Adaptive Learning and Macroeconomics

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

Bruce McGough
University of Oregon

February 2, 2020

Abstract

While rational expectations remains the benchmark paradigm in macro-economic modeling, bounded rationality, especially in the form of adaptive learning, has become a mainstream alternative. Under the adaptive-learning (AL) approach, economic agents in dynamic, stochastic environments are modeled as adaptive learners making decisions based on forecasting rules that are updated in real time as new data become available. Their decisions are then coordinated each period via the economy's markets and other relevant institutional architecture, resulting in a time-path of economic aggregates. In this way, the AL approach introduces additional dynamics into the model – dynamics that can be used to address myriad macroeconomic issues and concerns, including, for example, empirical fit and the plausibility of specific rational expectations equilibria.

In this chapter we introduce the basic concepts of adaptive learning in macroeconomics, and we provide a number of applications to illustrate the usefulness of the adaptive learning approach. We focus our attention on *reduced-form learning*, that is, the implementation of learning at the aggregate level; our companion chapter, Evans and McGough (2020a), covers *agent-level learning*, which includes pre-aggregation analysis of boundedly rational decision making.

In section 2 of this chapter we develop the needed language and mechanics using a simple cobweb model as our laboratory; we then turn, in section 3, to a more general multivariate setting to explore the implementation of adaptive learning in models commonly used for applied macroeconomics, including an examination of a New-Keynesian model. Section 4 develops the use of the AL approach as an equilibrium selection mechanism by examining the stability under learning of the sunspot equilibria present in models with indeterminacies. Section 5 provides a discussion of some of the many applications of, and extensions to adaptive learning in macroeconomics.

1 Introduction

Macroeconomic models are usually based in recursive, stochastic settings and can be summarized by dynamic systems that include expectational dependencies. In the simplest case of point expectations and representative agents, this could take the form

$$y_t = Q(y_{t-1}, y^e_{t+1}, w_t),$$

where y_t is a vector of endogenous variables at time t (unemployment, inflation, investment, etc.), y^e_{t+1} denotes the expectations formed at time t of the values these variables take at time $t+1$, and w_t is an exogenous vector of random factors at t. A more general formulation

$$E^*_t F(y_t, y_{t-1}, y_{t+1}, w_t) = 0 \tag{1}$$

allows for expectations of nonlinear interactions of future variables, where E^*_t denotes the subjective expectation of the representative agent. In many models heterogeneity across agents is also important, as we discuss later.

The presence of expectations, which typically arises from the (possibly implicit) assumption that the model's forward-looking agents solve dynamic programing problems, makes macroeconomics inherently different from natural science, and thus necessitates distinct solution concepts and methods.[1] Solving a model of the form (1) requires taking a stand on how agents form expectations. The benchmark "rational expectations" (RE) assumption in macroeconomics is that agents' expectations are formed optimally given their information; said differently, E^*_t is taken to be the mathematical expectations operator conditional on period t information. RE, as it is commonly implemented, requires that agents form their expectations against the objective conditional distribution of the model's variables, including y_{t+1}. However, this distribution is endogenous, and in particular depends on the manner in which all agents form expectations. Thus RE, as usually implemented in macroeconomics, is an equilibrium object: it is not (in general) rational to have rational expectations unless everyone else has rational expectations.

An alternative to RE is the adaptive learning approach, a perspective on expectation formation that views the standard assumption of RE as demanding too much knowledge of, and coordination by, economic agents: the endogenous distribution characterizing a rational expectations equilibrium (REE) is an unobserved solution to an infinite-dimensional system of equations that, in general, even the modeler can only approximate; furthermore, a mechanism would need to be specified for how agents could coordinate on the REE. Clearly we need a more realistic model of behavior. What form should this take?

One answer to this question is given by the *Cognitive Consistency Principle*: economic agents should be about as smart as economists. While this principle might be implemented in various ways, we focus here on the adaptive learning (AL) approach, which takes agents

[1] Evans and Honkapohja (2009b) and Woodford (2013) provide recent surveys that emphasize expectations, learning and bounded rationality in macroeconomics.

as revising their forecasting models and decision rules as new data become available. We will usually assume agents act as time-series econometricians, updating their model parameters over time, but the adaptive learning approach also includes, for example, selection among alternative behavioral rules based on recent forecast success. Reflecting our position that RE implicitly requires agents to have an unreasonable degree of knowledge about the structure of the economy, the adaptive learning approach instead typically assumes that agents make forecasts by simply regressing the variables being forecasted on relevant observed exogenous or lagged endogenous variables, and updating the estimated coefficients over time.

1.1 Reduced-form learning versus the agent-level approach

The implementation of AL in ad hoc macro models, i.e. models that take as primitive the relationships between aggregates, is straightforward in concept: simply replace the conditional expectations operator with a boundedly rational counterpart. Micro-founded dynamic stochastic general equilibrium (DSGE) environments allow for this method of implementation as well, but also invite a more nuanced approach that explicitly models agent behavior. Put more succinctly, in micro-founded models the AL approach can be implemented at the aggregate (reduced-form) level or at the agent level. In this chapter we focus on *reduced-form learning* (RFL), i.e on implementation at the aggregate level, and we leave agent-level considerations for our companion chapter, Evans and McGough (2020a). However, to motivate RFL, and to draw distinction between it and agent-level implementations, it is helpful to begin with a broader discussion that includes boundedly optimal decision making.

The modern, micro-founded DSGE approach to macroeconomic modeling and analysis includes a common architecture comprised of many agents making decisions in dynamic, stochastic environments, with the per-period determination of aggregates coordinated by markets and possibly other institutional arrangements. This per-period determination, referred to as *temporary equilibrium*, results in a time-indexed collection of equations and inequalities characterizing the model's endogenous aggregates as functions of the exogenous drivers.[2]

The common architecture described above must include assumptions detailing *how* agents make decisions. The benchmark is to adopt the *rational expectations hypothesis* which stipulates that agents forecast optimally given their information and make decisions optimally given their forecasts – in effect, to make their period t decisions agents are assumed to solve a dynamic programming problem, taking into account the stochastic evolution of all the relevant state variables, including the economy's endogenous aggregates. Temporary equilibrium is then obtained by imposing market clearing and aggregating.

Because agents' behavior depends on their expectations of future aggregates, which, themselves, depend on agents' behavior, an additional equilibrium concept is needed to guarantee internal consistency. A *rational expectation equilibrium* (REE) is a distribution

[2] The notion of temporary equilibrium was introduced by Hicks (1946).

over the collections of time-paths of endogenous variables that is consistent with rational decision-making in this sense that if agents condition on the distribution then their attendant decisions will aggregate to comport with the distribution.

Explicit, closed-form computation of an REE is very challenging – indeed, it is not possible under most circumstances. In the discrete-time case, the most common approximation approach involves simplifying the temporary equilibrium restrictions to obtain a system of non-linear, expectational difference equations – the reduced-form system – and then linearizing these equations around a non-stochastic steady state, thus yielding a system of linear, expectational difference equations – the linearized reduced-form system – and then finally using well-known techniques to solve this linear RE-model.

The adaptive learning literature questions the wisdom of adopting the rational expectations hypothesis, which itself is tantamount to the *assumption* that the economy is always in an REE. Instead, the AL literature prefers the view that an REE is a possible outcome that might, or might not, *emerge* as a result of more realistically modeled agent-level decision making.

A natural implementation of adaptive learning thus begins by discarding the rational expectations hypothesis, and instead developing explicit models of boundedly rationality, including alternate specifications of forecasting models, planning horizons and decision rules. The implied decisions are then coordinated and aggregated in temporary equilibrium, thus resulting in a dynamic system that can at least be simulated if not analytically assessed. This is the agent-level approach to adaptive learning, and includes, for example, Euler-equation learning (Evans and Honkapohja (2006)), long-horizon learning (Preston (2005)), and the shadow-price approach developed by Evans and McGough (2018c). Details of agent-level learning are explored in Evans and McGough (2020a).

An alternative and simpler approach is reduced-form learning. Under this implementation, the reduced-form system of expectational difference equations, derived under rational expectations, is taken as given. The modeler then simply replaces the conditional expectations operators in this reduced form system with a boundedly rational counterpart, and proceeds with analysis. In fact, this replacement is often done after the reduced-form system is linearized.

RFL can be viewed as a reasonable and effective short cut that can greatly simplify the analysis of a given DSGE model; and, if needed, it can often be justified via an agent-level approach; and it is particularly convenient when conducting empirical work. The agent-level approach is more appropriate for policy analysis, particularly in models with complex agent-level behaviors that are expected to impinge on policy outcomes.

1.2 The appeal of adaptive learning

By anchoring to behavioral assumptions that are more plausible than those required by the rational expectations hypothesis, the AL approach has considerable intellectual appeal. Additional attractive features include:

1. AL provides a test of the plausibility of an REE. Because the economy is self-referential in the sense that the evolution of the economy depends on the expectations of agents and vice versa, an REE is most naturally viewed as a possible emergent outcome of this evolution. AL provides a natural mechanism through which an REE might emerge, and thereby provides a "plausibility test" of the REE.

2. Many macro models have multiple REE. With multiple REE it is common for one or more of the REE to be (locally) stable under AL while stability fails to hold for other REE. In these circumstances AL acts as a selection criterion.

3. Under AL there are learning dynamics, at least in the transition to RE. Furthermore, natural relaxations of our adaptive learning rules can lead to persistence in these dynamics, which can help fit empirical regularities.

4. In most standard macro models, RE implies homogenous expectations about key economic aggregates. This is counterfactual: survey evidence always shows substantial heterogeneity. AL can be easily modified in simple and natural ways to allow for heterogeneous expectations.

5. The temporary equilibrium AL approach is readily implemented computationally, and is often more tractable than RE, especially in nonlinear environments.

1.3 Alternatives to adaptive learning

While our focus is on the AL approach summarized above, it is important to note that there are related approaches to expectation formation that are distinct but complementary. The eductive approach, introduced by Guesnerie (1992), considers conditions under which fully rational agents, with common knowledge both of the economic structure and of the rationality of other agents, would be able through mental reasoning to coordinate on an REE. The conditions for coordination, which can be quite strict, are related to the iterative formulation of expectational stability used by Evans (1985); see the discussion in Evans and Guesnerie (1993). Applications of the eductive stability or "strong rationality" approach can be found in Guesnerie (2005). For an exploration of the complementary insights of eductive and adaptive approaches in the context of the Real Business Cycle model, see Evans, Guesnerie, and McGough (2019).

A related model of expectation formation is the "k-level" approach of Nagel (1995).[3] As in the eductive approach, agents need structural knowledge of the economy and engage in higher level reasoning, but different agents may choose to use different levels of reasoning. Evans, Gibbs, and McGough (2019) provides a model that combines adaptive and k-level reasoning in which agents may switch their choice of k-level based on past performance. Finally, as discussed in Section 2.1, there is a range of behavioral approaches in which agents make decisions based on forecasts made using simple rules of thumb. To the extent that agents choose rules based on past performance, this is consistent with the spirit of the AL approach.

1.4 Chapter organization

This overview of adaptive learning in macroeconomics begins in Section 2 with a description of the key equilibrium and stability tools, and related techniques, within a particularly simple framework. Section 3 shows how these tools can be extended to a general multivariate linear setting, using the reduced-form short-cut, and illustrates results within the context of monetary policy in New Keynesian setting. Section 4 discusses models in which there can be multiple REE, shows how AL can be used as a selection criterion, and assesses whether there can exist stable "sunspot equilibria" in which economic fluctuations are driven by self-fulfilling expectations. Section 5 concludes with a survey of empirical papers that use AL to study policy and related issues.

