Adaptive Learning and the Restricted Perceptions Approach

William A. Branch

University of California, Irvine

November 10, 2023

Abstract

Rational expectations makes overly strong cognitive and informational assumptions while being empirically unrealistic. Adaptive learning is an alternative modeling paradigm that assumes agents generate their expectations like a good econometrician who specifies and estimates forecasting models. In many models, though, people will identify and estimate, using available data, a parametric family of (linear) forecast models that are almost always misspecified. Nevertheless, despite the misspecification, they do not make systematic forecast errors. In a restricted perceptions equilibrium (RPE), each agent uses an optimal forecast model among the candidates under consideration. Notably, like a rational expectations equilibrium, an RPE is a Nash equilibrium concept: the optimal forecast model for an individual depends on the behaviors and hence the models used by other agents. The restricted perceptions approach brings realism and cognitive consistency into models of expectation formation while preserving the model consistency that is the hallmark of rational expectations models—this chapter overviews the adaptive learning and restricted perceptions approach.

JEL Classification: D82; D83; E40; E50

Keywords: Expectations, learning, bounded rationality, misspecification.

1. Introduction

Modern macroeconomic models place people and firms in environments where they make repeated decisions over time and under uncertainty. Consequently, macroeconomic behavior and outcomes depend on people's expectations about future events. For instance, when people expect prices to rise, they act in ways that make prices rise. Businesses set higher prices in anticipation of higher costs, workers demand higher wages, and consumers buy now rather than later. One of the most critical and open issues is: how do people formulate their expectations?

Since the 1970s, the prevailing macroeconomic paradigm is rational expectations. The rational expectations hypothesis aligns subjective expectations with actual outcomes, a strong assumption for a theory of expectation formation that requires people to make correct predictions given the information available to them. At the same time, people's subjective beliefs could, by pure chance, equate to the actual probability distribution, which depends on those beliefs. A theory of purposeful behavior would hold that agents could calculate the correct expectations taking into account all information. However, calculating rational beliefs requires agents to correctly understand the complete economic structure, the exogenous state variables, and the preferences and beliefs of all other individuals.

Instead, many models depart from rational expectations through bounded rationality in beliefs or decision-making. Sargent (1993) warned researchers about the dangers of the "wilderness of bounded rationality." The appeal of rational expectations is that they are model consistent. Equating subjective expectations and actual outcomes within a model environment removes beliefs as an independent macroeconomic force. Once departing from model consistency, it is difficult to know where to add modeling discipline. Evans and Honkapohja (2001) argue in favor of a cognitive consistency principle: agents should be modeled like a good economist who specifies, estimates, and revises models. The cognitive consistency principle says that the model can still discipline beliefs by determining the state variables agents include in their forecasting models. An estimation process, though, will lead to real-time learning that may or may not converge to rational expectations. This adaptive learning approach encompasses a broad range of schemes for how people might forecast variables based on observed data.

Once adopting the cognitive consistency principle naturally leads us down a path where boundedly rational forecasting models do not nest rational expectations as a particular case. White (1994) begins his text on econometric inference by noting that all models are misspecified. The restricted perceptions approach is a burgeoning field of bounded rationality that studies the consequences of misspecified beliefs in macroeconomic models. Restricted perceptions are similar to rational expectations except the agents' own subjective model misses the truth in some dimension. With restricted perceptions, agents restrict attention to a (linear) parametric family of forecast models that are misspecified in some dimension, e.g., the number of lags, variables, and linearity, among others. However, it is possible to discipline beliefs by asking the agents to do their best, given their misspecification. In a restricted perceptions equilibrium, agents' perceived model is the projection of the endogenous variable on their restricted approximating model – that is, it is the best forecast given the agent's information, knowledge, and abilities. Thus, restricted perceptions preserve the discipline of rational expectations, while allowing for substantial departures. This chapter overviews the restricted perceptions approach and its applications.

Restricted perceptions grant a significant and plausible role to beliefs in driving macroeconomic fluctuations, exemplified notably in the context of sunspot phenomena. It is well-established that within the rational expectations framework, economic outcomes can become self-fulfilling prophecies, known as sunspot equilibria (Grandmont (1985)). Here, the core idea is that individuals might alter their expectations based on purely extrinsic events, termed 'sunspots', leading to increased consumption and production. However, these sunspot equilibria are often criticized for their fragility and instability, particularly when subject to adaptive learning processes. In contrast, under the lens of restricted perceptions, individuals' reactions to sunspot events are robust. These reactions might be influenced by sunspots because they act as proxies for overlooked real-world elements in their forecasting models. Such elements could include critical but subtle factors like nonlinear economic relationships or the concealed policy preferences of central banks. This nuanced approach allows for a more stable and realistic representation of how beliefs, even when imperfectly formed, can significantly impact macroeconomic dynamics.

Section 2 introduces the model, while section 3 studies the rational expectations solution and introduces the idea of adaptive learning and expectational stability ("E-stability"). In section 4, there is an overview of the restricted perceptions approach, and section 5 illustrates several economic applications. Section 6 surveys the literature and concludes.

2. Model

To provide a clear foundation, we start with a simple reduced-form model. Consider the relationship of an economic variable y_t to its subjective future expectations $\hat{E}_t y_{t+1}$, and an exogenous variable z_t, representing uncertain events. We will delve into concrete economic applications later in this chapter.

2.1 Temporary equilibrium and adaptive learning

In a celebrated contribution, Grandmont (1977) focused on the concept of a "temporary equilibrium." Here, economic transactions occur over time, and decisions are based on the expectations held at each moment. The idea of adaptive learning posits that these expectations, and consequent decisions, evolve in response to accumulating data. Our model encapsulates this in a linear expectational framework:

$$y_t = \alpha + \beta \hat{E}_t y_{t+1} + \gamma' z_t \tag{1}$$

where $\alpha, \beta \in \mathbb{R}, \hat{E}_t$ represents the expectations operator (details below), and z_t is an $(n \times 1)$ vector of stationary exogenous variables.[1] Throughout, assume that

$$z_t = \rho z_{t-1} + \varepsilon_t$$

is stationary with eigenvalues of ρ inside the unit circle and $\varepsilon_t \sim N(0, \Sigma_\varepsilon)$.

2.2 Rational vs. boundedly rational expectations

When $\hat{E}_t = E_t$, agents have rational expectations, perfectly predicting y_{t+1} except for unforeseeable shocks like ε_{t+1}. In this scenario, equation (1) becomes an expectational difference equation. A rational expectations equilibrium is a stable solution to this equation, where long time series for y_t are determined partly by its one-step-ahead expected value $E_t y_{t+1}$. This self-referential nature can lead to self-fulfilling phenomena like asset-price bubbles and sunspot shocks.

This chapter, however, focuses on bounded rationality models. These models allow the expectation operator to deviate from conditional mathematical expectations, mapping unobservable random variable to functions of observables. Rational expectations equate this operator with the orthogonal projection onto the space of measurable functions generated by observables, i.e., the conditional expectations operator. The adaptive learning approach, as detailed in Evans and Honkapohja (2001) maintains model consistency without assuming perfect alignment between expectations and outcomes. Agents in this framework possess a linear forecasting model, akin to rational agents, but with an evolving understanding of relationships within the mnodel, akin to an econometrician

[1] More general reduced form systems include lagged endogenous variables and non-linearities. Restricting attention to models in the form of (1) is for expositional ease. The restricted perceptions approach has been applied in fully general environments.

regressing y_t on known variables over time.

2.3 Restricted perceptions and model misspecification

The restricted perceptions approach extends beyond adaptive learning and rational expectations, accommodating situations where agents' forecasting model differ from those posited by rational expectations. This approach acknowledges that the actual data generating process might remain perpetually unknown to agents. Let the actual data generating process produce a probability distribution over outcomes, say $f(y^t|\phi)$, where ϕ is a parameter vector in the actual data generating process. Then an approximating model can also be described by a probability distribution over outcomes, but with an alternative parameterization: $f(y^t|\theta_j)$. One can index different approximating models by the parameterization θ_j. As a concrete example, restricted perceptions might underparameterize a model by omitting some variable(s) in z_t.

The rational expectations hypothesis equates the approximating, or subjective model, with the true model: $f(y^t|\phi) = f(y^t|\theta_j)$. Theoretically, there are things to like about rational expectations. There is model consistency in agents' expectations by equating outcomes with subjective beliefs. A well-specified learning process could lead the agents' approximating model to converge to the rational expectations equilibrium. The stability of rational expectations equilibria is the subject of Evans and Honkapohja (2001). However, the assumption is strong, and it may not be reasonable to expect alignment between subjective beliefs and actual outcomes in many settings. The learning models described in this chapter preserve many of the salient features of rational expectations by imposing a set of moment conditions on the approximating model so that, at least for long periods, the agent within the model would be unable to detect their misspecification. The case for these moment conditions, and their equilibrium implications, is the focus of the present chapter.

Sargent (1993) highlighted the dilemma faced by research into bounded rationality. On the one hand, rational expectations are an a priori strong and unreasonable assumption. On the other hand, the model disciplines rational expectations, and departing from it may land one in a "wilderness of bounded rationality." Evans and Honkapohja (2001) advocate for disciplining bounded rationality via the cognitive consistency principle, which holds that one should model agents' forecasting behavior as if they were a good econometrician. Applied econometricians specify, estimate, and revise econometric models. Thus, the early literature on adaptive learning assumed that agents forecast based on linear econometric models that nest rational expectations. The learning process occurs as agents revise their coefficient estimates by least squares as the economy generates new data over time. In a wide class of models, least-squares learning converges to rational expectations.

