as new data become available.

There are two natural ways to model stock prices under learning, depending on whether we want to treat (33) or (30) as the key equation that determines p_t , given expectations. If traders are “fundamentalists” then price will be set in accordance with (33), based on forecasts $E_t^* d_{t+j} = d_t \exp(j\hat{\mu}_t + j\hat{\sigma}_t^2/2)$, where $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ are the time t estimates of μ and σ^2 . This leads to

$$p_t = \frac{\exp(\hat{\mu}_t + \hat{\sigma}_t^2/2)}{1 + r - \exp(\hat{\mu}_t + \hat{\sigma}_t^2/2)} d_t.$$

This approach is investigated by Timmermann (1993) and Timmermann (1996) under the name “present-value learning.”

An alternative approach is to assume that p_t is determined by the expected rate of return over the coming period in accordance with (30). Traders would then also estimate a model for p_t , e.g. $p_t = a + \lambda d_{t-1} + \eta_t$. Using estimates a_t , λ_t , $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ traders form forecasts $E_t^* d_{t+1}$ and $E_t^* p_{t+1}$, with p_t determined by (30). This “self-referential learning” approach was also studied in Timmermann (1996).

For concreteness, consider the present-value learning approach. For the dividend process (31), “steady-state” learning rules suffice, i.e.

$$\begin{aligned} \hat{\mu}_t &= \frac{n-1}{n} \hat{\mu}_{t-1} + \frac{1}{n} (\ln(d_t) - \ln(d_{t-1})) \\ \hat{\sigma}_t^2 &= \frac{n-1}{n^2} [n\hat{\sigma}_{t-1}^2 + (\hat{\mu}_{t-1} - (\ln(d_t) - \ln(d_{t-1})))^2], \end{aligned}$$

where n is the sample size. Because the dividend process is exogenous, standard asymptotic statistical results apply, so that $\hat{\mu}_t \rightarrow \mu$ and $\hat{\sigma}_t^2 \rightarrow \sigma^2$ as $n \rightarrow \infty$. Thus, the price process converges asymptotically to RE. However,

during the learning transition there will be substantial excess volatility for a substantial period of time.

Calibrating the model using US annual dividend data from Standard and Poors, and using an initial prior in each simulation based on a sample size of 10, Timmermann (1993) finds that for sample sizes of around 40, the extent of gross violations of the Shiller-type volatility bounds is in the 30 to 50% range. This effect drops away rapidly as the sample size increases. This suggests that using learning to explain excess volatility requires occasional structural shifts in the dividend process, leading agents to reduce their effective sample size. On the other hand, the “predictability anomaly” – that excess returns are predictable by the lagged dividend yield – arises even with large sample sizes.

More recent work on learning and stock prices has extended the framework in several directions. Brock and Hommes (1998) introduce heterogeneous expectations using the dynamic predictor selection methodology discussed earlier. Branch and Evans (2008) and Adam, Marcet, and Nicolini (2008) both focus on self-referential learning.

Adam, Marcet, and Nicolini (2008) use the consumption-based version of model (30). Forecasts are given by $E_t^* p_{t+1} = b_t p_t$, where b_t is updated according to

$$b_t = b_{t-1} + (t + K)^{-1}(p_{t-1}/p_{t-2} - b_{t-1}),$$

for $K \geq 0$ given, modified by a projection facility that bounds estimates of b_t to ensure positive, finite prices. The dividend process is assumed known, with forecasts set at the true conditional expectation. The learning transition exhibits mean reversion of returns, excess volatility, and persistence of price-dividend ratios. A calibrated version of the model is shown to match many aspects of US data.

Branch and Evans (2008) examine learning within a mean-variance linear model. Agents choose between a risk-free asset that pays fixed rate of return $R = \beta^{-1} > 1$ and a risky stock. The supply of the stock is exogenous and random. Demand depends positively on expected excess returns but negatively risk, as measured by expected conditional variance of returns. Equating supply and demand for the risky asset leads to

$$p_t = \beta E_t^*(p_{t+1} + d_{t+1}) - \beta a \sigma_t^2 z_{st}$$

where z_{st} is asset supply, $a \geq 0$ is the absolute risk-aversion parameter and σ_t^2 is the time t estimate of the conditional variance. The dividend process

is assumed exogenous and known and forecasts $E_t^* d_{t+1}$ are set at the true conditional expectation. $E_t^* p_{t+1}$ is generated from estimates of the price process