2 Adaptive Learning in the Cobweb Model

Developing the AL approach requires some initial investment in tools. We begin by illustrating the key AL concepts using a simple model: the linear "cobweb" model used by Muth (1961) in his seminal formulation of RE. The cobweb model considers an isolated market for a perishable good. Demand for the good is linear, taking the form

$$D_t = A - Bp_t + v_{1t}, \text{ where } A, B > 0,$$

and there is a continuum of firms $i \in [0,1]$ that supply output in a competitive market. Production is subject to a one period delay, so that the supply of firm i depends on the expected price $p_t^e(i)$, and we assume that the cost structure is such that supply depends linearly on expected price. Here v_{1t} is a zero mean *iid* shock and, with suitable assumptions, p_t will be such that D_t is always positive.

Suppose that the cost of producing planned quantity $S_t^*(i)$, for each firm i, is given by $G_1 S_t^*(i) + \frac{1}{2}G_2 \left(S_t^*(i)\right)^2 - H'w_{t-1}S_t^*(i)$, and that there is also a zero mean *iid* shock v_{2t} that

[3] Closely related approaches are developed in Evans and Ramey (1992), Garcia-Schmidt and Woodford (2019) and Fahri and Werning (2019).

affects the actual quantity supplied, i.e. $S_t(i) = S_t^*(i) + v_{2t}$. Assuming firms are risk neutral and at time $t-1$ choose $S_t^*(i)$ to maximize expected profits, the quantity $S_t(i)$ supplied by firm i is

$$S_t(i) = Cp_t^e(i) + K + F'w_{t-1} + v_{2t}$$

where $C = G_2^{-1} > 0$, $K = -G_2^{-1}G_1$ and $F = G_2^{-1}H$. The exogenous observable vector w_t represents factors that affect marginal costs. It is assumed to follow a stationary stochastic process and is assumed to have zero mean and finite second moments. For example, w_t could follow a stationary VAR(1) process.

Let $p_t^e(i)$ denote the expectation of firm i of p_t, the market-clearing price at t. We assume that $p_t^e(i)$ is formed at the end of period $t-1$. This expectation can be interpreted as the mean of p_t, conditional on information available at $t-1$, computed using the subjective probability distribution of firm i. We also use the notation $E_{t-1}^* p_t(i)$ for $p_t^e(i)$, or, in the homogeneous expectations case, $E_{t-1}^* p_t = p_t^e$. In general one can allow for heterogeneous expectations across firms, with aggregate supply is given by $S_t = \int_0^1 S_t(i)di$. Assuming homogenous expectations case, period t price is determined by market clearing, i.e. $D_t = S_t$, which gives the equation

$$p_t = \mu + \alpha p_t^e + \delta' w_{t-1} + \eta_t, \tag{2}$$

where $\mu = B^{-1}(A-K)$, $\alpha = -B^{-1}C$, $\delta = -B^{-1}F$ and where $\eta_t = B^{-1}(v_{1t} - v_{2t})$ is *iid* exogenous with mean zero. Note that $\alpha < 0$ assuming demand and supply relations have the usual slopes. Equation (2) is the temporary equilibrium (TE) equation that determines market clearing price, given expectations and the exogenous shocks.

Under RE we replace $p_t^e = E_{t-1}^* p_t$ by $E_{t-1}p_t$, the true conditional distribution of p_t, and it is easily verified that there is a unique REE in which

$$E_{t-1}p_t = \bar{a} + \bar{b}' w_{t-1}, \text{ where } \bar{a} = (1-\alpha)^{-1}\mu \text{ and } \bar{b} = (1-\alpha)^{-1}\delta, \text{ and} \tag{3}$$

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

We can see, via the TE map (2), the sense in which RE must be viewed as an equilibrium concept. Indeed, suppose all agents $i \in [0,1]$ except for agent j had naive expectations, i.e. $p_t^e(i) = p_{t-1}$ for all $i \neq j$. Then $p_t^e = p_{t-1}$ so that

$$p_t = \mu + \alpha p_{t-1} + \delta' w_{t-1} + \eta_t,$$

from which we see that the rational forecast for agent j would be

$$p_t^e(j) = \mu + \alpha p_{t-1} + \delta' w_{t-1},$$

which is different from RE.

2.1 Behavioral Rules

The example at the end of the preceding section illustrates the flexibility of the temporary equilibrium approach, in that we can examine the implications of agents using forecasts that reflect plausible behavioral rules or rules of thumb. Economic "learning to forecast" experiments show that subjects often appear to use simple forecast rules. Examples include

$$\begin{aligned}
\text{Naive} &: \; p_t^e = p_{t-1} \\
\text{Trend-chasing} &: \; p_t^e = p_{t-1} + \vartheta \left(p_{t-1} - p_{t-2}\right), \text{ where } 0 < \vartheta < 1 \\
\text{Adaptive expectations} &: \; p_t^e = p_{t-1}^e + \lambda \left(p_{t-1} - p_{t-1}^e\right), \text{ where } 0 < \lambda < 1 \\
\text{Mean forecasts} &: \; p_t^e = N^{-1} \sum\nolimits_{j=1}^{N} p_{t-j}, \text{ for integer } N > 1.
\end{aligned}$$

For the apparent prevalence of these types of rules used by subjects in experiments see Hommes (2011) and Hommes (2013). In these experiments, serially correlated exogenous observables are typically not included. It should be noted that in some circumstances such behavioral rules can be fully optimal. For example, adaptive expectations (AE) is known to provide the optimal forecast, for appropriate λ, if p_t follows an IMA(1,1) process, i.e. if the first-difference of p_t is a first-order moving average process.

For any of the above forecast rules, we can solve for the implied TE path by inserting the expectation rule into the TE equation (2). Indeed, it is also straightforward computationally to obtain the implied TE path that arises when there are heterogeneous expectations and fixed proportions of agents that use, for example, each of the above four forecast rules.

From the AL viewpoint this falls short of our bounded rationality guidelines since each agent is using a fixed rule without trying to choose a well-performing forecast rule. However, some natural extensions of the behavioral approach clearly do satisfy the principle outlined above. If agents used AE but instead of using a fixed λ they attempted to use the best-performing value of λ, this would be an AL approach. Similarly what we call "mean forecasts" can be viewed as providing an optimal estimate of the population mean of the p_t. Finally if we allow agents to *choose* between the above four rules, and if their choices reflect relative forecast performance, then this more sophisticated version of forecasting also is in line with the the AL guiding principles. Some evidence of shifting proportions over time in the choice of rules can be found in Hommes (2013).

2.2 Least-squares Learning

The lead implementation of AL for our exposition here will be econometric learning, and more specifically least-squares (LS) learning. This implementation obeys the cognitive consistency principle: when economists need to forecast they usually proceed statistically, and the benchmark procedure is to use LS. Least-squares is a central forecasting procedure because of its optimality properties, e.g. the minimum mean-square-error linear forecast of a variable y given observed variables x is provided by the least-squares projection of y onto x. This linear projection can be estimated easily using the standard ordinary least squares regression of y onto x using the available data.

For the cobweb model under consideration here, we assume that agents (here, the firms) have a Perceived Law of Motion (PLM)

$$p_t = a + b' w_{t-1} + \eta_t \tag{4}$$

where a, b are unknown and η_t is a perceived white noise (exogenous *iid* zero mean) unobserved shock. Thus we are assuming that agents understand that specified observable, exogenous variables w_{t-1} impact equilibrium price in a linear way, but that agents do not know the RE values $\bar{a}, \bar{b}$. Under LS learning, the agents will estimate a, b using time-series data on p_t, w_{t-1}. Suppose that agents use their estimates to make forecasts and update these estimates over time as more data become available. Will their estimates converge over time to the REE values? To answer this we use the TE framework with expectations evolving in accordance with LS learning.

Start the system at time $t-1$ with initial data $\{p_s, w_{s-1}\}_{s=1}^{t-1}$ and assume that in $t-1$ agents first use these data to regress p on lagged w yielding estimates (a_{t-1}, b_{t-1}) of the unknown (a, b), according to

$$\begin{pmatrix} a_{t-1} \\ b_{t-1} \end{pmatrix} = \left(\sum_{s=1}^{t-1} z_{s-1} z'_{s-1} \right)^{-1} \left(\sum_{s=1}^{t-1} z_{s-1} p_s \right), \text{ where } z'_s = \begin{pmatrix} 1 & w'_s \end{pmatrix}. \tag{5}$$

The exogenous vector w_{t-1} is then realized, and time $t-1$ expectations of p_t are formed according to

$$p_t^e = a_{t-1} + b'_{t-1} w_{t-1}. \tag{6}$$

We then move to period t. Given expectations (6) and the white noise shock η_t, the temporary equilibrium price p_t is determined by (2). The data set then incorporates (p_t, w_{t-1}), the LS coefficients estimates (a_{t-1}, b_{t-1}) are updated to (a_t, b_t), given by

$$\begin{pmatrix} a_t \\ b_t \end{pmatrix} = \left(\sum_{s=1}^{t} z_{s-1} z'_{s-1} \right)^{-1} \left(\sum_{s=1}^{t} z_{s-1} p_s \right), \tag{7}$$

and the process continues. This recursion fully defines a path for a_t, b_t, p_t over time. The question of interest is whether

$$(a_t, b_t) \rightarrow (\bar{a}, \bar{b}) \text{ as } t \rightarrow \infty.$$

If so we say that the REE is stable under LS learning. The answer is given by the following Theorem, due to Bray and Savin (1986) and Marcet and Sargent (1989b).

Theorem 1 *Consider the dynamic system given by (2) and (6) with LS updating of (a_t, b_t) according to (5). If $\alpha < 1$ then $(a_t, b_t) \rightarrow (\bar{a}, \bar{b})$ as $t \rightarrow \infty$ with probability 1. If $\alpha > 1$ then (a_t, b_t) then convergence occurs with probability 0.*

Note that if supply and demand have their usual slopes then $\alpha < 0$ so that we always have convergence to RE with probability one. The case $\alpha < 0$ is often called the *negative expectational feedback* case. Some simple macro models take the form (2) but have *positive expectational feedback.* Convergence of LS learning to the REE still obtains if $0 < \alpha < 1$.

As an example with positive feedback, a simple ad hoc macro model combines an expectations augmented Phillips curve with a quantity theory aggregate demand equation and a money supply feedback rule is given by:

$$\begin{aligned} y_t &= \bar{y} + \kappa \left(p_t - p_t^e\right) + \zeta_t, \ \ \kappa > 0 \\ m_t + v_t &= p_t + y_t \\ m_t &= \bar{m} + u_t + \psi w_{t-1}, \end{aligned}$$

where w_t is an exogenous observable vector and ζ_t, v_t are exogenous white noise. When solved for p_t this takes the form (2) with $\alpha = (1+\kappa)^{-1}\kappa$, giving us the stable positive feedback case.

The positive result in Theorem 1 was proved from first principles by Bray and Savin (1986). Marcet and Sargent (1989b) provided a more general framework, based on stochastic approximation theorems, which also delivers the negative result. The latter theorems are discussed in detail in Evans and Honkapohja (2001). Although formal proofs based on these techniques use some sophisticated technical machinery, the stability condition can generally be obtained quickly using the *expectational stability* or *E-stability* principle, which we now describe.