However, it was White (1994) who opened his classic econometrics book by observing that all models are misspecified. Econometricians often face degrees of freedom problems that force them to specify parsimonious models. The actual data-generating process may be non-linear, but econometricians typically estimate linear models. A wide variety of shocks impact the economy, and only a subset of those exogenous variables may be observable by the econometrician.

2.4 An example environment

The environment is a mean-variance linear asset pricing model similar to De Long, Shleifer, Summers, and Waldmann (1990) as developed in Branch and Evans (2011a). There is a continuum of agents born each time t, indexed by $\omega_t \in \Omega$. Each agent lives for two periods, and discounts at the rate $0 < \beta < 1$. The number n_t of young agents is iid with $En_t^{-1} = 1$. Each agent receives an endowment of y when young and consumes only when old. The endowment is non-storable, but agents save using one of two assets: a riskless storage technology with gross return $R = \beta^{-1} > 1$ payable when old; a risky asset, in the form of a Lucas tree, which is in fixed outside supply s_0, with claims to the tree traded competitively at a price p_t. An agent of type ω holds s_{dt} shares in the risky asset. The risky asset pays a stochastic dividend q_{t+1}. Because the population of young agents is random, the per-capita supply of the asset $z_t = s_0/n_t$ is also random. Finally, young agents also face risk in the form of an idiosyncratic asset float shock $f_t(\omega)$ that randomly redistributes holdings of the asset among old agents. The asset float shock proxies for idiosyncratic variations in asset float because of, for instance, lock-up expirations. The focus here is on a small noise limit.

2.4.1 Temporary equilibrium

Preferences are of the CARA form

$$U\left(c_{t+1}\right) = -\exp\left\{-ac_{t+1}\right\}$$

with $a > 0$ is the coefficient of absolute risk aversion. Agents assume a normal distribution for $p_{t+1} + q_{t+1}$, which leads to the household portfolio decision for young agent ω in the form of a mean-variance optimization problem:

$$\max_{s_{dt}} - \exp\left\{-a\hat{E}_t c_{t+1} + \left(a^2/2\right) \mathrm{Var}_t^* c_{t+1}\right)$$

subject to

$$c_{t+1} = (y - p_t s_{dt}(\omega))\,\beta^{-1} + f_t(\omega) s_{dt}(\omega)\,(p_{t+1} + q_{t+1})$$

Here, $\hat{E}_t$ denotes the subjective conditional expectation and Var_t^* the subjective conditional variance. Taking conditional subjective expectations leads to the first-order condition

$$-\beta^{-1} p_t + f_t(\omega)\hat{E}_t\,(p_{t+1} + q_{t+1}) - a\sigma^2 s_{dt} = 0$$

Here we assume that subjective conditional variance expectations are homogeneous and, without a loss of generality, we take σ^2 to be exogenous and time-invariant.[2]

Market equilibrium requires that $s_0 = \int s_{dt}(\omega)d\omega$. Integrating the first-order condition across agents leads to

$$\frac{a\sigma^2 s_0}{n_t} = \int f_t(\omega)\hat{E}_t^{\omega}(p_{t+1} + q_{t+1})d\omega - \beta^{-1} p_t$$

Now, we focus on the small noise limit $f_t(\omega) \to 1$, so that the temporary equilibrium[3] asset-pricing condition becomes

$$p_t = \beta\hat{E}_t\,(p_{t+1} + q_{t+1}) - a\sigma^2\beta s_t$$

or, equivalently,

$$p_t = \beta\hat{E}_t\,(p_{t+1} + q_{t+1}) + \gamma s_t \tag{2}$$

where $\hat{E}(x) = \int E^{\omega}(x)d\omega$ is the aggregate (linear) subjective expectations operator and s_t is the stochastic share supply s_0/n_t. Finally, we assume that the share supply follows a stationary AR(1) process

$$s_t = \rho s_{t-1} + \varepsilon_t \tag{3}$$

Equations (1)-(3) provide the (temporary) equilibrium value for price p_t given the aggregate expectations operator $\hat{E}_t$. The structural model takes the same form as the

[2]Branch and Evans (2011a) focuses on an environment where the endogeneity of σ^2 plays a central role in generating asset price bubbles and crashes.

[3]A temporary equilibrium is a time t market equilibrium in which agents solve their optimization problem taking as given own subjective expectations over payoff-relevant variables whose determination is treated as exogenous by each of the agents.

expectational difference equation in (1) with $z_t' = \left(\hat{E}_t q_{t+1}, s_t\right)$. The next section describes a variety of behavioral assumptions that lead to misspecification and restricted perceptions equilibria.

3. Rational expectations equilibrium and E-stability

In order to solve the temporary equilibrium equation (1) we need to make specific assumptions about the expectations operator. This section assumes rational expectations, $\hat{E}_t = E_t$, solves for the rational expectations equilibria, and then uses the adaptive learning approach to assess the plausibility of the equilibrium.

3.1 A unique rational expectations equilibrium

Let the expectational feedback parameter $0 < \beta < 1$ in (1). In this case, there is a unique rational expectations equilibrium of the form

$$y_t = \bar{a} + \bar{b}' z_t$$

This is a minimal state variable solution as y_t is an exact linear function of the state variables in (1). One can find the rational expectations equilibrium values $\left(\bar{a}, \bar{b}\right)$ by the method of undetermined coefficients. From the linear form,

$$E_t y_{t+1} = \bar{a} + \bar{b}' \rho z_t$$

in which case, the true data-generating process is

$$y_t = \alpha + \beta \bar{a} + \left[\beta \bar{b}' \rho + \gamma'\right] z_t$$

Equating the linear form and the true data-generating process leads to

$$\begin{aligned} \bar{a} &= \alpha + \beta \bar{a} \\ \bar{b} &= \left[\beta \bar{b}' \rho + \gamma'\right] \end{aligned}$$

Or, after solving,

$$\begin{aligned}\bar{a} &= \frac{\alpha}{1-\beta}\\ \bar{b} &= (I-\beta\rho)^{-1}\gamma\end{aligned}$$

The method of undetermined coefficients illustrates how rational expectations aligns beliefs and outcomes. If agents believe that y_t is linearly related to the variables in z_t, then their expectations will be $E_t y_{t+1} = \bar{a} + \bar{b}'\rho z_t$. Then the temporary equilibrium (1), given those expectations, says that outcomes are

$$y_t = \alpha + \beta\bar{a} + \left[\beta\bar{b}'\rho + \gamma'\right] z_t$$

So, rational agents think that y_t is linearly related only to z_t – this is the minimal set of state variables – and, consequently, y_t is linearly related only to z_t in actuality. Beliefs and outcomes are aligned.

One way in which the adaptive learning approach is useful is to ask, if the agents did not know $\bar{a}$ and $\bar{b}$, but formed estimates of it through least-squares regressions of y_t on a constant and z_t, would those estimated values a_t, b_t eventually converge to $\bar{a}, \bar{b}$? Evans and Honkapohja (2001) develop the E-stability principle and demonstrate how an "E-stability condition" can determine whether a rational expectations equilibrium is learnable under a reasonable learning rule, e.g., least-squares learning. The approach is similar to the method of undetermined coefficients. Assume agents' beliefs come from an approximating model, called a perceived law of motion, of the form

$$y_t = a + b'z_t \Rightarrow \hat{E}_t y_{t+1} = a + b'\rho z_t \tag{4}$$

The perceived law of motion is a linear forecasting rule of the same form as the unique rational expectations equilibrium but for an arbitrary parameterization (a, b). Given those expectations, plugging into (1) leads to the actual law of motion implied by the perceived law of motion:

$$y_t = \alpha + \beta a + [\beta b'\rho + \gamma'] z_t \tag{5}$$

Equating the perceived and actual laws of motion leads to the same parameter values $(\bar{a}, \bar{b})$ uncovered by the method of undetermined coefficients.

The E-stability approach, though, notes that (5) can be rewritten

$$y_t = T(a,b)' X_t$$

where $X_t' = (1, z_t)$ and

$$T(a,b)' = [\alpha + \beta a, \beta b' \rho + \gamma']$$

A rational expectations equilibrium is a fixed point to the "T-map": $\left(\bar{a}, \bar{b}\right) = T\left(\bar{a}, \bar{b}\right)$. Up to this point, we are exactly describing the method of undetermined coefficients.

It turns out, though, that for a wide range of learning algorithms, convergence of real-time learning – in the sense that $\left(\bar{a}_t, \bar{b}_t\right) \to \left(\bar{a}, \bar{b}\right)$ – is governed by the local asymptotic stability of the E-stability ordinary differential equation:

$$\frac{d\,(a,b)}{d\tau} = T(a,b) - (a,b) \tag{6}$$

A rational expectations equilibrium is expectationally stable (E-stable) provided that $(\bar{a}, \bar{b})$ is a locally stable resting point to (6). That is, for (a,b) initially near $(\bar{a}, \bar{b})$ we have ${d(a,b)}/{d\tau} \to 0$ as $\tau \to 0$. The variable τ is notional time as (6) delivers learnability conditions but is not, strictly speaking, a real-time learning algorithm: see below. Notice that the rational expectations equilibrium $(\bar{a}, \bar{b})$ is a resting point of the E-stability equation. The T-map, intuitively, tells us what coefficients an econometrician would recover with a long data history if (a,b) remained constant. Thus, the E-stability o.d.e. says that a good learning algorithm will adjust the approximating model whenever the actual coefficients $T(a,b)$ differ from the perceived coefficients (a,b). If that adjustment moves towards the rational expectations equilibrium, it is E-stable.