$$p_t = k_t + c_t p_{t-1} + \varepsilon_t.$$

There is a fundamentals REE with fixed parameters $(\bar{k}, \bar{c}, \bar{\sigma}^2)$ and this REE is stable under LS learning. Interesting dynamics arise if agents instead estimate the parameters (k_t, c_t, σ_t^2) using constant-gain LS learning. There are occasional “escapes” to non-fundamental random-walk behavior of asset prices, in which agents’ estimates are close to $(k_t, c_t) = (0, 1)$. In this regime, there is bubble-like behavior in p_t . However, subsequent revisions in the estimates of risk eventually lead to crashes back to fundamentals values. Thus, learning about both returns and risk can lead to recurrent bubbles and crashes.

Exchange rate dynamics also exhibit a number of puzzles that learning models may help to resolve.

Chakraborty and Evans (2008) focus on the forward-premium puzzle. Letting s_t be the log of the price of foreign currency and F_t the log of the forward rate at t for foreign currency at $t + 1$, under RE and risk neutrality we have $\alpha = 0$ and $\beta = 1$ in the forward-premium regression

$$s_{t+1} - s_t = \alpha + \beta(F_t - s_t) + u_{t+1}.$$

However, in practice estimates of b are substantially less than one and often negative.

Using the benchmark monetary exchange rate model, Chakraborty and Evans (2008) argue that this anomaly can be explained by learning. The reduced form is

$$s_t = \theta E_t^* s_{t+1} + v_t,$$

where fundamentals v_t are modeled as $v_t = \rho v_{t-1} + \varepsilon_t$. The REE is $s_t = b v_{t-1} + c \varepsilon_t$, where $b = \bar{b} = (1 - \rho\theta)^{-1}\rho$. Agents estimate the parameter b by constant-gain LS and make forecasts $E_t^* s_{t+1} = b_t v_t$, where b_t is their time t estimate of b . For small gains b_t converges to a stochastic process distributed with high probability in a small neighborhood of \bar{b} . Surprisingly, this implies a strong downward bias in the forward-premium regression for the realistic case of ρ near 1. In particular, $\text{plim}(\hat{\beta}) \rightarrow 0$ as $\rho \rightarrow 1$, and for realistic sample sizes the median value of $\hat{\beta}$ is negative, in line with the data.

Another, potentially complementary, approach to exchange-rate modeling is based on dynamic predictor selection, see De Grauwe and Grimaldi (2006).

Further applications of learning to exchange rates include Kasa (2004), Kim (2008), Mark (2007), and Markiewicz (2008).

4 Concluding Remarks

The adaptive learning approach to macroeconomics treats economic agents – firms, households and policymakers – like econometricians when modeling how they make forecasts. Macroeconomic models under learning, in which estimated forecast rules are updated in accordance LS or other statistical rules, can be analyzed using stochastic approximation techniques, and E-stability provides a key tool for analyzing the dynamics under LS learning. Applications of learning in macroeconomics have expanded rapidly in recent years, and the learning approach has provided novel insights to central issues of monetary policy, business cycles, and asset pricing.

There are a number of current research areas that, for reasons of space, our review has not covered, but which are likely to become more prominent. Empirical work on forecasts, based on survey data, experiments, and indirect measures from asset markets, will help to assess the alternative models of learning and expectations formation. For recent papers, see Branch (2004), Adam (2007) and Pfajfar and Santoro (2007). In addition to the empirical topics covered in our review, there are applications of learning to several other areas. One new example is DSGE models, see Milani (2007) and Slobodyan and Wouters (2007). Alternatives to econometric approaches to learning include those based on genetic algorithm learning and evolutionary dynamics, e.g., see Arifovic (2000) and Georges and Wallace (2007).

There are also a number of unresolved conceptual questions that merit further investigation. One set of issues revolves around heterogeneous expectations, model selection and Bayesian model averaging. Heterogeneity in forecasting and multiple models are clearly evident in reality, but this contrasts with Bayesian approaches that suggest eventual convergence to a single model. Another major area concerns the use of structural information and forward-looking reasoning in forecasting by private agents. The econometric approach to learning currently emphasizes forecasting using reduced-form models, but in many cases it would be natural for agents to incorporate structural knowledge in their forecasting.