2.3 E-stability

The E-stability approach is described in Evans and Honkapohja (2001) and has been applied in a wide range of economic models. For the model (2) this works as follows. Start with PLM (4) and suppose (a, b) are fixed at some (possibly non-REE) value. For this PLM the expectation is given by $E_{t-1}^* p_t = a + b' w_{t-1}$, which, when inserted into (2), leads to the associated data-generating process, or *Actual Law of Motion* (ALM),

$$p_t = \mu + \alpha(a + b' w_{t-1}) + \delta' w_{t-1} + \eta_t. \tag{8}$$

This gives the mapping T: PLM $\rightarrow$ ALM

$$T(a, b) = (\mu + \alpha a, \delta + \alpha b). \tag{9}$$

Note that the REE $\bar{a}, \bar{b}$ is a fixed point of T. E-stability is defined by the ordinary differential equation (ODE)

$$\frac{d}{d\tau}(a, b) = T(a, b) - (a, b), \tag{10}$$

where τ is notional time. The REE $\bar{a}, \bar{b}$ is *E-stable* if it is a stable fixed point of this ODE. For the model at hand T is linear and the REE is E-stable when $\alpha < 1$. Note that this E-stability condition is precisely the condition given in the above theorem.

The intuition of E-stability is that under LS learning the parameters a_t, b_t are slowly adjusted, on average, in the direction of the corresponding ALM parameters. The *E-stability principle* is that E-stability quite generally governs stability of an REE under LS and closely related adaptive learning rules. The E-stability technique can be used in multivariate linear models, nonlinear models, and if there are multiple equilibria.

2.4 Recursive LS and stochastic approximation

The formal link between E-stability and stability under LS learning starts with the recursive formulation of LS. Letting $\phi_t' = (a_t, b_t')$ denote the vector of parameter estimates and $z_{t-1}' = (1, w_{t-1}')$ the vector of regressors, the temporary equilibrium equation (8), using time $t-1$ parameter estimates to make forecasts $E_{t-1}^* p_t = a_{t-1} + b_{t-1}' w_{t-1}$, can be written

$$p_t = T\left(\phi_{t-1}\right)' z_{t-1} + \eta_t, \tag{11}$$

where T is given by (9). Under RLS updating the LS estimates (7) can be written as

$$\begin{aligned} \phi_t &= \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1}(p_t - \phi_{t-1}' z_{t-1}) && (12) \\ R_t &= R_{t-1} + t^{-1}(z_{t-1} z_{t-1}' - R_{t-1}), && (13) \end{aligned}$$

which is known as the recursive least-squares (RLS) system. Here R_t is an estimate of the second moment matrix of the regressors. The RLS system constitutes a *stochastic recursive algorithm* (SRA), as shown by Marcet and Sargent (1989b).[4] There are general methods for analyzing the dynamics of SRAs; see Ljung (1977), Marcet and Sargent (1989b), Benveniste, Metivier, and Priouret (1990) and Evans and Honkapohja (2001). These methods, often called the stochastic approximation techniques, approximate SRAs by an associated ODE. In the context of (12)-(13) the ODE takes the form

$$\begin{aligned} d\phi/d\tau &= R^{-1} M(T(\phi) - \phi), && (14) \\ dR/d\tau &= M - R, && (15) \end{aligned}$$

where M is the unconditional second-moment matrix of the exogenous regressors z_t. Here τ is often interpreted as notional time, but as noted below, τ can be linked to real time t. The steps for obtaining the ODE approximation (14)-(15) are given in the references just cited. For a compact summary see Evans and Honkapohja (2009b).

[4] SRAs take the general form $\theta_t = \theta_{t-1} + \gamma_t Q(t, \theta_{t-1}, X_t)$, where θ_t is a vector of parameter estimates, X_t is the state vector, and γ_t is a deterministic sequence of "gains." To put (12)-(13) in this form one substitutes $R_t \equiv \tilde{R}_{t-1}$ and sets $\theta_t' = \left(\phi_t, \text{vec}\left(\tilde{R}_t\right)\right)$. Here vec stacks the columns of a matrix into vector form. Substituting for p_t from (11) the SRA form is obtained for suitable X_t and $\gamma_t = t^{-1}$.

Inspecting the differential equation (15), it is seen that $\lim_{\tau\to\infty} R(\tau) = M$. Hence $R^{-1}M \to I$ and local stability under (14)-(15) of the fixed point $\bar{\phi}' = (\bar{a}, \bar{b})$ is determined by local stability under the "small" ODE

$$d\phi/d\tau = T(\phi) - \phi, \tag{16}$$

which is identical to the E-stability equation (10). Note that the unique REE for model (2) is the unique fixed point $\bar{\phi} = T(\bar{\phi})$ of the system (10). Stochastic approximation results can be used to show that local convergence with probability one obtains if $\alpha < 1$ whereas there is convergence with probability zero if $\alpha > 1$.[5] This confirms for the REE in the cobweb model that E-stability governs stability under LS learning.

A variation of LS learning that can be sometimes be both useful and simpler, is *generalized stochastic gradient learning.*[6] In this case R_t is replaced in (14) by some fixed matrix and thus there is no additional updating equation (15). If the regressors are exogenous, then it would natural to replace R_t by M, the second-moment matrix of the regressors, if is known. In that case the RLS updating equation is just $\phi_t = \phi_{t-1} + t^{-1}M^{-1}z_{t-1}(p_t - \phi'_{t-1}z_{t-1})$ and the associated ODE reduces to (16), i.e. the E-stability equation itself.

2.5 The E-stability principle

The general definition of E-stability in economic models follows the same steps used in the cobweb model. Given an economic model, consider first an REE, identified as a stochastic process with parameter vector $\bar{\phi}$. Under adaptive learning (typically implemented by LS or a close variant), ϕ is estimated by economic agents and the estimates ϕ_t, which are updated over time, are used to form expectations and make decisions. An REE $\bar{\phi}$ is said to be E-stable if it is locally asymptotically stable under the differential equation (16).[7] To assess E-stability one simply determines the mapping from the PLM parameters ϕ to the ALM parameters $T(\phi)$ that describes the actual stochastic process when agents form expectations using ϕ. E-stability is then determined by checking whether the eigenvalues of $DT(\bar{\phi}) - I$ have negative real parts (or equivalently that eigenvalues of $DT(\bar{\phi})$ having real parts less than one). This, of course, provides the conditions for local stability of $\bar{\phi}$ under the E-stability ODE, and corresponds to local stability under LS learning. In some cases stability of $\bar{\phi}$ under the E-stability ODE is global, in which case we say that $\bar{\phi}$ is globally E-stable for the PLM being considered.

E-stability, which in many cases is a straightforward calculation, quite generally provides the correct condition for stability under LS learning. In many cases stochastic approximation

[5] In general, various technical assumptions need to be satisfied for application of the stochastic approximation theorems, and these are satisfied for the cobweb model.

[6] See Evans, Honkapohja, and Williams (2010) for an anaysis of generalised stochastic gradient learning.

[7] This definition of E-stability was introduced in Evans (1989) and Evans and Honkapohja (1992). Marcet and Sargent (1989b) emphasized the importance of SRAs to study LS learning.

results can be used to demonstrate that stability of an REE under LS learning is governed by E-stability. Even in cases in which for technical reasons the stochastic approximation results cannot be applied, numerical simulations indicate its validity.

A technical detail that can be important in practice arises especially when there are multiple REE associated with a given PLM so that E-stability of an REE is local but not global. In this case the claim that under LS learning E-stability implies convergence with probability one usually requires that the updating rule be augmented by a "projection facility" that prevents estimates from leaving a suitably defined compact set. Weaker convergence results (convergence probabilities near one or positive probability of convergence) can dispense with the projection facility.

It is important to note that E-stability depends on the PLM, i.e. on the specification of the forecasting model used by agents. If there is more than one natural representation of the REE, i.e. more than one PLM consistent with the REE, then in principle the E-stability conditions can depend on the PLM used by agents. For example, if agents overparameterize the dynamics of an REE this can lead to stricter E-stability conditions.[8]

Importantly, the E-stability Principle extends to cases in which agents estimate misspecified models, arising from PLMs that do not nest an REE because they omit a relevant variable, underparameterize the dynamics or misspecify the functional form. Because in practice econometric forecasters recognize that the forecast models they use are, at least to some degree, misspecified, we turn to this case now, in which the relevant solution concept is a restricted perceptions equilibrium (RPE).[9]

2.6 Restricted perceptions equilibria

If the PLM of the agents is linear but misspecified, the E-stability principle extends in a natural way to assess the stability under LS learning of an RPE. The RPE itself can be calculated using a set of orthogonality conditions.

Continuing with the cobweb model (2), suppose w_t is a 2×1 covariance stationary exogenous vector $w_t' = (w_{1t}, w_{2t})$, where for convenience we assume $Ew_t = 0$, so that

$$p_t = \mu + \alpha E_{t-1}^* p_t + \gamma_1 w_{1,t-1} + \gamma_2 w_{2,t-1} + \eta_t. \tag{17}$$

To capture misspecification, we assume agents condition their forecasts only on w_1, so that their PLM is given by

$$\text{PLM: } p_t = a + c w_{1,t-1} + \varepsilon_t, \tag{18}$$

where the perceived disturbance ε_t is treated by agents as unpredictable white noise. Then

[8] For these casess one can distinguish between weak and strong E-stability conditions, where the latter refers to E-stability with repsect to the overparameterized specification. See Evans and Honkapohja (2001).

[9] For a general discussion of RPEs see Branch (2006).

$E^*_{t-1}p_t = a + cw_{1,t-1}$ and the corresponding ALM is

$$\text{ALM: } p_t = \mu + \alpha a + (\gamma_1 + \alpha c)\, w_{1,t-1} + \gamma_2 w_{2,t-1} + \eta_t. \tag{19}$$

Because the PLM is misspecified the ALM law of motion does not lie in the space of PLMs considered. The RPE coefficients (a, c) that minimize the mean square error (MSE), and the associated E-stability conditions, are obtained from a T-map based on the *projected ALM,*

$$\text{Projected ALM: } p_t = T_a + T_c w_{1,t-1} + \varepsilon_t,$$

which projects the ALM process for p_t (19) onto the variables $(1, w_{1,t-1})$.

The coefficients (T_a, T_c) are given by the least-squares *orthogonality conditions* that the forecast error $p_t - T_a - T_c w_{1,t-1}$ must be uncorrelated with both regressors $(1, w_{1,t-1})$. This leads to the conditions

$$\begin{aligned} E\left(p_t - T^a - T^c w_{1,t-1}\right) &= 0 \\ E\left((p_t - T^a - T^c w_{1,t-1}) w_{1,t-1}\right) &= 0, \end{aligned}$$

where p_t is given by (19). Using $Ew_{1,t} = Ew_{2,t} = E\eta_t = 0$ the two conditions are given by the *projected T-map*

$$T^a = \mu + \alpha a \text{ and } T^c = \gamma_1 + \frac{\omega_{12}}{\omega_{11}}\gamma_2 + \alpha c,$$

where $\omega_{11} = \text{var}\,(w_{1t})$ and $\omega_{12} = \text{cov}\,(w_{1t}, w_{2t})$.