We can compute those E-stability conditions by examining the eigenvalues of the Jacobian to the right-hand side of (6), evaluated at the rational expectations equilibrium. E-stability then gives us conditions under which a small deviation from $\bar{a}, \bar{b}$ would converge back to its rational expectations equilibrium values. The minimal state variable solution has eigenvalues:

$$\begin{aligned} DT_a(\bar{a}, \bar{b}) &= \beta \\ DT_b(\bar{a}, \bar{b}) &= \beta\rho \end{aligned}$$

Recalling that ρ is a diagonal matrix whose values are less than one, it follows that the key condition is that $DT_a(\bar{a}, \bar{b}) = \beta < 1$. We conclude, when $0 < \beta < 1$ and there is a unique rational expectations equilibrium, it is also learnable.

3.2 Multiple rational expectations equilibria

Suppose instead that $|\beta| > 1$. The model is now indeterminate with a continuum of possible rational expectations equilibria. It turns out that the full class of solutions is of the form

$$y_t = \bar{a} + \bar{b}' z_t + \bar{c} y_{t-1} + \bar{d}\eta_t$$

where η_t is an extrinsic random variable satisfying $E_{t-1}\eta_t = 0$, i.e. a sunspot variable. Sunspot equilibria give agents' beliefs an independent role in the economy as they allow dependence on a self-fulfilling sunspot variable. In many cases, however, sunspot variables are not stable under learning.

The full class of solutions now extends beyond the minimal number of variables and includes lags of y and the extrinsic sunspot shock η_t. To see that this is an equilibrium, consider a simplified version of (1) where $\alpha = 0, \gamma = 0$. Then,

$$y_t = \beta \hat{E}_t y_{t+1}$$

Notice that $y_t = 0$ is a rational expectations equilibrium in this simplified setting: if $y_t = 0$ then $E_t y_{t+1} = 0$ and $y_t = \beta \times 0 = 0$. Suppose instead that the agents' perceive that

$$y_t = c y_{t-1} \Rightarrow \hat{E}_t y_{t+1} = c \hat{E}_t y_t = c^2 y_{t-1}$$

The minimal state variable solution $y_t = 0$ arises with $c = 0$. Without imposing anything on c, though, the actual law of motion becomes

$$y_t = \beta c^2 y_{t-1} \equiv T(c) y_{t-1}$$

The method of undetermined coefficients solves for the rational expectations equilibrium by computing $\bar{c} = T(\bar{c})$. In this simplified case, the rational expectations equilibrium values for c are either 0 or $1/\beta$. So, we have two different types of rational expectations equilibria. Only one – the minimal state variable solution – is learnable. Take the equilibrium with $\bar{c} = 1/\beta$, and notice that $T'(\bar{c}) = T'(\beta^{-1}) = 2\beta\beta^{-1} = 2 > 1$. So if agents believe that $y_t = c y_{t-1}$ with c near $1/\beta$, their learning will push them away from $1/\beta$. Sunspot rational expectations equilibria are not E-stable.

The remainder of the chapter examines the implications of instances where the approximating model does not nest the actual model. The focus is on reasonable restrictions on the approximating models that lead to a restricted perceptions equilibrium. Novel economic phenomena will arise, including the existence of E-stable sunspot equilibria.

3.3 Why E-stability?

Under least-squares learning, agents update (a_t, b_t, c_t) each period using all data up to that point. That is, coefficient estimates come from a regression model

$$y_t = \theta' x_t + e_t$$

where $x_t = (1, z_t', y_{t-1})'$ and $\theta' = (a, b', c)$. The full-sample least squares estimate based on data from $t = 1, ..., T$ is

$$\theta_T = \left(\sum_{t=1}^{T} x_t x_t' \right) \left(\sum_{t=1}^{T} \right)^{-1}$$

It turns out, though, that with the right initialization it is possible to write the least-squares estimator recursively as

$$\begin{aligned} \theta_t &= \theta_{t-1} + t^{-1} R_t^{-1} x_t \left(y_t - x_t' \theta_{t-1} \right) \\ R_t &= R_{t-1} + t^{-1} \left(x_t x_t' - R_{t-1} \right) \end{aligned}$$

where R_t is the sample second-moment matrix of the regressors. Using induction, it is possible to prove the equivalence of these formulas when R is initialized using the first k observations as $R_k = k^{-1} \sum_{i=1}^{k} x_i x_i'$.

The E-stability principle draws a connection between the conditions that govern how θ_t in recursive least squares behaves asymptotically with the conditions of the E-stability ordinary differential equation. Readers are referred to Evans and Honkapohja (2001) for technical details. Here, we can infer the connection.

The previous analysis showed how we can write the actual law of motion implies by agents' beliefs using the T-map:

$$y_t = T(\theta)' x_t$$

now using our new notation for θ and x. Plug this into the recursive updating equation for θ_t:

$$\theta_t = \theta_{t-1} + t^{-1} R_t^{-1} x_t \left[x_t' T(\theta_{t-1}) - x_t' \theta_{t-1} \right]$$

Now, let's rewrite again,

$$\frac{\theta_t - \theta_{t-1}}{t^{-1}} = R_t^{-1} x_t x_t' \left[T\left(\theta_{t-1}\right) - \theta_{t-1} \right]$$

while noticing that the term in the brackets is now the right hand side of the E-stability ordinary differential equation. The remainder of the derivation takes a continuous time limit and observes that the recursive second moment updating equation implies that $R_t^{-1} x_t x_t' \to I$.

The E-stability concept is important with other forms of learning besides adaptive and least-squares. Guesnerie (1992) introduced the idea of eductive, or rationalizable, reasoning into macroeconomic models by assuming that agents form their expectations by considering the range of possible actions and outcomes that other rational agents in the market might choose. This process involves a sophisticated level of strategic thinking, where each agent's expectations and actions are contingent on their beliefs about other agents' beliefs and behaviors. The agents, through this interactive process, coordinate their beliefs within the model, leading to outcomes that are consistent with rational behavior under specific common knowledge assumptions on higher order beliefs. Building upon this concept, Evans and Guesnerie (1999) and Evans and Guesnerie (2000) introduce Eductive E-stability to assess the stability of eductive reasoning. Their approach emphasizes the role of logical, forward-looking reasoning and the coordination among agents in achieving equilibrium outcomes. Eductive E-stability delves into the conditions under which the economy's equilibrium is not just learned adaptively based on historical data, but is also 'educed' or logically inferred by agents who are actively engaged in deducing the model's implications. This approach bridges the gap between the theoretical underpinnings of rational expectations and the practical aspects of how agents might realistically form and coordinate their expectations, highlighting a more dynamic and interactive process of equilibrium attainment in economic models.

E-stability is a simple way to calculate learnability conditions for least-squares learning. Yet, it is built on powerful machinery.

4. Restricted perceptions equilibrium

A restricted perceptions equilibrium (RPE) allows for misspecification in the approximating models entertained by agents. An RPE maintains cross-equation restrictions like the rational expectations hypothesis by requiring that the approximating model is the optimal linear projection within its class. This section describes a few common examples of misspecification and restricted perceptions equilibria. The first example is the case where the approximating models are under-parameterized. In the second example, the actual model is non-linear, but agents forecast with a linear model. The third example assumes that the actual model contains hidden state variables unobserved by agents.

4.1 Under-parameterization

As a first example, consider the case of (1) with bivariate exogenous shocks:

$$y_t = \beta \hat{E}_t y_{t+1} + \gamma_1 z_{1t} + \gamma_2 z_{2t}$$

and

$$z_{jt} = \rho_j z_{jt-1} + \varepsilon_{jt}, j = 1, 2$$

The innovations ε_{jt} are mean-zero with variances σ_j^2 and $E\varepsilon_{1t}\varepsilon_{2t} = \sigma_{12}$.

In a complete information environment, the rational expectations equilibrium is

$$y_t = \frac{\gamma_1}{1 - \beta\rho_1} z_{1t} + \frac{\gamma_2}{1 - \beta\rho_2} z_{2t}$$

The effects of the shocks have a direct effect parameterized by γ_1. The shocks also have an indirect effect that arises through the self-referential features of the model: y_t depends on expectations about its future values. The rational expectations equilibrium is the expected present value of the direct effects.

Branch and Evans (2006) note that in some situations, the cognitive consistency principle would lead agents to specify parsimonious forecasting models. In many environments, econometricians face a degree of freedom limitation that leads them to pare down the number of lags or exogenous variables. Forecasters often find, in environments with structural change of unknown form, that simple, parsimonious models perform better. In this simple univariate example, complexity is not an issue, but it provides an analytical example of how under-parameterization can alter the equilibrium dynamics.

Suppose that agents have a perceived law of motion (approximating model) that conditions only on z_{1t}:

$$y_t = b_1 z_{1t} + \epsilon_t \Rightarrow \hat{E}_t y_{t+1} = b_1 \rho_1 z_{1t} \tag{7}$$

where ϵ_t is a perceived noise variable. The actual law of motion is, then,

$$y_t = [\beta b_1 \rho_1 + \gamma_1] \, z_{1t} + \gamma_2 z_{2t} \tag{8}$$

The method of undetermined coefficients does not help pin down the value of b_1. Because z_{2t} is serially correlated and possibly correlated with z_{1t} the rational expectations value of $b_1 = \gamma_1/(1-\beta\rho_1)$ will not give the best forecast of y_t, in a least-squares sense. Instead, we need to compute b_1 from the linear projection of y_t onto the restricted space of variables, z_1.