References

- ADAM, K. (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *Economic Journal*, 117, 603–636.
- ADAM, K., G. W. EVANS, AND S. HONKAPOHJA (2006): “Are Hyperinflation Paths Learnable?,” *Journal of Economic Dynamics and Control*, 30, 2725–2748.
- ADAM, K., A. MARCET, AND J. P. NICOLINI (2008): “Stock Market Volatility and Learning,” Working paper.
- ARIFOVIC, J. (2000): “Evolutionary Algorithms in Macroeconomic Models,” *Macroeconomic Dynamics*, 4, 373–414.
- BENHABIB, J., AND R. E. FARMER (1996): “Indeterminacy and Sector-Specific Externalities,” *Journal of Monetary Economics*, 37, 421–443.
- BENHABIB, J., S. SCHMITT-GROHE, AND M. URIBE (2001): “The Perils of Taylor Rules,” *Journal of Economic Theory*, 96, 40–69.
- BERNANKE, B., AND M. WOODFORD (eds.) (2005): *The Inflation-Targeting Debate*. University of Chicago Press, Chicago, USA.
- BRANCH, W. A. (2004): “The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations,” *Economic Journal*, 114, 592–621.
- BRANCH, W. A., AND G. W. EVANS (2006): “Intrinsic Heterogeneity in Expectation Formation,” *Journal of Economic Theory*, 127, 264–295.
- (2007): “Model Uncertainty and Endogenous Volatility,” *Review of Economic Dynamics*, 10, 207–237.
- (2008): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” Working paper.
- BROCK, W. A., AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65, 1059–1095.

- (1998): “Heterogeneous Beliefs and Routes to Chaos in a Simple Asset pricing Model,” *Journal of Economic Dynamics and Control*, 22, 1235–1274.
- BULLARD, J. (2006): “The Learnability Criterion and Monetary Policy,” *Federal Reserve Bank of St. Louis Review*, 88, 203–217.
- BULLARD, J., AND S. EUSEPI (2005): “Did the Great Inflation Occur Despite Policymaker Commitment to a Taylor Rule?,” *Review of Economic Dynamics*, 8, 324–359.
- (2008): “When does Determinacy imply Expectational Stability?,” manuscript.
- BULLARD, J., AND K. MITRA (2002): “Learning About Monetary Policy Rules,” *Journal of Monetary Economics*, 49, 1105–1129.
- CHAKRABORTY, A., AND G. W. EVANS (2008): “Can Perpetual Learning Explain the Forward Premium Puzzle?,” *Journal of Monetary Economics*, 55, 477–490.
- CHIAPPORI, P., AND R. GUESNERIE (1991): “Sunspot Equilibria in Sequential Market Models,” in Hildenbrand and Sonnenschein (1991), pp. 1683–1762.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1–40.
- CLARIDA, R., J. GALI, AND M. GERTLER (1999): “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature*, 37, 1661–1707.
- COGLEY, T., AND T. J. SARGENT (2005): “The Conquest of US Inflation: Learning and Robustness to Model Uncertainty,” *Review of Economic Dynamics*, 8, 528–563.
- COOLEY, T. F. (ed.) (1995): *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, N.J.
- DE GRAUWE, P., AND M. GRIMALDI (2006): *The Exchange Rate in a Behavioral Finance Framework*. Princeton University Press, Princeton.

- DUFFY, J., AND W. XIAO (2007): “Instability of Sunspot Equilibria in Real Business Cycle Models under Adaptive Learning,” *Journal of Monetary Economics*, 54, 879–903.
- EUSEPI, S., AND B. PRESTON (2008): “Expectations, Learning and Business Cycle Fluctuations,” mimeo.
- EVANS, G. W. (1989): “The Fragility of Sunspots and Bubbles,” *Journal of Monetary Economics*, 23, 297–317.
- EVANS, G. W., E. GUSE, AND S. HONKAPOHJA (2008): “Liquidity Traps, Learning and Stagnation,” *European Economic Review*, forthcoming.
- EVANS, G. W., AND S. HONKAPOHJA (1992): “On the Robustness of Bubbles in Linear RE Models,” *International Economic Review*, 33, 1–14.
- (1994): “On the Local Stability of Sunspot Equilibria under Adaptive Learning Rules,” *Journal of Economic Theory*, 64, 142–161.
- (1995): “Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models,” *Econometrica*, 63, 195–206.
- (1996): “Least Squares Learning with Heterogeneous Expectations,” *Economics Letters*, 53, 197–201.
- (1999): “Learning Dynamics,” in Taylor and Woodford (1999), chap. 7, pp. 449–542.
- (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2003a): “Adaptive Learning and Monetary Policy Design,” *Journal of Money, Credit and Banking*, 35, 1045–1072.
- (2003b): “Existence of Adaptively Stable Sunspot Equilibria near an Indeterminate Steady State,” *Journal of Economic Theory*, 111, 125–134.
- (2003c): “Expectations and the Stability Problem for Optimal Monetary Policies,” *Review of Economic Studies*, 70, 807–824.
- (2005): “Policy Interaction, Expectations and the Liquidity Trap,” *Review of Economic Dynamics*, 8, 303–323.

- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.
- (2007): “Expectations, Learning and Monetary Policy: An Overview of Recent Research,” mimeo.
- EVANS, G. W., S. HONKAPOHJA, AND R. MARIMON (2007): “Stable Sunspot Equilibria in a Cash-in-Advance Economy,” *The B.E. Journal of Macroeconomics* (Advances), 7, Iss. 1, Article 3.
- EVANS, G. W., S. HONKAPOHJA, AND K. MITRA (2007): “Anticipated Fiscal Policy and Learning,” mimeo.
- EVANS, G. W., S. HONKAPOHJA, AND P. ROMER (1998): “Growth Cycles,” *American Economic Review*, 88, 495–515.
- EVANS, G. W., S. HONKAPOHJA, AND N. WILLIAMS (2008): “Generalized Stochastic Gradient Learning,” mimeo.
- EVANS, G. W., AND B. MCGOUGH (2005a): “Indeterminacy and the Stability Puzzle in Non-Convex Economies,” *The B.E. Journal of Macroeconomics* (Contributions), 5, Iss. 1, Article 8.
- (2005b): “Monetary Policy, Indeterminacy and Learning,” *Journal of Economic Dynamics and Control*, 29, 1809–1840.
- (2008): “Learning to Optimize,” mimeo.
- FARMER, R. E. (1999): *The Economics of Self-Fulfilling Prophecies, Second edition*. MIT Press, Cambridge, Mass.
- FARMER, R. E., AND J.-T. GUO (1994): “Real Business Cycles and the Animal Spirits Hypothesis,” *The Journal of Economic Theory*, 63, 42–72.
- GALÍ, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton New Jersey.
- GEORGES, C., AND J. C. WALLACE (2007): “Learning Dynamics and Non-linear Misspecification in an Artificial Financial Market,” manuscript.
- GIANNITSAROU, C. (2003): “Heterogenous Learning,” *Review of Economic Dynamics*, 6, 885–906.

- (2006): “Supply-Side Reforms and Learning Dynamics,” *Journal of Monetary Economics*, 53, 291–309.
- GRANDMONT, J.-M. (1998): “Expectations Formation and Stability of Large Socioeconomic Systems,” *Econometrica*, 66, 741–781.
- GUESNERIE, R. (2005): *Assessing Rational Expectations 2: Eductive Stability in Economics*. MIT Press, Cambridge, Mass.
- HILDENBRAND, W., AND H. SONNENSCHNEIN (eds.) (1991): *Handbook of Mathematical Economics, Vol. IV*. North-Holland, Amsterdam.
- HONKAPOHJA, S., AND K. MITRA (2004): “Are Non-Fundamental Equilibria Learnable in Models of Monetary Policy?,” *Journal of Monetary Economics*, 51, 1743–1770.
- (2006): “Learning Stability in Economies with Heterogeneous Agents,” *Review of Economic Dynamics*, 9, 284–309.
- HONKAPOHJA, S., K. MITRA, AND G. W. EVANS (2002): “Notes on Agents’ Behavioral Rules Under Adaptive Learning and Recent Studies of Monetary Policy,” manuscript, www.valt.helsinki.fi/RAKA/seppo.html, and <http://darkwing.uoregon.edu/~gevens>.
- HOWITT, P., AND R. P. MCAFEE (1992): “Animal Spirits,” *American Economic Review*, 82, 493–507.
- HUANG, K., Z. LIU, AND T. ZHA (2008): “Learning, Adaptive Expectations and Technology Shocks,” Working paper.
- KASA, K. (2004): “Learning, Large Deviations, and Recurrent Currency Crises,” *International Economic Review*, 45, 141–173.
- KIM, Y. S. (2008): “Exchange Rates and Fundamentals under Adaptive Learning,” mimeo.
- KREPS, D., AND K. WALLIS (eds.) (1997): *Advances in Economics and Econometrics: Theory and Applications, Volume I*. Cambridge University Press, Cambridge.
- LJUNGQVIST, L., AND T. J. SARGENT (2003): *Recursive Macroeconomic Theory, Second edition*. MIT Press, Cambridge, MA.