The fixed point of the map $T\,(a, c) = (T_a, T_c)$ is given by

$$\bar{a} = (1-\alpha)^{-1}\mu \text{ and } \bar{c} = (1-\alpha)^{-1}\left(\gamma_1 + \omega_{11}^{-1}\omega_{12}\gamma_2\right).$$

The corresponding E-stability ODE is

$$\frac{d}{d\tau}\begin{pmatrix} a \\ c \end{pmatrix} = \begin{pmatrix} \mu + \alpha a \\ \gamma_1 + \frac{\omega_{12}}{\omega_{11}}\gamma_2 + \alpha c \end{pmatrix} - \begin{pmatrix} a \\ c \end{pmatrix},$$

and it is easily verified that the E-stability condition is once again just $\alpha < 1$. The forecasts $E^*_{t-1}p_t = \bar{a} + \bar{c}w_{1,t-1}$ define the RPE expectations, and substituting this expression into (17) gives the RPE solution

$$p_t = \bar{a} + \bar{c}w_{1,t-1} + \gamma_2 w_{2,t-1} + \eta_t.$$

As a second example of an RPE suppose that the cobweb model is nonlinear. More specifically, suppose that demand is given by $D\,(p_t, v_t)$ where $D_p < 0$, i.e. the demand curve slopes downward, and v_t is an *iid* shock independent of all supply shocks; and assume that supply for each firm is given by $S\left(E^*_{t-1}p_t, w_{t-1}\right)$, where the supply curve depend positively on $E^*_{t-1}p_t$ and w_{t-1} is a vector of covariance stationary VAR(1) cost shocks. Assume homogeneous expectations across firms. The market equilibrium condition that demand equals total firm supply implies

$$p_t = F\left(E^*_{t-1}p_t, w_{t-1}, v_t\right)$$

for a suitable nonlinear F. In this example we assume that when making forecasts firms do make use of the entire vector w_{t-1}, but we assume that they do so using an estimated linear PLM $p_t = a + b' w_{t-1}$. Under AL their estimates are updated over time using (recursive) LS. For this PLM, the corresponding ALM is given by

$$p_t = F\left(a + b' w_{t-1}, w_{t-1}, v_t\right) \equiv \hat{F}\left(a, b, w_{t-1}, v_t\right).$$

Again, the RPE and its stability properties under learning are obtained using the projected PLM and the E-stability principle. The projected T-map is given by $T(a,b) = \left(T^a, T^b\right)(a,b)$, where

$$T^a(a,b) = E_{w,v}\hat{F}(a,b,w,v) \text{ and } T^b(a,b) = \Sigma_w^{-1} E_{w,v}\left(w\hat{F}(a,b,w,v)\right)$$

where Σ_w is the covariance matrix of the supply shocks w and $E_{w,v}$ is the expectation taken over the stationary distributions of w, v. The fixed point $\left(\bar{a}, \bar{b}\right)$ of T identifies the RPE associated with the linear forecast model, and it is E-stable if it is locally asymptotically stable under the ODE (10). Evans and McGough (2019) show that for normally sloped supply and demand curves then, at least for supply shocks with sufficiently small bounded support, there is a unique RPE and it is stable under adaptive learning.

There are many examples and applications of RPE in the literature, e.g. see Marcet and Sargent (1989a), Evans and Honkapohja (2001), Ch. 3 and 11, Evans and Ramey (2006), Branch and Evans (2006a), Adam (2007), Guse (2008), Slobodyan and Wouters (2012a) and Hommes and Zhu (2014). In specific settings RPE are often given specific names, e.g. Evans and McGough (2020b) show existence of "near-rational sunspot equilibria" in nonlinear models when agents have forecasts that depend linearly on continuously-measured extraneous "sunspot" variables.

2.7 Constant-gain learning

Under least-squares learning, all past data points count equally, and so when forecasts are updated each period the most recent data point has weight t^{-1}. This is reflected in the t^{-1} term in the RLS equations (12)-(13). An alternative that is natural if there is concern that structural change may be occurring over time, in some unmodelled way, is to weight recent data points more heavily than past data points, e.g. to downweight past data points geometrically. This is accomplished by replacing t^{-1} by γ for fixed $0 < \gamma < 1$, where γ represents the weight on the most recent data point; this in effect gives past data points from time $t - i$ weight $(1-\gamma)^i$. More generally the term t^{-1} can be replaced by a sequence $\gamma_t > 0$, called the gain sequence, with the main cases being the standard decreasing gain sequence $\gamma_t = t^{-1}$ and the constant gain sequence $\gamma_t = \gamma$.

Typical constant gains for quarterly data in macro models with one-period ahead forecasts are $\gamma \in [0.01, 0.05]$, while for models in which agents make long-horizon forecasts the gains used are smaller. See, for example, Branch and Evans (2006b) and Eusepi and Preston (2011). From the agents' point of view the optimal gain to use depends on the extent of the actual or perceived structural change: a large γ will more quickly track resulting changes in optimal forecast parameters, while a small gain is more effective at filtering out noise.

In practice, quantitative and estimated empirical models with adaptive learning typically assume a constant gain. This appears to fit the data well and provides a continuing role for adaptive learning. An advantage of this approach is that one can view the economy as a stationary stochastic process, rather than one exhibiting transitional dynamics. Under constant gain learning, agents estimates are always being revised and hence there is "perpetual learning." Even if the PLM nests the REE, instead of convergence to the RE parameters, estimates will converge to a stochastic process. For small constant gains γ, the estimates will typically converge to a stochastic process centered at the RE parameters when the RE is E-stable.

Analytical results are available for the case of small constant gain. These results are most easily stated in the case in terms of the constant-gain version $\theta_t = \theta_{t-1} + \gamma Q(t, \theta_{t-1}, X_t)$ of the general SRA mentioned in a footnote in Section 2.4. Here for $\theta_t' = \left(\phi_t, \text{vec}\left(\tilde{R}_t\right)\right)$, where $\tilde{R}_{t-1} \equiv R_t$, the associated ODE takes the form

$$d\theta/d\tau = h\left(\theta\left(\tau\right)\right), \text{ where } h_\phi(\phi) = \tilde{R}^{-1}M\left(T\left(\phi\right) - \phi\right) \text{ and } h_{\tilde{R}}(\tilde{R}) = M - \tilde{R}. \qquad (20)$$

The ODE $d\theta/d\tau = h\left(\theta\left(\tau\right)\right)$ give the *mean dynamics* of the discrete-time SRA.

Informally, making the identification $\tau = \gamma t$,[10] it can be shown that, starting from given initial parameter estimates θ_0 near the RE value $\bar{\theta}$, for constant gains γ sufficiently small the unconditional expected value of θ_t can be approximated, over a given time range $0 \leq \tau \leq \mathcal{L}$, by

$$E\theta_t \approx \tilde{\theta}\left(\gamma t, \theta_0\right),$$

where $\tilde{\theta}\left(\tau, a\right)$ is the solution to $d\theta/d\tau = h\left(\theta\left(\tau\right)\right)$ over the same time range. Furthermore, under additional technical assumptions that ensure that a well-defined stationary distribution is reached for large t, it can be shown that for γ sufficiently small and γt sufficiently large $\theta_t \overset{a}{\sim} N(\bar{\theta}, \gamma C)$ for a suitable matrix C. Here $\overset{a}{\sim}$ means "approximately distributed as." See Evans and Honkapohja (2001) for the method for computing C.

The mean dynamics given by (20) can be useful for showing the global dynamics of the PLM coefficient estimates, from a variety of initial expectations, as well as for verifying local stability under AL. However, as emphasized in Sargent (1999), even when starting near the REE, particular sequences of random shocks can lead the economy to follow "escape" paths leading to big deviations from the REE steady state for extended periods of time. This possibility of recurrent large deviations, known as "escape dynamics," was documented and studied in Cho, Williams, and Sargent (2002). Using the theory of large deviations, Williams

[10] This holds also for the decreasing gain case for appropriate assumptions on the gain sequence, including $\gamma_t = t^{-1}$, where instead $\tau = \sum_{i=1}^{t} \gamma_i$.

(2019) provides tools for characterizing the frequency and direction of escape dynamics. Applications in which escape dynamics play a prominent role include Kasa (2004), McGough (2006), Cho and Kasa (2008), Branch and Evans (2011) and Branch and Evans (2017).

As noted in the introduction to this section, a motivation for the use of constant-gain learning is the presence of structural change: in non-stationary environments agents have incentive to view apparent outlier data as possibly indicating a shift in the economic environment. By placing more weight on more recent data, *constant-gain learning* (CGL) provides a natural model of this potential alertness.

To illustrate constant-gain learning in the presence of structural change, we modify the cobweb model to include a proportional sales tax τ: whence if p_t^e is the expected market price then expected marginal (net) revenue is $(1-\tau)p_t^e$, yielding the TE equation

$$p_t = \mu + \alpha(1-\tau)p_t^e + \delta' w_{t-1} + \eta_t.$$

Figure 1 illustrates learning dynamics for an unanticipated, permanent (and admittedly large) tax increase of 50% in period 500. The upper panels provide the dynamics of beliefs (a_t, b_t), which were initialized in REE. The lower-left panel shows the price dynamics over the whole simulation and the lower-right panel shows the price dynamics (black) in a 100-period window centered at the time of the tax hike; also included in this panel, for comparison purposes, are the price dynamics under REE (red). All horizontal, dashed lines correspond to REE means.

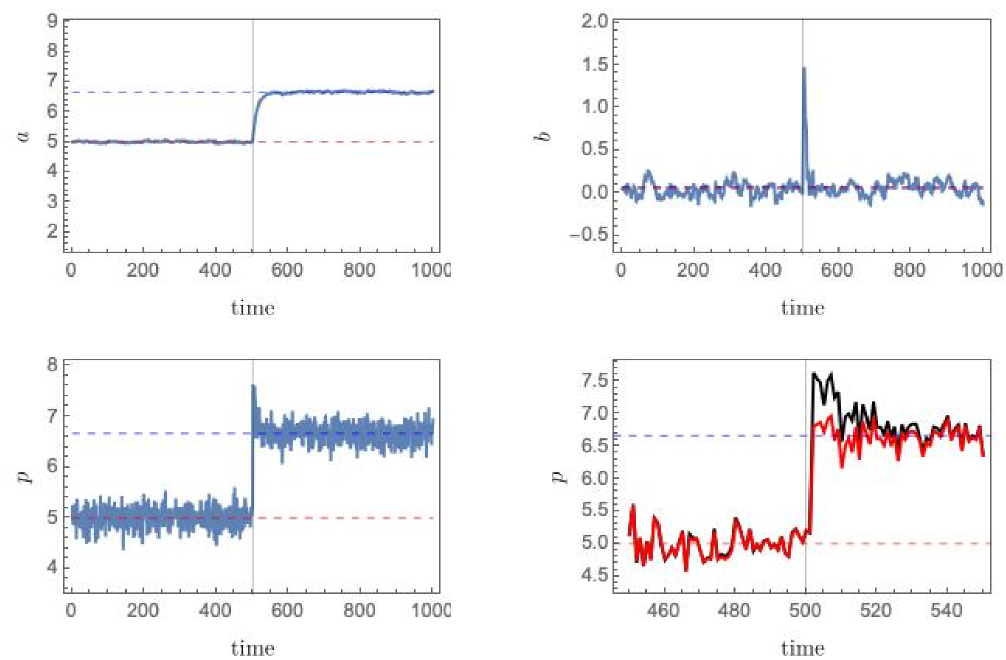

Figure 1: Cobweb model with tax increase.

We see that, prior to the shock, beliefs stay near the REE values, and, as indicted in the lower-right panel, the learning and REE price paths are very nearly identical. The tax hike in period 500 causes a sharp decrease in supply, resulting in a price spike under both learning and RE. The REE price path simply stays aligned with the new REE price process; the learning path, however, is more complex. The CGL algorithm initially attributes the price spike to an increased responsiveness of prices to w, as seen in the upper right-panel. This misperception causes further curtailment of supply, resulting in an over-shooting of the price path. Newly generated data quickly disabuse agents of their erroneous beliefs, as the constant term rises to its new steady-state level – see upper-left panel – and the learning price path returns to match its REE analog.

Even in stationary environments CGL can provide interesting dynamics over-and-above those observed under rational expectations. For example, the heavier weight placed on newly observed data can induce excess volatility, particular when the model's expectational feedback is strong. As an illustration, we again use the cobweb model. When $\delta = 0$ and the PLM is simply given by a constant, the excess volatility, defined as the unconditional variance of price under learning relative to its variance under REE, can be derived analytically, and in doing so it can be shown that the learning dynamics are stationary provided $\alpha \in (1 - 2/\gamma, 1)$.