Beliefs in a restricted perceptions equilibrium will satisfy the least-squares orthogonality condition that delivers the approximating model as the best linear model in its restricted class. In the current example,

$$Ez_{1t} \left(y_t - b_1 z_{1t} \right) = 0$$

Solving for b_1:

$$b_1 = \frac{Ey_t z_{1t}}{Ez_{1t}^2}$$

Alternatively, after plugging in the actual law of motion (8):

$$b_1 = [\beta b_1 \rho_1 + \gamma_1] + \gamma_2 \frac{Ez_{1t} z_{2t}}{Ez_{1t}^2} \equiv T(b_1)$$

A few comments:

- Notice that the least-squares projection of y_t on z_{1t} implies that the coefficient b_1 consists of two terms. The first is the actual coefficient in the data generating process (8). The second term is the omitted variable bias that emerges when z_{1t}, z_{2t} are correlated.

- The least-squares orthogonality condition implies that $b_1 = T(b_1)$, an expression similar to the previous section. However, there is no one-to-one mapping from the perceived law of motion to the actual law of motion; since the approximating model omits the z_{2t} term, it is under-parameterized. Instead, the T-map here is a "projected T-map." The interpretation of this T-map is that if agents held a perceived law of motion of the form (7), with fixed b_1, then with a sufficiently long sample, the regression of y_t on z_{1t} would estimate a coefficient $T(b_1)$.

Thus, a restricted perceptions equilibrium b_1^* is a fixed point to the T-map:

$$b_1^* = T(b_1^*)$$

Simple calculations show that

$$b_1^* = \frac{\gamma_1}{1-\beta\rho_1} + \frac{\gamma_2}{1-\beta\rho_1}\frac{\sigma_{12}}{\sigma_1^2}\frac{1-\rho_1^2}{1-\rho_1\rho_2}$$

The RPE value for the belief parameter b_1^* equals its REE value plus a term that captures the omitted variable bias. If the omitted variable is uncorrelated with the regressor, i.e., $\sigma_{12} = 0$, then there is no bias. The bias is increasing in the size of the omitted variable's direct effect, γ_2, the strength of the correlation (relative to the variance of z_{1t}), and the persistence of the omitted variable ρ_2.

Branch and Evans (2006) asked whether it is possible for under-parameterization/parsimony will lead to an equilibrium with heterogeneous expectations. The approach was to give agents a choice between all possible under-parameterized forecasting models. Agents could, for instance, select models based on their mean-squared forecast errors. Defining a misspecification equilibrium as a restricted perceptions equilibrium where agents only choose the best-performing models, heterogeneity will arise when each model delivers equivalent mean-squared errors. Branch and Evans (2006) call this intrinsic heterogeneity.

Let n denote the fraction of agents who forecast with (7). Then $1-n$ agents have expectations $\hat{E}^2 y_{t+1} = b_2\rho_2 z_{2t}$. The actual law of motion is

$$y_t = [n\beta b_1\rho_1 + \gamma_1]\, z_{1t} + [(1-n)\beta b_2\rho_2 + \gamma_2]\, z_{2t}$$

The indirect effect depends on the population distribution across the two models, i.e., n. The idea of a misspecification equilibrium is that n, b_1, b_2 are all equilibrium objects. So, there are now a pair of orthogonality conditions for $j = 1, 2$:

$$E z_{jt}\,(y_t - b_j z_{jt}) = 0$$

and a selection rule:

$$n = \begin{cases} 1 \text{ if } F(1) > 0 \\ 0 \text{ if } F(0) < 0 \\ \hat{n} \in (0,1) \text{ if } F(\hat{n}) = 0 \end{cases}$$

where

$$F(n) = E(y_t - b_2 z_{2t})^2 - E(y_t - b_1 z_{1t})^2$$

is the relative forecast accuracy of the z_1 model vis a vis the z_2 model.

It turns out, that when $0 < \beta < 1$, then in equilibrium $n = 0$, $n = 1$, or both.[4]. When there is negative feedback in the model with $-1 < \beta < 0$, there is the possibility of intrinsic heterogeneity. This case requires that the indirect effect is strong enough – i.e., sufficiently negative β – and that the two exogenous variables are correlated and sufficiently volatile.

4.2 Linear beliefs in a non-linear model

Applied forecasting models are typically linear. Macroeconomic models, particularly their DSGE variants, are solved as log-linear approximations around a steady state. However, most macroeconomic environments produce a non-linear relationship between equilibrium outcomes and state variables, including expectations.

Another restricted perceptions approach assumes that the economy's agents form expectations via a linear forecasting model and that the equilibrium law of motion is non-linear (c.f. Branch and McGough (2005); Hommes, Sorger, and Wagener (2013)). Consider a non-linear version of (1):

$$y_t = G(y^e_{t+1}) + v_t \tag{9}$$

where y is univariate, G is continuous and real-valued, and v_t is white noise with compact support. The earlier OLG example, with a more general preference structure, could lead to a pricing relationship like (9). The increasing returns model of Evans and Honkapohja (2001) is another example environment. Implicitly assumed in (9) is that agents hold point

[4] More concretely, in the case of multiple equilibria, there is a knife-edge equilibrium with $0 < n < 1$, but this equilibrium is unstable under learning.

expectations, i.e. $\hat{E}_t G(y_{t+1}) = G(y^e_{t+1})$. Among other technical assumptions, Branch and McGough (2005) impose that the function G is symmetric about a steady-state α.

Solving stochastic non-linear rational expectations models like (9) are complicated. Most researchers approximate the solution by solving for a solution to a first or second-order expansion around α. The restricted perceptions approach, on the other hand, assumes that the data generating process (9) is non-linear, but agents hold a linear perceived law of motion:

$$y_t = a + b\,(y_{t-1} - a) \Rightarrow y^e_{t+1} = a + b^2\,(y_{t-1} - a) \tag{10}$$

A restricted perceptions equilibrium determines the coefficients a, b as the optimal least-squares projection of y_t onto the space of linear models of the form (10). In this case, the coefficient for a will reflect the unconditional mean of y, and b will equal the unconditional first-order autocorrelation, where the former is taken with respect to the asymptotic distribution for y_t and computing the latter is taken from the joint asymptotic distribution over (y_t, y_{t-1}). Branch and McGough (2005) provide a general existence and uniqueness result for symmetric G.

As an example, consider the function $G(y) = F(y - \alpha) + \alpha$, where

$$F(y) = \begin{cases} y^\beta & \text{if } y \geq 0 \\ -\left(-y\right)^\beta & \text{else} \end{cases}$$

where $0 < \beta < 1$. The actual law of motion, then, is

$$y_t = G\left[\alpha + b^2\,(y_{t-1} - \alpha)\right] + v_t \tag{11}$$

Suppose that $\beta = 1/3, \alpha = a = 1.1, v_t \in [-1, 1]$. Branch and McGough (2005) compute that $b^* \approx 0.74$.

Figure 1 plots the restricted perceptions equilibrium and outcomes. The solid line is the function G plotted in phase space. The dashed line is the perceived law of motion (10). The circles represent the outcomes from a typical 1000-period simulation. Given linear beliefs (10) with $(a, b) = (1.1, 0.74)$ and resulting data generating process (11), leads to a stochastic process with a linear trend consistent with the linear beliefs. Although the agents mistakenly believe that the underlying process is linear, the actual non-linear process produces realizations that self-confirm those beliefs within the restricted perceptions equilibrium.

Figure 1: Restricted perceptions equilibrium: linear forecasting model in a non-linear world.

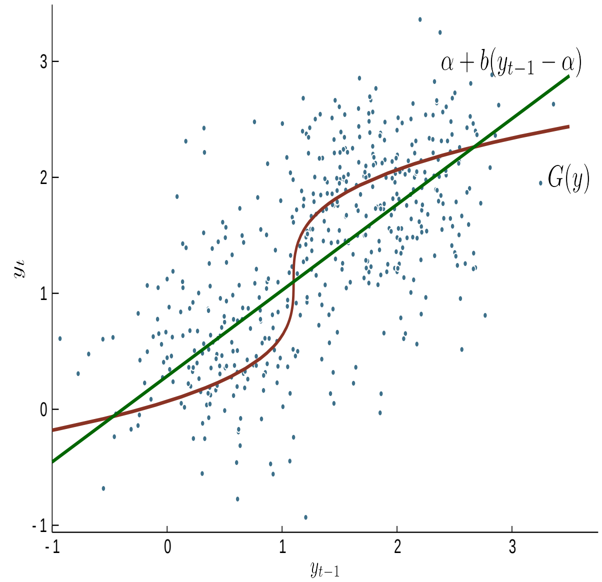

The applications section will consider restricted perceptions where linear beliefs in a non-linear world produce novel economic phenomena near rational sunspot equilibria. However, the restricted perceptions approach also provides a computationally and economically intuitive method for solving non-linear stochastic models.

4.3 Hidden state variables

In part, random shocks drive macroeconomic and asset-pricing models. A model is a simple formalization of a complex economic process disturbed by a significant number of exogenous forces. It is unlikely that people will observe all of those shocks. What are the consequences of hidden variables to a model?

4.3.1 Unobservable variables and rational expectations

With rational expectations, agents' beliefs come from expectations conditional on all available information. But, what if some information is not available? This is not just a theoretical curiousity: central banks set policy interest rates given their own views on the inflation-unemployment trade-off. These views and preferences of central bankers change over time. We do not observe them. Interestingly, whether hidden variables matter under rational expectations depends on the timing of when y_t is part of the information used by agents when forecasting $E_t y_{t+1}$. If we think economies work by people taking actions given their beliefs and then markets clear and determine y_t, then we would not expect even rational agents to observe y_t at the time they predict $E_t y_{t+1}$.