- LUCAS, JR., R. E. (1972): “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 4, 103–124.
- (1973): “Some International Evidence on Output-Inflation Trade-offs,” *American Economic Review*, 63, 326–334.
- MARCET, A., AND J. P. NICOLINI (2003): “Recurrent Hyperinflations and Learning,” *American Economic Review*, 93, 1476–1498.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.
- MARIMON, R. (1997): “Learning from Learning in Economics,” in Kreps and Wallis (1997), chap. 9, pp. 278–315.
- MARK, N. (2007): “Changing Monetary Rules, Learning, and Real Exchange Rate Dynamics,” manuscript.
- MARKIEWICZ, A. (2008): “Monetary Policy, Model Uncertainty, and Exchange Rate Volatility,” manuscript.
- MCCALLUM, B. T. (2007): “E-stability vis-à-vis Determinacy Results for a Broad Class of Linear Rational Expectations Models,” *Journal of Economic Dynamics and Control*, 31, 1376–1391.
- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal Of Monetary Economics*, 54, 2065–2082.
- MUTH, J. F. (1961): “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315–335.
- 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.
- (2007): “Robust Monetary Policy with Imperfect Knowledge,” *Journal of Monetary Economics*, 54, 1406–1435.

- PFAJFAR, D., AND E. SANTORO (2007): “Heterogeneity and Learning in Inflation Expectation Formation: An Empirical Assessment,” mimeo.
- PRESTON, B. (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1, 81–126.
- (2006): “Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy,” *Journal of Monetary Economics*, 53, 507–535.
- PRIMICERI, G. E. (2006): “Why Inflation Rose and Fell: Policy-Makers’ Beliefs and U. S. Postwar Stabilization Policy,” *Quarterly Journal of Economics*, 121, 867–901.
- SARGENT, T. J. (1973): “Rational Expectations, the Real Rate of Interest and the Natural Rate of Unemployment,” *Brookings Papers on Economic Activity*, 2, 429–472.
- (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98, 5–37.
- SARGENT, T. J., N. WILLIAMS, AND T. ZHA (2006): “Shocks and Government Beliefs: The Rise and Fall of American Inflation,” *American Economic Review*, 96, 1193–1224.
- SCHMITT-GROHE, S., AND M. URIBE (1997): “Balanced-Budget Rules, Discretionary Taxes and Aggregate Instability,” *Journal of Political Economy*, 105, 976–1000.
- SLOBODYAN, S., AND R. WOUTERS (2007): “Learning in an Estimated DSGE Model,” mimeo.
- TAYLOR, J. (1993): “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series in Public Policy*, 39, 195–214.
- TAYLOR, J., AND M. WOODFORD (eds.) (1999): *Handbook of Macroeconomics, Volume 1*. Elsevier, Amsterdam.

- TIMMERMANN, A. G. (1993): “How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices,” *Quarterly Journal of Economics*, 108, 1135–1145.
- (1996): “Excessive Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning,” *Review of Economic Studies*, 63, 523–557.
- VAN NIEUWERBURGH, S., AND L. VELDKAMP (2006): “Learning Asymmetries in Real Business Cycles,” *Journal of Monetary Economics*, 53, 753–772.
- WALSH, C. (2003): *Monetary Theory and Policy*. MIT Press, Cambridge, Massachusetts, second edn.
- WILLIAMS, N. (2004a): “Adaptive Learning and Business Cycles,” working paper.
- (2004b): “Escape Dynamics in Learning Models,” Discussion paper, working paper, Princeton University.
- WOODFORD, M. (1990): “Learning to Believe in Sunspots,” *Econometrica*, 58, 277–307.
- (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ.




Figure 1: 2-cycle and steady state




Figure 2: SSE and multiple steady states




Figure 3: Inflation as a function of expected inflation




Figure 4: Multiple steady states with global Taylor rule




Figure 5: Expectation dynamics under normal policy