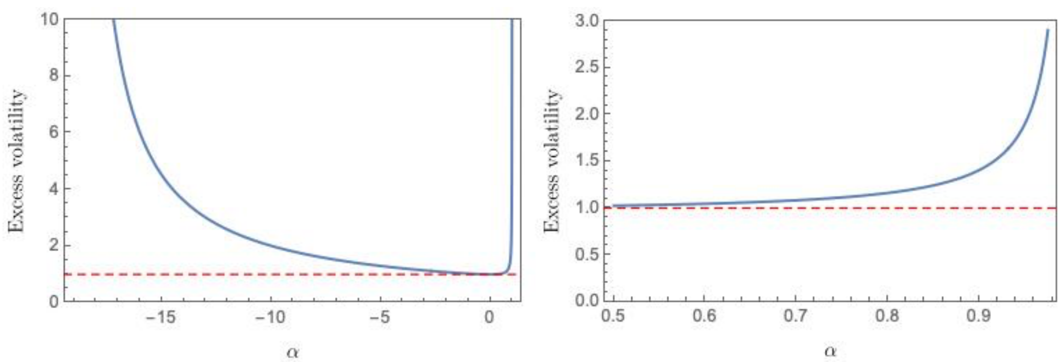

Figure 2: Excess volatility under constant gain learning

Figure 2 provides the relevant plot, allowing the expectational feedback parameter α to vary: the left panel provides the full range of α and the right panel provides a closer look for $\alpha \in (0.5, 1)$.[11] Notice that for values of α near $1 - 2/\gamma$ or one the excess volatility can be arbitrarily large. Via simulations, we have found that the same pattern emerges for examined calibrations with $\delta > 0$ and in which agents' PLM includes the shock w as a regressor.

In simple asset-pricing models, with strong positive expectational feedback, constant gain learning can lead to substantial excess volatility of asset prices, in line with the above theoretical results and the well-known empirical results of Shiller (1981). See, for example, Bullard, Evans, and Honkapohja (2009). Constant gain learning also may be able to explain

[11] Here the gain is chosen as $\gamma = 0.1$.

some well-known empirical puzzles in other areas of asset-pricing, including foreign exchange rates and the yield curve: see Chakraborty and Evans (2008) and Sinha (2016).

2.8 Heterogeneous expectations and multiple forecasting models

Although in our examples we have focused on representative-agent models in which agents have homogeneous expectations, the AL approach can be extended to incorporate heterogeneous expectations in a variety of ways. Evans, Honkapohja, and Marimon (2001) show how it is possible to allow for random inertia and heterogeneous gains across agents in updating expectations.

Another way to incorporate heterogeneous expectations is to assume that several distinct forecasting models are available and to use discrete-choice (or "dynamic predictor selection") models to determine the proportion of agents using each type of model based on a measure of fitness (e.g. estimated mean-square error) and an intensity of choice parameter. This approach was developed in Brock and Hommes (1997) and is a major focus of Hommes (2013), in which the set of forecast rules available are usually taken to be behavioral rules of the type described in Section 2.1. For an application of the discrete-choice approach to survey data see Branch (2004).

A distinct but related approach uses replicator dynamics to model the evolution of the proportions of a set of forecast rules over time. For examples see Guse (2010) and Branch and McGough (2008). Another AL approach is to allow for two types of agents, in which one group of agents draws on the forecasts of other agents. See Granato, Guse, and Wong (2008).

Genetic-algorithm learning (or "social learning"), first applied in economics by Arifovic (1994), allows by design for considerable heterogeneity of expectations and decisions. Applications include Arifovic (1996), who studies foreign exchange-rate dynamics, and the Arifovic, Bullard, and Duffy (1997) examination of economic growth and development. For a recent application to monetary economics see Arifovic, Bullard, and Kostyshyna (2013).

Branch and Evans (2006a) extend the dynamic predictor selection methodology of Brock and Hommes (1997) with the LS-learning AL approach. Agents are assumed to choose between two or more forecasting models that are underparameterized, e.g. in their excluding different relevant variables. Coefficient parameters of each model are updated over time in accordance with LS-learning, and the proportions of agents using each model are based on dynamic predictor selection and a model of fitness. Branch and Evans (2006a) show existence of a "misspecification equilibrium," in which agents forecast optimally given their choices, with forecast model parameters and predictor proportions determined endogenously. When there is negative expectational feedback there can be intrinsic heterogeneity, in which multiple forecast rules across agents are employed. When there is positive feedback, as in Branch and Evans (2007), there can exist multiple stable misspecification equilibria, and pooling and regime-switching may occur endogenously under learning.

When multiple forecasting models are on the table, another approach is to assume that agents use econometric methods to choose between them. In the context of foreign exchange rates, Markiewicz (2012) develops a model learning approach, based on the BIC model selection criterion, and shows that it can explain a change in the GBP/USD exchange rate volatility. A general analysis of the model validation approach is given in Cho and Kasa (2015).

When a discrete number of alternative forecasting rules are under consideration, an alternative approach is that individual agents may combine them, e.g. using Bayesian model-averaging. This approach is explored in Gibbs (2017), where among other results it is shown that multiple equilibria can arise. Bayesian model averaging is used in Evans, Honkapohja, Sargent, and Williams (2013) in the context of the cobweb model. Two models are under consideration by agents: one nests the REE, whereas the other model incorrectly assumes that a key parameter is subject to random-walk drift. They show that in the positive feedback case it is possible for Bayesian model averaging to lead agents select the misspecified model. Using large-deviation tools Cho and Kasa (2017) develop this possibility further in the context of a simple asset-pricing model.

3 Multivariate Linear Models

While the cobweb model (2) was convenient for introducing adaptive learning concepts and tools, macroeconomic models usually have more complex dynamic structures. Typically these include both forward-looking and backward-looking endogenous variables. For example, the risk-neutral asset-pricing model takes the form

$$p_t = \beta E_t^* p_{t+1} + d_t,$$

where p_t is the price of an equity and d_t is the dividend, assumed following an exogenous stochastic process. The same equation form arises for simple PPP models of the foreign exchange rate and for special cases of the overlapping generations model of money. In many models costs of adjustment are also important. For example, Sargent's linear-quadratic version of the Lucas-Prescott model of investment takes the form

$$k_t = \alpha + \beta E_t^* k_{t+1} + \delta k_{t-1} + \gamma u_t,$$

where k_t is the industry capital stock and u_t is exogenous. In this case, as is typically true of serious multivariate macroeconomic models, the key endogenous variables are both forward and backward looking.

An important new feature that can arise in the context of forward-looking linear models is the possibility of *indeterminacy*, i.e. the existence of multiple stationary REE. A linear model is said to be *determinate* if there is a unique nonexplosive RE solution.

Most dynamic stochastic general equilibrium (DSGE) macro models are, in addition, nonlinear, but they are frequently analyzed under RE using a linearization around a steady state, and typically take the standard form

$$y_t = ME_t y_{t+1} + Ny_{t-1} + Pv_t, \quad (21)$$
$$v_t = Fv_{t-1} + \varepsilon_t, \quad (22)$$

where y_t is a vector of endogenous variables, in deviation from mean form, and v_t is an exogenous, observable, stationary VAR(1) process driven by the white noise shock ε_t. The usual *minimal state variable* (MSV) RE solution takes the form

$$y_t = \bar{a} + \bar{b} y_{t-1} + \bar{c} v_t, \quad (23)$$

with here $\bar{a} = 0$.

3.1 E-stability in multivariate linear models

From the adaptive learning perspective, the ideal would be to build up every economic model from the agent level, with explicit assumptions about how agents formulate their decisions rules and make forecasts in nonlinear infinite or long-horizon settings. A convenient short-cut that can often be very informative, is to simply start from the linearized RE model and replace $E_t y_{t+1}$ by one-step ahead forecasts $E_t^* y_{t+1}$. This is the reduced-form learning approach. Under adaptive learning agents believe that $y_t = a + by_{t-1} + cv_t$ and use LS learning to update over time their estimates (a_t, b_t, c_t) of the parameters.

The E-stability principle can be used to evaluate stability of the RE solutions to (21)-(22) under learning. The analysis for an MSV solution proceeds as follows. For the PLM

$$y_t = a + by_{t-1} + cv_t$$

the corresponding forecasts are $E_t^* y_{t+1} = (I+b)a + b^2 y_{t-1} + (bc + cF)v_t$. Here, for convenience and without loss of generality, we assume that F is known. Inserting the forecasts into the model yields the ALM

$$y_t = M(I+b)a + (Mb^2 + N)y_{t-1} + (Mbc + NcF + P)v_t,$$

which gives the mapping from PLM to ALM:

$$T(a, b, c) = (M(I+b)a, Mb^2 + N, Mbc + NcF + P).$$

The REE $(\bar{a}, \bar{b}, \bar{c})$ is a fixed point of $T(a, b, c)$. If the E-stability ODE

$$d/d\tau(a, b, c) = T(a, b, c) - (a, b, c)$$

is (locally asymptotically) stable at the REE it is said to be E-stable. Evans and Honkapohja (2001), Chapter 10, show that the corresponding E-stability conditions can be stated in terms of the derivatives

$$DT_a = M(I + \bar{b}),\ DT_b = \bar{b}' \otimes M + I \otimes M\bar{b},\ DT_c = F' \otimes M + I \otimes M\bar{b}, \tag{24}$$

where $\otimes$ denotes the Kronecker product. If all eigenvalues of these matrices have real parts less than one then the REE is E-stable.

For a given numerical values of M, N, P, F it is straightforward to compute the MSV solutions, to determine whether the model is determinate or indeterminate, and to use (24) to assess each MSV solution for local stability under learning. In an indeterminate model with multiple stationary MSV solutions, E-stability can be used as a selection criterion. For some applications general analytical results can be obtained for determinacy and E-stability. The ease with which E-stability conditions can be assessed numerically is an advantage of the reduced-form short-cut to assessing stability under adaptive learning in linearized models.

3.2 Example: monetary policy

An early application of E-stability to New Keynesian (NK) models, using the reduced-form adaptive learning approach, was the assessment by Bullard and Mitra (2002) of alternative interest-rate rules by monetary policymakers. The basic linearized NK model (see e.g. Clarida, Gali, and Gertler (1999)) takes the form

$$x_t = -\varphi\left(i_t - E_t^* \pi_{t+1}\right) + E_t^* x_{t+1} + v_{xt} \tag{25}$$

$$\pi_t = \beta E_t^* \pi_{t+1} + \lambda x_t + v_{\pi t}, \tag{26}$$

where $0 < \beta < 1$ is the discount factor and $\varphi, \lambda > 0$. Here x_t is the output gap, π_t is inflation and i_t is the nominal interest rate, all of which are expressed in deviation from their steady-state values.[12] These equations are often called the NK IS and NK PC equations. The exogenous IS and inflation shocks $v_{xt}, v_{\pi t}$ are typically assumed to be independent, stationary AR(1) processes. These model equations must be supplemented by a monetary policy rule. Typically this takes the form of a “Taylor” rule, which sets the interest rate in response to either contemporaneous, lagged or expected inflation and output. Since monetary policy is often treated as forward-looking a simple such rule takes the form

$$i_t = \alpha_\pi E_t^* \pi_{t+1} + \alpha_x E_t^* x_{t+1}, \tag{27}$$

where $\alpha_\pi, \alpha_x > 0$. Sometimes an exogenous shock is also included in the interest-rate setting equation. Alternative rules are current-data rules

$$i_t = \alpha_\pi \pi_t + \alpha_x x_t, \tag{28}$$

[12] For this to be a suitable approximation the (steady state) target inflation rate for the monetary authorities needs to be close to zero.

and backward-looking rules,

$$i_t = \alpha_\pi \pi_{t-1} + \alpha_x x_{t-1}. \tag{29}$$

The current-data rule (28) has been criticized on the grounds that, first, policymakers in fact respond mainly to the expected state of the economy in the near future, and, second, accurate information on current output and inflation is only available with a lag.