Now, let's return to the linear model (1) with $0 < \beta < 1$ and z_t univariate. The unique rational expectations equilibrium is

$$y_t = (1 - \beta\rho)^{-1} z_t$$

But, using the Wold decomposition on the assumed process for $z_t = \rho z_{t-1} + \varepsilon_t$, we have

$$z_t = (1 - \rho L)^{-1} \varepsilon_t$$

where L is the lag operator ($L z_t = z_{t-1}$), and so substituting for z_t

$$\begin{aligned} y_t &= (1 - \beta\rho)^{-1}(1 - \rho L)^{-1}\varepsilon_t \\ \Leftrightarrow (1 - \rho L) y_t &= (1 - \beta\rho)^{-1}\varepsilon_t \\ \Leftrightarrow y_t &= \rho y_{t-1} + (1 - \beta\rho)^{-1}\varepsilon_t \end{aligned}$$

So, an equivalent representation of the rational expectations equilibrium takes the form of a first-order autoregressive process (AR(1)).

The question asked by Marcet and Sargent (1989) is what does a rational expectations equilibrium look like if z_t is unobservable or a hidden variable to economic agents? The answer, it turns out, depends on whether y_t is contemporaneously observable.

Suppose that agents perceive the process to follow an AR(1)

$$y_t = b y_{t-1} + d\varepsilon_t$$

If agents observe y_t contemporaneously, then

$$E_t y_{t+1} = b y_t$$

and the actual law of motion becomes

$$y_t = \beta b y_t + \gamma z_t$$

or,

$$y_t = by_{t-1} + \frac{\gamma}{1-\beta b}\varepsilon_t$$

So, the rational expectations equilibrium takes the usual form when y_t is observable and $b = \rho$. After all, when $b = \rho$ then our actual law of motion is identical to the rewritten version of the unique equilibrium.

Now suppose that y_t is observed with a lag: $E_{t-1}y_{t+1} = b^2 y_{t-1}$. In this case, the rational expectations equilibrium follows an AR(∞), it depends on an infinite number of lags. To see this, note that with the AR(1) perceived law of motion and where y_t is not currently observable, the actual law of motion is

$$y_t = \beta b^2 y_{t-1} + \gamma z_t$$

or again substituting with $z_t = (1-\rho L)^{-1}\varepsilon_t$,

$$(1-\rho L)\left(1-\beta b^2 L\right) y_t = \gamma \varepsilon_t$$

When agents cannot observe z_t or contemporaneous y_t and guess an AR(1) equilibrium, the actual data-generating process is an AR(2), with 2 lags of y_t. More generally, if agents have an AR(p) perceived law of motion, then

$$E_{t-1}y_{t+1} = b_1\left(b_1 y_{t-1} + \cdots + b_p y_{t-p}\right) + b_2 y_{t-1} + \cdots + b_p y_{t-p+1}$$

and the true data-generating process is an AR($p+1$):

$$(1-\rho L)\left[1-\beta\sum_{j=1}^{p}\left(b_1 b_j + b_{j+1}\right)L^j\right] y_t = \gamma\varepsilon_t$$

It follows that with some components of z_t hidden and y_t observed with a lag, the rational expectations equilibrium is an AR(∞). In this case, if the agents had an infinitely long history of y_t they could reverse engineer the equation with infinite lags of y_t to filter and recover the hidden z_t shocks.

4.3.2 Unobservable variables and restricted perceptions

Of course, it is not practical to estimate the coefficients of an AR(∞) with finite data histories. Hommes and Zhu (2014) find a restricted perceptions equilibrium with an AR(1) forecasting equation. Following the generalization in Branch, McGough, and Zhu (2022), this section traces the implications of an under-parameterized AR(1) perceived law of motion.

Assume that agents hold beliefs consistent with an AR(1) perceived law of motion:

$$y_t = by_{t-1} + \epsilon_t \Rightarrow \hat{E}_t y_{t+1} = b^2 y_{t-1} \tag{12}$$

Then the actual law of motion is

$$y_t = \beta b^2 y_{t-1} + \gamma z_t$$

In a restricted perceptions equilibrium, the coefficient b needs to be sure that the approximating model delivers the best statistical fit in an ordinary least-squares sense. In other words, b will satisfy the least-squares orthogonality condition:

$$Ey_{t-1}\left(y_t - by_{t-1}\right) = 0$$

Or solving for b,

$$b = \frac{Ey_t y_{t-1}}{Ey_{t-1}^2} \equiv T(b)$$

The right hand side is the first-order autocorrelation coefficient for y_t. Notice that because the model is self-referential, the right-hand side depends on b: if agents believe that y_t has an autocorrelation coefficient b, then $\hat{E}_t y_{t+1} = b^2 y_{t-1}$, and the resulting actual autocorrelation coefficient is b. The restricted perceptions equilibrium is a weaker version of rational expectations: rather than trying to match the full distribution, agents' beliefs just need to be consistent with this one empirical moment in the data.

To calculate the T-map we need to compute various moments based off of the actual law of motion:

$$\begin{aligned} Ey_t y_{t-1} &= E\left(\beta b^2 y_{t-1} + \gamma z_t\right) y_{t-1} \\ &= \beta b^2 Ey^2 + \gamma^2 Ez^2 + 2\beta b^2 \gamma \rho Eyz \end{aligned}$$

using that $Ez_t y_{t-1} = E(\rho z_{t-1} + \varepsilon_t) y_{t-1} = \rho Eyz$. Then similarly calculating Eyz, and Ey^2, straightforward computations show that the T-map is

$$b \to \frac{\beta b^2 + \rho}{1 + \beta b^2 \rho} = T(b)$$

Hommes and Zhu (2014) and Branch, McGough, and Zhu (2022) prove the existence of an RPE $\hat{b}$ that is a fixed point to the T-map. This equilibrium value for $\hat{b}$ is a complicated expression of β and ρ.

Interestingly, the restricted perceptions equilibrium identified here may not be unique, and other equilibria can depend on sunspots. The applications portion of this chapter studies the complete set of solutions.

5. Applications

This section turns to a few novel applications that arise from the restricted perceptions approach. These are non-Ricardian beliefs, sunspot equilibria that are stable under learning, and random-walk beliefs that lead to asset price bubbles and inflation scares.

5.1 (Non-)Ricardian beliefs

Several adaptive learning papers develop non-Ricardian beliefs in New Keynesian type models (c.f. Evans, Honkapohja, and Mitra (2009); Eusepi and Preston (2018); Woodford (2013)). A boundedly rational consumption function combines the Euler equations with a household intertemporal budget constraint without assuming the household correctly understands the government's budget constraint and debt solvency. Instead, the household also must forecast the government's debt and primary surplus evolution. In these environments, the government's level of debt and surplus are state variables. Branch and Gasteiger (2022) study the consequences of under-parameterization in forecasting in this environment without directly imposing Ricardian beliefs. The under-parameterization here reflects that most applied analyses of the effects of fiscal policy typically only include a single fiscal variable. Here, though, it is a convenient formalization for how non-Ricardian beliefs could emerge endogenously.

Woodford (2013) introduces a special case of a purely real New Keynesian model without imposing Ricardian beliefs. The key equations are:

$$b_{t+1} = \beta^{-1}(b_t - s_t) \tag{13}$$

$$y_t = v_t + (1-\beta) b_t \tag{14}$$

$$v_t = (1-\beta)(b_{t+1} - b_t) + \hat{E}_t v_{t+1} \tag{15}$$

$$s_t = \phi_b b_t + z_t \tag{16}$$

The variables in the linearized equations denote log deviations from steady-state. Equation (13) is the government's flow budget constraint, and it relates the stock of one-period government debt issued at time t, b_{t+1}, to the difference between the beginning of period debt and the primary surplus s_t. The parameter $0 < \beta < 1$ is the discount rate and β^{-1} is equal to the steady-state gross real interest rate. The second equation (14) is the consumption function after imposing the goods market clearing condition that $c_t = y_t$. The variable v_t is a continuation value that reflects the annuitized value of future tax liabilities and returns on debt holdings that enter the household consumption function. So aggregate output depends on this forward-looking variable as well as the existing stock of debt. Consumption is, therefore, has a standard permanent income form, here written recursively. Equation (15) provides the recursion that determines the continuation value v_t. The question of Ricardian equivalence is whether the household's expectations $\hat{E}_t v_{t+1}$ correctly anticipate future surpluses and debt issuances. If so, then y_t will not depend on government debt b_t, and Ricardian equivalence holds. The final equation is the fiscal rule with the policy coefficient $1 - \beta < \phi_b < 1$ and white noise policy shocks z_t.

Since b_{t+1}, depending on b_t, s_t, is the relevant state variable, the full information equilibrium will depend on b_{t+1} if it is observable or separately on b_t, s_t. Suppose that agents do not know b_{t+1} and the flow government budget constraint. Moreover, the agents prefer parsimonious models and include a single fiscal variable in their forecast equation. Then there are two potential forecasting models:

$$v_t = \psi^s s_{t-1} + \epsilon_t \tag{17}$$

$$v_t = \psi^b b_{t-1} + \epsilon_t \tag{18}$$

The first model (17) depends only on the flow budget surplus, while the model (18) depends on the previous period's government debt level. Loosely speaking, n parameterizes the extent of non-Ricardian beliefs. A restricted perceptions equilibrium jointly pins down the belief parameters in (17)-(18) and the aggregate variables (13)-(16).

Rather than imposing that all agents forecast with the same model, let

$$\hat{E}_t v_{t+1} = n\psi^s s_t + (1 - n)\psi^b b_t$$

where n is the fraction of households that forecast with the surplus model (17). In a restricted perceptions equilibrium, there are a pair of least-squares orthogonality conditions:

$$
\begin{aligned}
Es_{t-1}\left[v_t - \psi^s s_{t-1}\right] &= 0 \\
Eb_{t-1}\left[v_t - \psi^b b_{t-1}\right] &= 0
\end{aligned}
$$

Branch and Gasteiger (2022) prove the existence, given n, of a unique RPE.

Consider the special case of $n = 0$; all households forecast with the debt model. Then it turns out the RPE is

$$
\begin{aligned}
y_t &= -\left(\beta^{-1} - 1\right) z_t \\
\psi^b &= -\left(\beta^{-1} - 1\right)\left(1 - \phi_b\right)
\end{aligned}
$$

Notice that aggregate output does not depend on the debt stock b_t. It does depend on the innovation to the primary surplus z_t, but since $\beta \approx 1$, the effect is negligible. Branch and Gasteiger (2022) call this a form of weak Ricardian equivalence.

Now let $n = 1$; all households forecast with the surplus model. Then Woodford (2013) finds that

$$
\begin{aligned}
y_t &= \left[\frac{(1-\beta)(1+\beta-\phi_b)}{\beta(1+\beta)+\phi_b}\right] b_t - \left[\frac{\beta\left(1-\beta^{-1}\right)}{\beta(1+\beta)+\phi_b}\right] z_t \\
\phi^s &= -\frac{\beta^{-1}(1-\beta)\left(1-\beta^2-\phi_b\right)}{\beta(1+\beta)+\phi_b} < \beta^{-1} - 1
\end{aligned}
$$

Generalizing, Branch and Gasteiger (2022) show that the unique RPE, given n is of the form

$$y_t = \xi_1(n) b_t + \xi_2(n) z_t$$

with

$$\xi_1(n) \neq 0 \Leftrightarrow n > 0$$

So Ricardian equivalence is fragile. Only if every agent forecasts with the stock of debt will they correctly understand the consequences of fiscal policy. For any $n > 0$ – including $n \to 0$ – then neither type of agent holds Ricardian beliefs. In this sense, Ricardian equivalence is fragile and occurs in a highly restrictive self-confirming equilibrium.

5.2 Sunspots

Section 3 showed how belief-driven fluctuations could arise through indeterminacy and sunspot equilibria. There are two drawbacks to rational sunspots as a model of self-confirming fluctuations: first, often, the sunspot equilibria are unstable under learning; second, the indeterminacy regions in the DSGE model often coincide with empirically unrealistic parameterizations. The restricted perceptions approach, however, can overcome both limitations of rational sunspots.

5.2.1 Statistical sunspots

Branch, McGough, and Zhu (2022) show that with unobservable variables, the restricted perceptions approach can generate sunspot equilibria in determinate models that are stable under learning.

The approach is to extend (12) to include dependence on an extrinsic shock:

$$y_t = by_{t-1} + d\xi_t \Leftrightarrow \hat{E}_t y_{t+1} = b^2 y_{t-1} + d\left(b + \phi\right)\xi_t \tag{19}$$

where

$$\xi_t = \phi\xi_{t-1} + v_t$$

is an extrinsic noise term uncorrelated with the hidden variable z_t. To ease exposition, assume that $\varepsilon_t \sim N\left(0, \sigma_\varepsilon^2\right), v_t \sim N\left(0, \sigma_v^2\right), \sigma_{v,\varepsilon} = 0$. If $d \neq 0$, then the economy depends on this extrinsic noise. The question is whether ξ_t can matter in a self-confirming way, and we can interpret it as a sunspot. Branch, McGough, and Zhu (2022) study how ξ_t can matter in a restricted perceptions equilibrium.

The least-squares orthogonality condition is

$$E\left(y_t - by_{t-1} - d\xi_t\right)\left(y_{t-1}, \xi_t\right)' = 0$$

The sunspot will matter for y_t when there is a non-zero correlation between y_t and ξ_t:

$$d = \left(1 - b\phi\right)\frac{Ey_t\xi_t}{E\xi_t^2}$$

But, y_t depends in turn on (b, d). For this reason, Branch, McGough, and Zhu (2022) call ξ_t a statistical sunspot; its presence arises because of a self-confirming statistical correlation.

Proceeding as in Section 4.3.2, the $d-$component of the T-map is

$$d \to \frac{d\beta\,(b+\phi)\,(1-b\phi)}{1-\beta b^2\phi}$$

There are two fixed points to the T_d component of the T-map: $d^* = 0$ and

$$b^* = \frac{1-\beta\phi}{\beta\,(1-\phi^2)}$$

When $d^* = 0$, the PLM (19) is identical to (12) and the earlier RPE results. When $d^* \neq 0$, then the RPE value can be found as a solution to the T_b component (see Branch, McGough, and Zhu (2022) for details):

$$(d^*)^2 = \xi(b^*, \beta, \rho, \phi)$$

Branch, McGough, and Zhu (2022) show that for fixed ϕ – the serial correlation of the sunspot variable – that

1. There exists a unique RPE with $(b, d) = (\hat{b}, 0)$, and $\hat{b}$ is the equilibrium identified by Hommes and Zhu (2014).

2. There exist threshold values $\tilde{\beta}(\phi), \tilde{\rho}(\beta, \phi)$ so that sunspot RPE $(b, d) = (b^*, \pm d^*)$ exist $\Leftrightarrow \tilde{\beta} < \beta < 1$ and $\tilde{\rho} < \rho < 1$.

3. When sunspot RPE exists, the sunspot RPE is E-stable, while the fundamentals RPE with $d = 0$ is E-unstable.

Figure 2 illustrates the results. The solid lines correspond to the fixed points of the two T-map components. The vector field indicates the E-stability dynamics. An RPE occurs where the contours intersect. There are three RPE. There is the fundamental RPE with $d = 0$. There are also two symmetric sunspot RPE. Notice, in particular, that the sunspot RPE is E-stable. Thus, with hidden variables in an environment with a unique rational expectations equilibrium, we expect the economy to exhibit dependence on self-confirming sunspots.

5.2.2 Near rational sunspots

Evans and McGough (2020b) find existence of E-stable near rational sunspot equilibria in non-linear models where the steady-state is locally indeterminate. Their approach, like

Figure 2: Statistical sunspots.

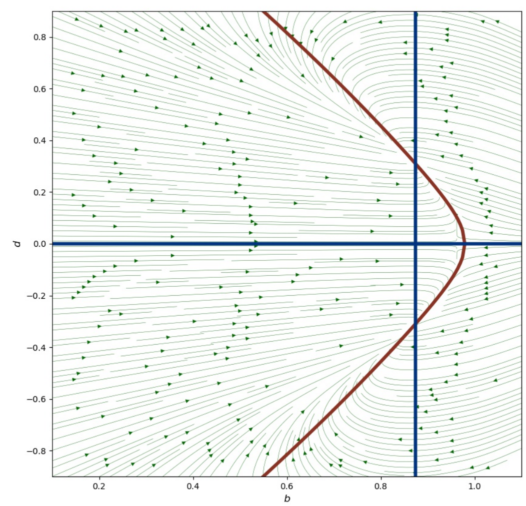

section 4.2, assumes that agents forecast with a linear model while the data-generating process is non-linear.

Suppose, as in Evans and McGough (2020b) that

$$y_t = F\left(y^e_{t+1}\right)$$

with

$$F(y) = \theta y + \mu y^3$$

Agents have a linear perceived law of motion:

$$y_t = a + d\eta_t$$

where η_t is a serially correlated process

$$\eta_t = \lambda\eta_{t-1} + e_t$$

The coefficient λ satisfies a "resonant frequency condition" so that a sunspot equilibrium in the linearized version corresponds to the sunspot equilibrium found in Section 3

Now in a restricted perceptions equilibrium, the coefficients (a, d) are found from the least-squares projection of the actual non-linear law of motion onto the space spanned by

$(1, \eta_t)$. The main results in Evans and McGough (2020b) can be illustrated in their simple cubic example with $\theta = -5$ and $\mu = 2$. The $\theta < 0$ is equivalent to the negative feedback case in (1). Figure 3 plots the existence and E-stability of near rational sunspots.

Figure 3: Near rational sunspot equilibria.