Bullard and Mitra (2002) assess the NK model for determinacy, i.e. whether there is a unique non-explosive RE solution, and for stability under learning, which is examined using E-stability for the three alternative policy rules and for different choices of α_π, α_x. Determinacy can be assessed by putting the system in standard first-order matrix form and comparing the number of eigenvalues outside the unit circle to the number of "free" (current endogenous) variables.[13] The E-stability conditions in Section 3.1 can be used to assess stability of an REE under adaptive learning. These are distinct conditions. For the NK model (25)-(26), with either policy rule (27) or (28), inserting the policy equation leads to a bivariate forward-looking model[14] of the form

$$y_t = M E_t y_{t+1} + P v_t, \tag{30}$$

where $y_t' = (x_t, \pi_t)$ and $v_t' = (v_{xt}, v_{\pi t})$, and where M and P are functions of the structural parameters. In this case the model is determinate if both roots of M lie inside the unit circle, while if one or both roots lie outside the unit circle the model is indeterminate. When the model is determinate the unique stationary solution takes the form

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

for suitable 2×2 matrix $\bar{c}$. When the model is indeterminate there continues to be an MSV solution of this form, but there are also other stationary solutions.

E-stability of the MSV solution to (30) is also straightforward to compute. Because $\bar{b} = 0$ in the MSV solution in the current cases, the PLM estimated by agents takes the form

$$y_t = a + c v_t.$$

Note that we include intercepts in the PLM because agents will generally need to estimate the means of the x_t and π_t as well as the impact of v_t. The E-stability conditions reduce to checking that all eigenvalues of

$$DT_a = M \text{ and } DT_c = F' \otimes M$$

have real parts less than one.

[13] If these are equal the model is determinate. If there are fewer eigenvalues outside the unit circle, the model is indeterminate: there are multiple stationary REE. For details see, e.g., Blanchard and Kahn (1980) and Sims (2001).

[14] We are assuming this is the complete model. A debt-dynamics equation would need to be included if there is "active" fiscal policy.

Bullard and Mitra show that for the rule (28), both determinacy and E-stability hold provided $\lambda(\alpha_\pi - 1) + (1-\beta)\alpha_x > 0$. In particular, if the "Taylor principle" $\alpha_\pi > 1$ is satisfied then the model is determinate and the MSV solution is stable under learning. However, matters are more subtle in the case of forward-looking and backward-looking rules, and Bullard and Mitra (2002) computed the results for standard calibrations of φ, β, λ.

For the policy rule (27) we have that for $\alpha_\pi > 1$, but not too large, and with $\alpha_x > 0$ small, both determinacy and E-stability hold. However, there is a large region with $\alpha_\pi > 1$ and $\alpha_x > 0$ relatively large in which the model is indeterminate and the MSV solution is stable under learning. Of course in the indeterminate case other stationary REE exist as solutions. Subsequent research by Honkapohja and Mitra (2004) and Evans and McGough (2005b) established that in this case there exist "sunspot solutions," i.e. solutions that in addition depend on extraneous variables, which are stable also stable under learning. Below in Section 4 we discuss sunspot equilibria more generally.

Finally, for the policy rule (29) the MSV solution instead takes the form (23). Bullard and Mitra found that for $\alpha_\pi > 1$ and $\alpha_x > 0$ small the MSV solution is determinate and E-stable. However for $\alpha_\pi > 1$ and $\alpha_x > 0$ large the REE solutions are explosive and for $\alpha_\pi < 1$ and $\alpha_x > 0$ for an intermediate range of values, the model is determinate, but the stationary MSV solution is not E-stable. This last result is an illustration of the fact that, while in a wide range of models determinacy implies stability under learning, this is not a fully general result.[15]

There is an extensive literature on monetary policy and adaptive learning within the linear framework (25)-(26) and in hybrid versions that allow for habit formation and indexation that bring in lagged endogenous variables. This literature includes alternative implementations of optimal policy, simultaneous estimation of structural parameters by monetary policymakers, policymaker uncertainty about structural parameters, implications of constant-gain learning and/or misspecified PLMs, policy allowing for expectational heterogeneity, and the relative desirability of price level vs. inflation targeting.[16] See the survey papers Evans and Honkapohja (2003a), Evans and Honkapohja (2009a) and Evans and Honkapohja (2013). For a recent and comprehensive survey of monetary policy under learning that emphasizes the long-horizon "anticipated utility" approach, see Eusepi and Preston (2018).

In the aftermath of the Great Recession, a major policy issue in the US, beginning around 2014, was normalization of monetary policy: how quickly should interest rates be returned to normal levels from the near-zero rates that extended from January 2009 through December 2015? A neo-Fisherian view, advanced for example by Cochrane (2015), argued for

[15] For further discussion of this issue see McCallum (2007) and Bullard and Eusepi (2014).

[16] The possibility that some implementations of optimal policy may not be stable under AL was emphasized by Evans and Honkapohja (2003c) and Evans and Honkapohja (2006). For optimial policy rules that are stable under AL and allow for structural parameter uncertainty see Evans and McGough (2007). McGough, Rudebusch, and Williams (2005) examine monetary policy conducted using long rates. Honkapohja and Mitra (2019) study price-level targeting when when agents update their assessment over time of the credibility of the policy.

immediately increasing the policy rate to the level consistent with the targeted steady-state inflation rate and fixing the policy rate at that level. Rational expectations, and uniqueness of the steady state under an interest-rate peg, would appear to guarantee that inflation would then return to its targeted level. Evans and McGough (2018b) and Evans and McGough (2018a) argued that under AL such a policy would lead to unstable paths and which could, in particular, lead to renewed recession. Complementary arguments are provided by Garcia-Schmidt and Woodford (2019). The argument that fixed interest-rate pegs have undesirable properties under learning can be traced back to Howitt (1992).

4 Multiple Equilibria: Sunspots

A series of papers including Shell (1977), Azariadis (1981), Cass and Shell (1983), Grandmont (1985) and Azariadis and Guesnerie (1986) established that in simple nonstochastic nonlinear overlapping generations models of money there can exist multiple REE taking the form of stationary sunspot equilibria (SSEs) or regular periodic cycles. Here the term "sunspot" refers to an exogenous stochastic process that is "extrinsic" in the sense that is unconnected to fundamental shocks (the latter, also called "intrinsic" shocks, include taste, productivity, and other relevant random shocks). The SSEs shown to exist typically took the form of 2-state Markov processes.

Nonstochastic nonlinear models can also exhibit multiple steady states, and these can have SSEs taking values near different steady states. More recently RBC-type models with nonconvexities arising from externalities, increasing returns and monopolistic competition, distortionary taxes, etc., have been developed that possess an *indeterminate* steady state, in which there exist SSEs local to the indeterminate steady state. As already noted the indeterminate case can arise in New Keynesian (NK) models for certain specification of the interest-rate rule, which implies the existence of SSEs. Finally, a recent strand of the asset-price bubbles literature emphasizes cases in which there are nonexplosive asset-bubbles, and these may be viewed as SSEs.

An advantage of the AL approach is that it provides a way of assessing the plausibility of SSEs: if agents have PLMs that condition on an observed sunspot process, will the SSE be (at least locally) stable under leaning? The seminal paper by Woodford (1990) showed that SSEs in an overlapping-generations model could indeed be stable under AL. The work of Evans (1989), Evans and Honkapohja (1994), Evans and Honkapohja (2003b) and others showed that SSEs in general may or may not be stable under learning, and that stability can be readily assessed using the E-stability principle.

While much of the early work focused on simple nonlinear models, the existence of indeterminate multivariate RBC-type models and NK models has shown how the E-stability tools can be extended to examine SSEs in linearized models. We develop our discussion in this context, focusing on two special cases: (i) a purely forward-looking univariate or multivariate model, and (ii) a forward and backward-looking univariate model.

4.1 SSEs in a forward-looking model

We start with the simplest possible case of a univariate purely forward-looking linear model

$$y_t = \beta E_t^* y_{t+1}, \tag{31}$$

where $\beta \neq 0$, and for simplicity we have normalized the intercept to zero and omitted exogenous shocks. Since we have not specified the underlying economic model, we will take the reduced-form learning viewpoint. Provided $|\beta| < 1$ there is a unique nonexplosive REE given by the MSV solution

$$y_t \equiv 0, \tag{32}$$

all t. For the PLM $y_t = a$ the corresponding ALM is $y_t = \beta a$ and it follows that the REE (32) is E-stable and hence stable under AL if $\beta < 1$.

If instead $|\beta| > 1$ then (32) remains a solution, and for $\beta < -1$ it remains stable under AL for the PLM $y_t = a$. However, in this case there are other stationary solutions. In particular there are solutions of the form

$$y_t = \beta^{-1} y_{t-1} + \varepsilon_t, \tag{33}$$

where ε_t is, for example, an *iid* exogenous observable white-noise shock. Here ε_t is an extraneous variable, often called a sunspot, and the solutions (33) are called sunspot solutions. It turns out these sunspot solutions are not stable under AL for PLMs of the form $y_t = a + by_{t-1} + \varepsilon_t$.[17]

However, there is another representation of SSEs that *can* be stable under AL. For the case $|\beta| > 1$ consider the SSE solutions

$$\begin{aligned} y_t &= \eta_t \text{ where} \\ \eta_t &= \beta^{-1}\eta_{t-1} + \varepsilon_t. \end{aligned}$$

Here the sunspot η_t is a stationary AR(1) process, assumed observable.[18] Consider now PLMs of the form

$$y_t = a + b\eta_t. \tag{34}$$

The $T-$map is $T(a, b) = (\beta a, b)$ and the eigenvalues of DT are 1 and β. The eigenvalue of 1 is a reflection of the fact that in our linear set-up there is actually a continuum of SSE depending on η_t, i.e. $y_t = \eta_t'$ is also an SSE if η_t' is a scalar multiple of η_t.[19] It follows that

[17] The sunspot representations (33) are sometimes called "general-form" representations, and (34) are then called "common-factor" represeentations.

[18] ε_t is in general a martingale difference sequence, which of course includes *iid* continuously measured white noise processes. However also allowable are finite-state processes that generate finite-state Markov processes for η_t.

[19] This is an artifact of the linear set-up. Also the AR(1) coefficient β^{-1} in (34) is sometimes called the "resonance-frequency." This too is an artifact of the linear specification. In nonlinear models an interval of suitable coefficients is consistent with sunspot equilbria.

if $\beta < -1$ the set of sunspots of the form (34) is E-stable[20] and hence stable under AL, whereas for $\beta > 1$ SSEs are not stable under AL.

This approach extends to multivariate frameworks including those with intrinsic shocks, i.e. to models of the form

$$\begin{aligned} y_t &= ME_t^* y_{t+1} + Pv_t, \\ v_t &= Fv_{t-1} + \tilde{v}_t. \end{aligned}$$

A special case is the standard bivariate NK model with forward-looking interest-rate rule. Here v_t is a vector of observable stationary exogenous shocks and we assume the roots of F are inside the unit circle. If the roots of M are inside the unit circle the model is determinate and there is as unique stationary solution, taking the MSV form $y_t = \bar{c}v_t$. In the indeterminate case, in which one or more roots of M are outside the unit circle there are SSEs taking the form

$$y_t = \bar{c}v_t + \bar{d}\zeta_t$$

for suitable exogenous sunspot processes ζ_t. For example if M has one eigenvalue ϕ with magnitude greater than one, then ζ_t could be a sunspot represented as an AR(1) process with damping coefficient ϕ^{-1}.

As noted in Section 3.2, the NK model with forward-looking interest-rate rules can lead to SSEs that are stable under AL. See Honkapohja and Mitra (2004) and Evans and McGough (2005b) for further discussion.

4.2 SSEs with predetermined variables

Consider now a univariate model taking the form

$$y_t = \beta E_t^* y_{t+1} + \delta y_{t-1} + v_t, \tag{35}$$

where for simplicity we omit an intercept and we now assume v_t is a white noise exogenous process. To keep the analysis simple and generic we continue to take a reduced-form approach to AL in this set-up. We assume $\beta \neq 0, \delta \neq 1$ and $\beta + \delta \neq 1$.

Associated with this model is the quadratic $\beta b^2 - b + \delta = 0$ and we restrict attention to the case in which the roots b_1, b_2 are real. The model is determinate if exactly one of the roots is smaller than one in magnitude and one is larger than one in magnitude. In this case the unique stationary solution takes the MSV form

$$y_t = b_1 y_{t-1} + (1 - \beta b_1)^{-1} v_t$$

[20] Informally, a set is stable if trajectories initialized near the set converge to some point inside the set.

where $|b_1| < 1 < |b_2|$.[21] For PLMs of the form $y_t = a + by_{t-1} + kv_t$ it is straightforward to work out the E-stability conditions and it can be shown that in the determinate case the MSV solution above is E-stable.[22]

In the indeterminate case there are *common-factor* sunspot solutions, depending on a stationary sunspot ζ_t, that take the form

$$\begin{aligned} y_t &= b_1 y_{t-1} + d\zeta_t + (1-\beta b_1)^{-1} v_t, \text{ where} \\ \zeta_t &= b_2 \zeta_{t-1} + \varepsilon_t \end{aligned}$$

for a martingale difference sequence ε_t. Furthermore, it can be shown that there are regions of the parameter space for which these sunspot solutions are E-stable and stable under AL. For a systematic treatment of the stability of SSEs in this model (35) see Evans and McGough (2005c).

4.3 Discussion

In his IMF Blog of Dec. 11, 2011, Olivier Blanchard stated that *" ... the world economy is pregnant with multiple equilibria – self-fulfilling outcomes of pessimism or optimism, with major macroeconomic implications."* This view makes imperative understanding when and how sunspot equilibria, which represent and characterize the class of stationary multiple equilibria, are consistent with the modern DSGE paradigm. Studying the existence and attainability of SSEs, assessed by local stability under AL, are natural tools for assessing the importance of this perspective.

Indeed, considerable work has been done exploring SSEs and their stability under AL in a variety of modern settings, including linearized RBC-type models with nonconvexities, NK models under alternative policy rules, endogenous growth frameworks, and asset-pricing models. Assessments of SSEs in indeterminate RBC-type models are given in Evans and McGough (2005a), Duffy and Xiao (2007) and McGough, Meng, and Xue (2013). Evans, Honkapohja, and Romer (1998) provide an example of stable SSEs in an endogenous growth model. Evans, Honkapohja, and Marimon (2007) show existence of stable SSEs in a cash-in-advance monetary model with seigniorage- and tax-financed government spending. Zanna (2009) shows existence of stable SSEs in a class of small open economy models. Existence and stability under learning of SSEs in regime-switching models has been examined by Branch, Davig, and McGough (2013). Shea (2013) finds stable SSEs in a model of learning by doing with short-sighted managers. For an example of the new generation of bubble models, with stochastically stationary bubbles, Miao, Shen, and Wang (2019) provide an assessment of stability under learning of the alternative solutions and a corresponding reassessment

[21] Here $b_1 = (2\beta)^{-1}\left(1 - \sqrt{1-4\beta\delta}\right)$ and $b_2 = (2\beta)^{-1}\left(1 + \sqrt{1-4\beta\delta}\right)$.

[22] The precise information set needs to be specified. A common assumption is that y_{t-1} and v_t are in the time t information set, but that y_t itself is not observed when $E_t^* y_{t+1}$ is formed. Including y_t in the time t information set can in some cases alter the E-stability conditions.

of the policy implications. Branch, McGough, and Zhu (2019) show that stable "statistical sunspot" equilibria can exist in models with a unique REE if agents use misspecified forecasting models.

There are also several experimental papers that look at the attainability of SSEs in laboratory settings. See, in particular, Marimon, Spear, and Sunder (1993), Duffy and Fisher (2005) and Arifovic, Evans, and Kostyshyna (2019). Finally, recent theoretical research show that existence and assessment of near-rational SSEs can proceed in a systematic way in nonlinear general equilibrium settings using standard tools. In particular Evans and McGough (2020b) show that indeterminacy and stability of the MSV solution in the linearized model imply the existence of stable sunspot equilibria in the linearized model, and stable near-rational sunspot equilibria in the nonlinear model. These findings are collectively referred to as the *MSV Principle*.

5 Other Applications and Extensions

The implications of the adaptive learning approach for macroeconomics are potentially major and wide-ranging. A large body of work has aimed to implement AL in calibrated or estimated models, and in a number of cases link the analysis to survey data.

Orphanides and Williams (2007) show that efficient monetary policies that take account of AL by private agents, and misperceptions of natural rates by policymakers, indicate the need for greater policy inertia, a larger response to inflation, and a smaller response to the perceived unemployment gap than would be optimal under RE. Furthermore policies that would be optimal under RE can perform poorly under AL.

Using US data and Bayesian estimation techniques, Milani (2007) showed that the fit of DSGE models could be substantially increased by replacing RE with AL, and that the resulting parameter estimates indicated that in the AL model mechanical sources of persistence like habit-formation and indexation, are less important than they appear to be under RE. The implications of incorporating AL in applied medium-scale DSGE models has been explored in Slobodyan and Wouters (2012b) and Slobodyan and Wouters (2012a). The latter paper models agents as using small forecasting models. The estimated model is used to explain a decline in the mean and volatility of inflation, and the results are linked to survey evidence on inflation expectations. Using survey data on expectations as well as aggregate macro data, Milani (2011) estimates both the structural parameters of a macro model and expectations shocks, and concludes that expectation shocks can account for about half of business cycle fluctuations.

Using a long-horizon AL real business cycle model calibrated to US data, 1948:I to 2007:IV, Eusepi and Preston (2011) show that the AL model dynamics have several important differences from RE. For example, a smaller productivity innovation variance is needed to fit output volatility, and under AL the IRFs show hump-shaped patterns corresponding to initial overoptimism about future returns to investment and overpessimism concerning future wages. In addition the expectation forecast errors exhibit some features found also in the Survey of Professional Forecasters.

A number of papers have used AL to address various aspects of macroeconomic policymaking. Cogley and Sargent (2005), Bullard and Eusepi (2005), Primiceri (2006) and Sargent, Williams, and Zha (2006) address the rise and fall of inflation in the US over the 1960-1990 period. For a cross-country study of disinflations see Gibbs and Kulish (2017). Hyperinflations in South America were studied by Marcet and Nicolini (2003) and Sargent, Williams, and Zha (2009) using calibrated and estimated models, respectively. Income distribution dynamics in an incomplete markets model with AL is studied in Giusto (2014).

Branch and Evans (2010) and Adam, Marcet, and Nicolini (2016) show how AL can extend the range of asset-price dynamics, and Adam, Marcet, and Nicolini (2016) argue that most of the stylized facts, including several prominent puzzles, can be fully or partly explained using an AL approach. See also Lansing (2010). Additional results, using an "internal rationality" approach, are provided in Adam, Marcet, and Beutel (2017).

Another major policy issue, particularly since the Great Recession, has been the role of the zero lower bound (ZLB). Benhabib, Schmitt-Grohe, and Uribe (2001) showed that under RE a global Taylor-type interest-rate rule that is active at the targeted steady state implies the existence of a second (indeterminate) steady state at a lower rate of inflation (or deflation). Using Euler-equation learning in a nonlinear model, Evans, Guse, and Honkapohja (2008) showed that while the targeted steady state is locally stable under AL, the unintended low-inflation steady state is unstable under learning. However, they argue that the global dynamics reveal the potential for major recessions if there is a sufficiently large pessimistic shock to output and inflations expectations. If the negative expectations shock is large enough the economy enters a deflationary trap region in which output and inflation fall over time with interest rates at the ZLB. However, fiscal policy can be effective in this situation. These results are extended by Benhabib, Evans, and Honkapohja (2014) using a long-horizon model, and show that aggressive fiscal policy can be effective in escaping the stagnation trap even though agents take full account of the tax consequences of fiscal policy. Evans, Honkapohja, and Mitra (2016) develop further policy implications in a set-up in which lower bounds on inflation and consumption imply the existence of a third low output stagnation steady state.

Models in which agents are engaged in adaptive learning also raise new econometric issues concerning identification and the asymptotic distribution for the "external" estimation problem of economists making inferences from the data concerning structural parameters of the model. These issues are discussed in Chevillon, Massmann, and Mavroeidis (2010) and Chritopeit and Massmann (2018).

6 Conclusions

Adaptive learning in macroeconomics is an active field of research with a wide diversity of approaches. The AL approach can be applied to virtually any macroeconomic model in which household and business forecasts play an important role and optimal dynamic decision-making is central. In principle the AL approach can be fully grounded in agent-level decision-making, with explicit aggregation. However, various analytical tools and short-cuts can also provide key results, and the AL approach is well-suited for incorporation into computational models and numerical simulations.

References

Adam, K. (2007): "Experimental Evidence on the Persistence of Output and Inflation," *Economic Journal*, 117, 603–636.

Adam, K., A. Marcet, and J. Beutel (2017): "Stock Price Booms and Expected Capital Gains," *American Economic Review*, 107, 2352–2408.

Adam, K., A. Marcet, and J. P. Nicolini (2016): "Stock Market Volatility and Learning," *Journal of Finance*, 71, 33–82.

Arifovic, J. (1994): "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, 18, 3–28.

——— (1996): "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy*, 104, 510–541.

Arifovic, J., J. Bullard, and J. Duffy (1997): "The Transition from Stagnation to Growth: An Adaptive Learning Approach," *Journal of Economic Growth*, 2, 185–209.

Arifovic, J., J. Bullard, and O. Kostyshyna (2013): "Social Learning and Monetary Policy Rules," *Economic Journal*, 123, 38–76.

Arifovic, J., G. W. Evans, and O. Kostyshyna (2019): "Are Sunspots Learnable? An Experimental Investigation in a Simple Macroeconomic Model," *Journal of Economic Dynamics and Control,* forthcoming.

Azariadis, C. (1981): "Self-Fulfilling Prophecies," *Journal of Economic Theory*, 25, 380–396.

Azariadis, C., and R. Guesnerie (1986): "Sunspots and Cycles," *Review of Economic Studies*, 53, 725–737.

Benhabib, J., G. W. Evans, and S. Honkapohja (2014): "Liquidity Traps and Expectation Dynamics: Fiscal Stimulus or Fiscal Austerity?," *Journal of Economic Dynamics and Control*, 45, 220–238.

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

Benveniste, A., M. Metivier, and P. Priouret (1990): *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, Berlin.

Blanchard, O., and C. Kahn (1980): "The Solution of Linear Difference Models under Rational Expectations," *Econometrica*, 48, 1305–1311.

Branch, W. (2006): "Restricted Perceptions Equilibria and Learning in Macroeconomics," in Colander (2006), pp. 135–160.