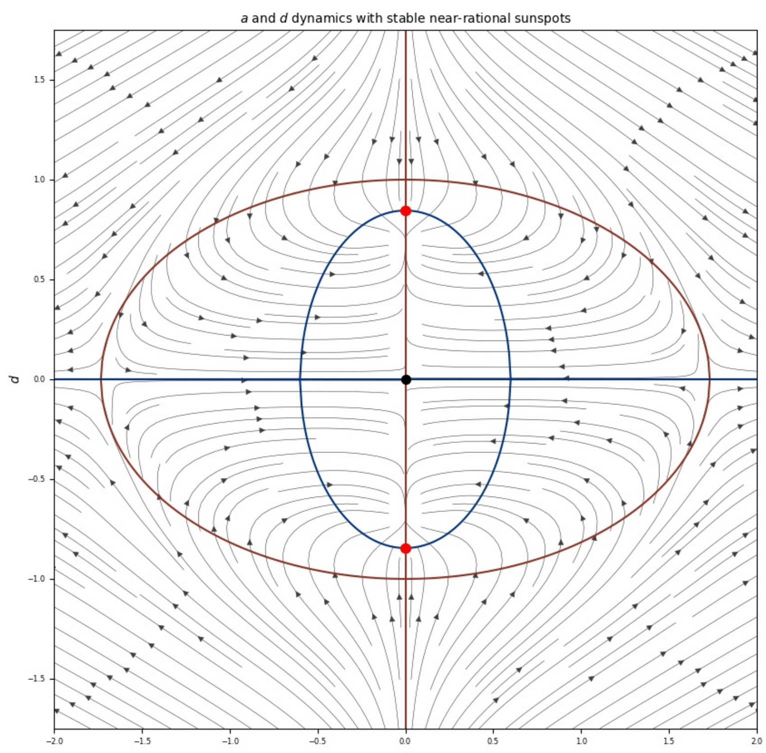

Figure 3 illustrates the determination of near rational sunspots. The outer circle and vertical line correspond to fixed points to the T_a component of the T-map: given a value for the coefficient d on the sunspot, these are all the optimal values for a from the least-squares orthogonality condition. The inner circle and horizontal axis are the fixed points to T_d: given a value for the intercept (or, mean) a, these are all the optimal values for the coefficient d. The intersection of the contours are RPE: $(\bar{a}, \bar{d}) = \left(T_a(\bar{a}, \bar{d}), T_d(\bar{a}, \bar{d})\right)$. A near rational sunspot equilibrium simply requires that $d \neq 0$. Therefore, intersections in Figure 3 with $d \neq 0$ are near rational sunspot equilibria. In Figure 3 these near-rational sunspot equilibria are denoted intersections along the vertical axis and marked with circles. There are two symmetric near-rational sunspot equilibria.

Towards the beginning of this chapter, we saw that sunspot equilibria are generally unstable in linear models. Figure 3 also illustrates the E-stability. The vector field lines denote, for initial coefficient values, the expected learning path. The steady state is saddle-path stable: if agents believe there are no sunspots, $d = 0$, then their adaptive learning process will bring them to an equilibrium without sunspots $(\bar{a}, \bar{d}) = (0, 0)$. However, it is the near-rational sunspots that are stable under learning. All of the paths (besides the aforementioned saddle path) point towards the near-rational sunspots.

Forecasting with a linear model – as most economists do – in a non-linear setting – as most models are – produces learnable sunspots that would not arise in a fully rational setting.

5.3 Random-walk beliefs

In the presence of hidden variables, the restricted perceptions approach finds equilibria that are more serially correlated than under rational expectations. Moreover, these serially correlated and volatile equilibria are stable under learning. In other settings, restricted perceptions equilibria can exist that are self-confirming but not stable under learning. Nonetheless, learning dynamics can still be drawn toward the RPE and exert influence over dynamics for a finite period. In linear self-referential models, these RPE take the form of random-walk beliefs.

Continuing with the general model (1), suppose that z_t is univariate white noise and that agents perceive that y follows a random walk without drift:

$$y_t = y_{t-1} + \epsilon_t \Rightarrow \hat{E}_t y_{t+1} = y_{t-1}$$

The actual law of motion, then, is

$$y_t = \alpha + \beta y_{t-1} + \gamma z_t$$

It is convenient instead to rewrite the perceived and actual laws of motion as moving average processes (MA). Notice that the perceived law of motion is also

$$(1 - L)y_t = \epsilon_t \Rightarrow y_t = (1 - L)^{-1}\epsilon_t \equiv g(L)\epsilon_t$$

and $g(L)$ is a lag polynomial with an infinite number of lags of ϵ_t.[5] Similarly, from the

[5]The Wold decomposition rewrites an AR(1) as an MA(∞). Take $z_t = \rho z_{t-1} + \varepsilon_t = \rho\left(\rho z_{t-2} + \varepsilon_{t-1}\right) +$

actual law of motion

$$\begin{aligned}(1-\beta L)y_t &= \alpha+\gamma z_t \\ \Leftrightarrow y_t &= \frac{\alpha}{1-\beta}+\gamma(1-\beta L)^{-1}z_t \equiv \mu+f(L)z_t\end{aligned}$$

Putting it altogether, we have infinitely-ordered moving average MA(∞) processes for both perceived and actual laws of motion:

$$\text{PLM: } y_t = g(L)\epsilon_t \tag{20}$$

$$\text{ALM: } y_t = \mu+f(L)z_t \tag{21}$$

where $g(L)=(1-L)^{-1}, \mu=\alpha/(1-\beta), f(L)=\gamma(1-\beta L)^{-1}$. In (21), the unconditional mean of y_t, is $\alpha/(1-\beta)$, the same value as in the rational expectations equilibrium. However, when z_t is iid (as in this example), y_t is not serially correlated. Here the random-walk beliefs induce serial correlation that would not exist under rational expectations.

For large values of $0<\beta<1$, the serial correlation induced by random-walk beliefs is nearly self-fulfilling. In fact, as $\beta\to 1$, the moving average structure is identical under the perceived law of motion (20) and the actual law of motion (21). In a similar setting, Sargent (1999) notes that random-walk beliefs can track constants well. In the rational expectations equilibrium, y_t equals μ plus iid innovations. The perceived law of motion (20) does not have a constant and uses higher-order moments to track low-frequency movements in the unconditional mean. This latter point provides key intuition for why random-walk beliefs, in models with strong expectational feedback, can be expected to arise under learning.

Section 3 showed that the iid rational expectations equilibrium is E-stable. It is natural to ask, what kind of real-time learning dynamic could generate nearly self-fulfilling random-walk beliefs? The answer is a variant of recursive least-squares called "constant gain learning."

Suppose that agents have an AR(1) perceived law of motion:

$$y_t=a+by_{t-1}+\epsilon_t \Rightarrow \hat{E}_t y_{t+1}=a(1+b)+b^2y_{t-1} \tag{22}$$

Notice that this PLM nests both the rational expectations equilibrium and random-walk beliefs. Let $\theta_t=(a_t,b_t)'$, $x_t=(1,y_{t-1})'$ and R_t be the sample estimate of the unconditional

ε_t, and so on. Then $z_t=\sum_{j=0}^{\infty}\rho^j z_{t-j}=z_t\left[1+\rho L+\rho L^2+\cdots\right]=(1-\rho L)^{-1}z_t$.

covariance matrix $Ex_t x_t'$. Then recursive estimates of a_t, b_t are updated via the stochastic recursive algorithm:

$$\begin{aligned}\theta_t &= \theta_{t-1} + \phi_t R_t^{-1} x_t \left(y_t - \theta_{t-1}' x_t\right) \\ R_t &= R_{t-1} + \phi_t \left(x_t x_t' - R_{t-1}\right)\end{aligned}$$

At $t = 0$, agents have priors over a_0, b_0. They form expectations via (22), shocks occur, and new data are determined by (1). Then θ_t, R_t are updated, and the process repeats. The variable ϕ_t is called a gain sequence. Recursive least-squares arises when $\phi_t = 1/t$. In this case, the learning estimates place equal weight on all data points. When $\phi_t = \phi$, it is called constant gain learning. With a constant gain, the learning algorithm places geometrically declining weights on past data. Constant gain learning is advisable when agents are concerned about structural change of an unknown form.

The E-stability principle says that the recursive least-squares estimates $\theta_t \to \theta^*$, the REE values for (a, b), with probability 1 as $t \to \infty$. Under constant gain learning, however, the estimates θ_t do not settle down to the REE values. Because of the time-invariant gain, it turns out that θ_t can converge in distribution: for large t and as $\phi \to 0$, $\theta_t \sim N(\mu, \phi V)$. So constant gain learning beliefs, for large samples, are distributed around the rational expectations equilibrium with a variance proportional to the constant gain ϕ.

Even more insightful, is a result in Evans and Honkapohja (2001), that as $\phi \to 0$ and t large, the random learning path $\theta(\tau)$ converges weakly to the solution $\tilde{\theta}(\tau, \theta_0)$ – for any θ_0 in a neighborhood of the rational expectations equilibrium – to the mean-dynamics:

$$\begin{aligned}\dot{\theta} &= R^{-1} M(\theta) \left(T(\theta) - \theta\right) \\ \dot{R} &= M(\theta) - R\end{aligned}$$

and where τ maps the discrete time sequence for θ_t into a continuous time path $\theta(\tau)$. The mean-dynamics o.d.e. are found via a continuous-time interpolation of the recursive least-squares algorithm and appealing to a law of large numbers. The result tells us that the solution to the mean-dynamics o.d.e. delivers the expected path for the real-time learning estimates θ_t following a sequence of shocks that drive learning to θ_0. Those sequences of shocks that "initialize" the mean dynamics are called escape dynamics; see Williams (2019).

A numerical example demonstrates the possibility for learning to be drawn toward random-walk beliefs. Set $\beta = 0.95, \alpha = 0.1, \sigma_z = 0.1$. For $\theta_0 \approx (2, 0)$, E-stability tells us that learning will converge to the rational expectations equilibrium $\theta^* = (2, 0)$. It is during the transition path that random-walk beliefs can arise. Suppose that $\theta_0 = (2, 0.5)$. Figure 4 illustrates the mean dynamic paths for $a(\tau), b(\tau)$. Notice that eventually, learning converges to the rational expectations equilibrium. The transition path, however, features $(a, b) \approx (0, 1)$ for a finite stretch of time: these are random-walk beliefs as in (20). Eventually, the E-stability of the rational expectations equilibrium restores beliefs to the rational expectations equilibrium values.