Branch, W., T. Davig, and B. McGough (2013): "Adaptive Learning in Regime-Switching Models," *Macroeconomic Dynamics*, 17, 998–1022.

Branch, W., B. McGough, and M. Zhu (2019): "Statistical Sunspots," Working paper, University of Oregon.

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

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

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

——— (2010): "Asset Return Dynamics and Learning," *Review of Financial Studies*, 23, 1651–1680.

——— (2011): "Learning about Risk and Return: A Simple Model of Bubbles and Crashes," *American Economic Journal: Macroeconomics*, 3, 159–191.

——— (2017): "Unstable Inflation Targets," *Journal of Money, Credit and Banking*, 49, 785–806.

Branch, W. A., and B. McGough (2008): "Replicator Dynamics in a Cobweb Model with Rationally Heterogeneous Expectations," *Journal of Economic Behavior and Organization*, 65, 224–244.

Bray, M., and N. Savin (1986): "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, 54, 1129–1160.

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

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.

——— (2014): "When Does Determinacy Imply Expectational Stability?," *International Economic Review*, 55, 1–22.

Bullard, J., G. W. Evans, and S. Honkapohja (2009): "A Model of Near-Rational Exuberance," *Macroeconomic Dynamics*, 14, 166–188.

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

Cass, D., and K. Shell (1983): "Do Sunspots Matter?," *Journal of Political Economy*, 91, 193–227.

Chakraborty, A., and G. W. Evans (2008): "Can Perpetual Learning Explain the Forward Premium Puzzle?," *Journal of Monetary Economics*, 55, 477–490.

Chevillon, G., M. Massmann, and S. Mavroeidis (2010): "Inference in Models with Adaptive Learning," *Journal of Monetary Economics*, 57, 341–351.

Cho, I.-K., and K. Kasa (2008): "Learning Dynamics and Endogenous Currency Crises," *Macroeconomic Dynamics*, 12, 257–285.

——— (2015): "Learning and Model Validation," *Review of Economic Studies*, 82, 45–82.

——— (2017): "Gresham's Law of Averaging," *American Economic Review*, 107, 3589–3616.

Cho, I.-K., N. Williams, and T. J. Sargent (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

Chritopeit, N., and M. Massmann (2018): "Estimating Structural Parameters in Regression Models with Adaptive Learning," *Econometric Theory*, 34, 68–111.

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

Cochrane, J. H. (2015): "Do Higher Interest Rates Raise or Lower Inflation?," Working paper, University of Chicago Booth School of Business.

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.

Colander, D. (2006): *Post Walrasian Macroeconomics*. Cambridge, Cambridge, U.K.

Duffy, J., and E. Fisher (2005): "Sunspots in the Laboratory," *American Economic Review*, 95, 510–529.

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 (2011): "Expectations, Learning and Business Cycle Fluctuations," *American Economic Review*, 101, 2844–2872.

——— (2018): "The Science of Monetary Policy: An Imperfect Knowledge Perspective," *Journal of Economic Literature*, 56, 3–59.

Evans, G. W. (1985): "Expectational Stability and the Multiple Equilibria Problem in Linear Rational Expectations Models," *The Quarterly Journal of Economics*, 100, 1217–1233.

——— (1989): "The Fragility of Sunspots and Bubbles," *Journal of Monetary Economics*, 23, 297–317.

Evans, G. W., C. G. Gibbs, and B. McGough (2019): "A Unified Model of Learning to Forecast," mimeo, University of Oregon.

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

Evans, G. W., R. Guesnerie, and B. McGough (2019): "Eductive Stability in Real Business Cycle Models," *Economic Journal*, 129, 821–852.

Evans, G. W., E. Guse, and S. Honkapohja (2008): "Liquidity Traps, Learning and Stagnation," *European Economic Review*, 52, 1438–1463.

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.

——— (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.

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

——— (2009a): "Expectations, Learning and Monetary Policy: An Overview of Recent Research," chap. 2, pp. 27–76. Central Bank of Chile, Santiago.

——— (2009b): "Learning and Macroeconomics," *Annual Review of Economics*, 1, 421–451.

——— (2013): "Learning as a Rational Foundation for Macroeconomics and Finance," in Frydman and Phelps (2013), chap. 2, pp. 68–111.

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

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

Evans, G. W., S. Honkapohja, and K. Mitra (2016): "Expectations, Stagnation and Fiscal Policy," Discussion paper 11428, CEPR.

Evans, G. W., S. Honkapohja, and P. Romer (1998): "Growth Cycles," *American Economic Review*, 88, 495–515.

Evans, G. W., S. Honkapohja, T. J. Sargent, and N. Williams (2013): "Bayesian Model Averaging, Learning and Model Selection," chap. 7. Oxford University Press.

Evans, G. W., S. Honkapohja, and N. Williams (2010): "Generalized Stochastic Gradient Learning," *International Economic Review*, 51, 237–262.

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.

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

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

——— (2018a): "Equilibrium Selection, Observability and Backward-Stable Solutions," *Journal of Monetary Economics*, 98, 1–10.

——— (2018b): "Interest-Rate Pegs in New Keynesian Models," *Journal of Money, Credit and Banking*, 50, 939–965.

——— (2018c): "Learning to Optimize," mimeo, University of Oregon.

——— (2019): "Equilibrium Stability in a Nonlinear Cobweb Model," mimeo, University of Oregon.

——— (2020a): "Agent-level Adaptive Learning," *Oxford Research Encyclopedia of Economics and Finance,* in preparation.

——— (2020b): "Stable Near-rational Sunspot Equilibria," *Journal of Economic Theory,* forthcoming.

Evans, G. W., and G. Ramey (1992): "Expectation Calculation and Macroeconomic Dynamics," *American Economic Review*, 82, 207–224.

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

Fahri, E., and I. Werning (2019): "Monetary Policy, Bounded Rationality, and Incomplete Markets," *American Economic Review*, 109, 3887–3928.

Frydman, R., and E. e. Phelps (eds.) (2013): *Rethinking Expectations: The Way Forward for Macroeconomics*. Princeton University Press, Princeton, New Jersey.

Garcia-Schmidt, M., and M. Woodford (2019): "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis," *American Economic Review*, 109, 86–120.

Gibbs, C. G. (2017): "Forecast Combination, Non-linear Dynamics, and the Macroeconomy," *Economic Theory*, 63, 653–686.

Gibbs, C. G., and M. Kulish (2017): "Disinflations in a Model of Imperfectly Anchored Expectations," *European Economic Review*, 100, 157–174.

Giusto, A. (2014): "Adaptive Learning and Distributional Dynamics in an Incomplete Markets Model," *Journal of Economic Dynamics and Control*, 40, 317–333.

Granato, J., E. A. Guse, and M. S. Wong (2008): "Learning from the Expectations of Others," *Macroeconomic Dynamics*, 12, 345–377.

Grandmont, J.-M. (1985): "On Endogenous Competitive Business Cycles," *Econometrica*, 53, 995–1045.

Guesnerie, R. (1992): "An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis," *American Economic Review*, 82, 1254–1278.

——— (2005): *Assessing Rational Expectations 2: Eductive Stability in Economics*. MIT Press, Cambridge, Mass.

Guse, E. (2008): "Learning in a Misspecified Multivariate Self-Referential Linear Stochastic Model," *Journal of Economic Dynamics and Control*, 32, 1517–1542.

——— (2010): "Heterogeneous Expectations, Adaptive Learning, and Evolutionary Dynamics," *Journal of Economic Behavior and Organization*, 74, 42–57.

Hicks, J. R. (1946): *Value and Capital, Second edition.* Oxford University Press, Oxford UK.

Hommes, C. (2013): *Behavioral Rationality and Heterogeneous Expectations in Complex Economics Systems.* Cambridge University Press, Cambridge, UK.

Hommes, C. H. (2011): "The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab," *Journal of Economic Dynamics and Control*, 35, 1–24.

Hommes, C. H., and M. Zhu (2014): "Behavioral Learning equilibria," *Journal of Economic Theory*, 150, 778–814.

Honkapohja, S., and K. Mitra (2004): "Are Non-Fundamental Equilibria Learnable in Models of Monetary Policy?," *Journal of Monetary Economics*, 51, 1743–1770.

——— (2019): "Price Level Targeting with Evolving Credibility," *Journal of Monetary Economics,* forthcoming.

Howitt, P. (1992): "Interest Rate Control and Nonconvergence to Rational Expectations," *Journal of Political Economy*, 100, 776–800.

Kasa, K. (2004): "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review*, 45, 141–173.

Lansing, K. (2010): "Rational and Near-rational Bubbles without Drift," *Economic Journal*, 120, 1149–1174.

Ljung, L. (1977): "Analysis of Recursive Stochastic Algorithms," *IEEE Transactions on Automatic Control*, 22, 551–575.

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

Marcet, A., and T. J. Sargent (1989a): "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, 97, 1306–1322.

——— (1989b): "Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

Marimon, R., S. E. Spear, and S. Sunder (1993): "Expectationally Driven Market Volatility: An Experimental Study," *Journal of Economic Theory*, 61, 74–103.

MARKIEWICZ, A. (2012): "Model Uncertainty and Exchange Rate Volatility," *International Economic Review*, 53, 815–843.

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

MCGOUGH, B. (2006): "Shocking Escapes," *Economic Journal*, 116, 507–528.

MCGOUGH, B., Q. MENG, AND J. XUE (2013): "Expectational Stability of Sunspot Equilibria in Non-convex Economies," *Journal of Economic Dynamics and Control*, 37, 1126–1141.

MCGOUGH, B., G. RUDEBUSCH, AND J. WILLIAMS (2005): "Using Long-Term Interest Rate as the Monetary Policy Instrument," *Journal of Monetary Economics*, 52, 855–879.

MIAO, J., Z. SHEN, AND P. WANG (2019): "Monetary Policy and Rational Asset Price Bubbles: Comment," *American Economic Review*, 109, 1969–1990.

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

——— (2011): "Expectation Shocks and Learning as Drivers of the Business Cycle," *Economic Journal*, 121, 379–401.

MUTH, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, 315–335.

NAGEL, R. (1995): "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85, 1313–1326.

ORPHANIDES, A., AND J. C. WILLIAMS (2007): "Robust Monetary Policy with Imperfect Knowledge," *Journal of Monetary Economics*, 54, 1406–1435.

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

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

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.

——— (2009): "The Conquest of South American Inflation," *Journal of Political Economy*, 117, 211–256.

Shea, P. (2013): "Learning By Doing, Short-Sightedness, and Indeterminacy," *Economic Journal*, 123, 738–763.

Shell, K. (1977): "Monnaie et Allocation Intertemporelle," Working paper, CNRS Seminaire de E.Malinvaud, Paris.

Shiller, R. J. (1981): "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends," *America Economic Review*, 71, 421–436.

Sims, C. A. (2001): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

Sinha, A. (2016): "Learning and the Yield Curve," *Journal of Money, Credit and Banking*, 48, 513–547.

Slobodyan, S., and R. Wouters (2012a): "Estimating a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models," *American Economic Journal: Macroeconomics*, 4, 65–101.

——— (2012b): "Learning in an Estimated DSGE Model," *Journal of Economic Dynamics and Control*, 36, 26–46.

Williams, N. (2019): "Escape Dynamics in Learning Models," *Review of Economic Studies*, 86, 882–912.

Woodford, M. (1990): "Learning to Believe in Sunspots," *Econometrica*, 58, 277–307.

——— (2013): "Macroeconomic Analysis without the Rational Expectations Hypothesis," *Annual Review of Economics*, 5, 303–346.

Zanna, L.-F. (2009): "PPP Rules, Macroeconomic (In)stability and Learning," *International Economic Review*, 50, 1103–1128.