Figure 4: Mean dynamics: random walk beliefs

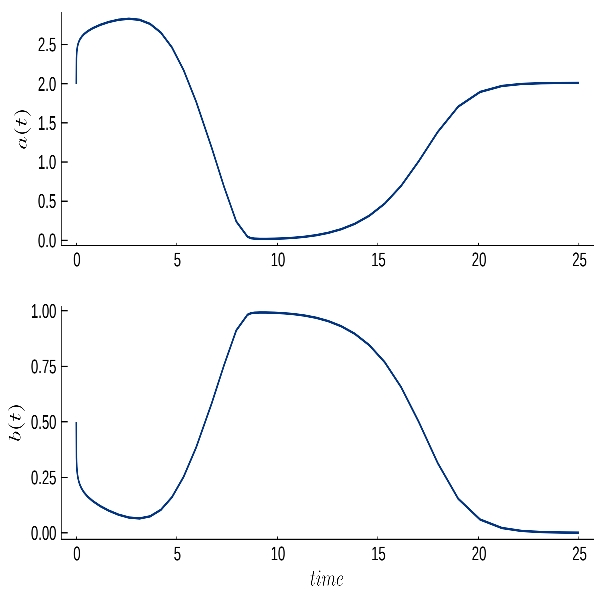

What drives beliefs toward random-walk beliefs? It is the same intuition described above that random-walk beliefs can introduce self-fulfilling serial correlation and approximate well low-frequency drift. Imagine a positive sequence of shocks to z_t. The agents' econometric model detects the serial correlation when regressing y_t on a constant and a lag. The self-referential nature of the model (1) then induces more serial correlation, which the agents' model again detects. The random-walk beliefs are nearly self-confirming. Eventually, a new sequence of shocks counters these beliefs, and agents learn the actual process again. This type of sequence of shocks – Sargent calls it the "most likely unlikely" sequence – triggers an escape from the rational expectations equilibrium. Then the mean dynamics show a temporary form of restricted perceptions.

Random-walk beliefs that arise endogenously have novel applications that arise by creating additional persistence and volatility:

1. Branch and Evans (2011a) incorporated adaptive learning into a mean-variance asset-pricing that takes the form of (1). The shock z_t is a shock to share supply, i.e., asset float. In that framework, the right sequence of shocks to share supply – say because of lock-up expirations – led to a constant gain learning process drawn towards random-walk beliefs. At that moment, agents mistakenly perceive all stock price innovations as permanent, increasing their demand and increasing prices further. The result is self-fulfilling bubbles or crashes in stock prices.

2. Branch and Evans (2017) studied constant gain learning in various New Keynesian monetary models with non-zero long-run inflation targets. If agents imperfectly understand the central bank's inflation target, then to forecast inflation, the agents need estimates of both the conditional mean and persistence of inflation. That gives rise to the possibility of endogenous random-walk learning dynamics. The key result in that study is that higher long-run inflation targets increase the likelihood of random-walk beliefs. Thus, increasing the target can trigger random-walk beliefs and an inflation scare. Once at the higher target, the emergence of random-walk beliefs can lead to an inflation scare or even a collapse to the zero lower bound. Lower values for the inflation target produce stable learning dynamics.

6. Literature Review and Conclusion

Early contributions to the restricted perceptions approach include Marcet and Sargent (1989), Evans, Honkapohja, and Sargent (1993), and Marcet and Sargent (1995). The papers Marcet and Sargent (1989, 1995) explore hidden variables and their implications for expectations. Evans, Honkapohja, and Sargent (1993) assume agents mistakenly fit an AR(1) regression model to data in an economy with complex, deterministic dynamics.

Evans and Honkapohja (2001) offered the first example of an under-forecasting model in a cobweb model and a forward-looking model like (1) that includes a lag. The E-stability properties of the latter are different with restricted perceptions than with rational expectations. Branch and Evans (2006) extended the under-parameterization framework to allow the agents to select their model optimally in a random-utility setting, generalizing Brock and Hommes (1997) to a stochastic environment. The key result of that paper is that heterogeneous expectations can arise endogenously. Multiple equilibria can arise in a monetary model, Branch and Evans (2007) similar to the Ricardian beliefs example discussed here, that generate time-varying inflation volatility. Similarly, Branch and Evans

(2009) find that endogenous choice of under-parameterized models generates empirically realistic shifting means and variances in stock returns. Other approaches to predictor selection appear in Markiewicz (2012) and Cho and Kasa (2015).

The restricted perceptions approach sometimes goes under alternative names. For instance, the consistent expectations equilibrium of Hommes and Sorger (1997) asks forecasts to be consistent with a finite number of sample autocorrelations. Hommes and Zhu (2014) describe a first-order consistent expectations equilibrium that they call a behavioral learning equilibrium. Finally, Cho and Kasa (2015) have an under-parameterization example that they label a self-confirming equilibrium. The restricted perceptions approach is encompassing, and these alternative definitions are refinements.

Recent work incorporates restricted perceptions into empirically realistic business cycle and monetary models. The key distinction of these classes of models is the assumption of infinitely-lived agents that solve complicated intertemporal optimization problems. Evans, Evans, and McGough (2022) present a general existence result for restricted perceptions equilibria when they make specific bounded optimality assumptions. Branch and McGough (2011) find that restricted perceptions and heterogeneous beliefs can amplify volatility in a real business cycle model. Branch and Evans (2011b) find hysteresis effects in a New Keynesian model with agent choice over under-parameterized models.

Besides the non-linear models referenced earlier, other papers have developed the restricted perceptions approach in non-linear models. Evans and McGough (2020a) study learning with linear under-parameterized forecasts in a non-linear cobweb model. Similarly, Shin (2020) develops an asset-pricing model with costly market participation and under-parameterized forecasting equations.

References

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

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

——— (2009): “Asset Return Dynamics and Learning,” Review of Financial Studies, forthcoming.

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

Branch, W. A., and G. W. Evans (2011b): "Monetary Policy and Heterogeneous Expectations," Economic Theory, this volume.

Branch, W. A., and G. W. Evans (2017): "Unstable Inflation Targets," Journal of Money Credit and Banking.

Branch, W. A., and E. Gasteiger (2022): "Endogenously (Non-)Ricardian Beliefs," Discussion paper.

Branch, W. A., and B. McGough (2005): "Consistent Expectations and Misspecification in Stochastic Non-linear Economies," Journal of Economic Dynamics and Control, 29, 659–676.

Branch, W. A., and B. McGough (2011): "Business Cycle Amplification with Heterogeneous Expectations," Economic Theory, this volume.

Branch, W. A., B. McGough, and M. Zhu (2022): "Statistical Sunspots," Theoretical Economics, 17(1), 291–329.

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

Cho, I.-K., and K. Kasa (2015): "Learning and Model Validation," Review of Economic Studies, 82, 45–82.

De Long, B., A. Shleifer, L. Summers, and R. Waldmann (1990): "Noise Trader Risk in Financial Markets," Journal of Political Economy, 98(4), 703–738.

Eusepi, S., and B. Preston (2018): "Fiscal Foundations of Inflation: Imperfect Knowledge," American Economic Review, 108(9), 2551–2589.

Evans, D., G. W. Evans, and B. McGough (2022): "The RPEs of RBCs and other DSGEs," Journal of Economic Dynamics and Control, forthcoming.

Evans, G. W., and R. Guesnerie (1999): "Coordination on Saddle Path Solutions: the Eductive Viewpoint. 1 – Linear Univariate Models," Working paper.

——— (2000): "Coordination on Saddle Path Solutions: the Eductive Viewpoint. 2 – Linear Multivariate Models," Working paper.

Evans, G. W., and S. Honkapohja (2001): Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, New Jersey.

Evans, G. W., S. Honkapohja, and K. Mitra (2009): "Anticipated Fiscal Policy and Learning," Journal of Monetary Economics, 56, 930–953.

Evans, G. W., S. Honkapohja, and T. J. Sargent (1993): "On the Preservation of Deterministic Cycles When Some Agents Perceive Them to be Random Fluctuations," Journal of Economic Dynamics and Control, 17, 705–721.

Evans, G. W., and B. McGough (2014): "Learning to Optimize," Working paper.

——— (2020a): "Equilibrium Stability in a Nonlinear Cobweb Model," Economics Letters, 193.

——— (2020b): "Stable Near-rational Sunspot Equilibria," Journal of Economic Theory, 186(3).

Grandmont, J.-M. (1977): "Temporary General Equilibrium Theory," Econometrica, 45(3), 535–572.

——— (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.

Hommes, C. H., and G. Sorger (1997): "Consistent Expectations Equilibria," Macroeconomic Dynamics, 2, 287–321.

Hommes, C. H., G. Sorger, and F. Wagener (2013): "Consistency of linear forecasts in a nonlinear stochastic economy," in Global analysis of dynamic models in economics and finance: essays in honour of Laura Gardini, ed. by G. Bischi, C. Chiarella, and I. Sushko. Springer.

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

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

——— (1995): Speed of Convergence of Recursive Least Squares: Learning with Autoregressive Moving-Average Perceptionschap. 6, pp. 179–215. Basil Blackwell, Oxford.

Markiewicz, A. (2012): "Model Uncertainty and Exchange Rate Volatility," International Economic Review, 53(3), 815–844.

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

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

Shin, M. (2020): "Endogenous Participation, Risk, and Learning in the Stock Market," Macroeconomic Dynamics, pp. 1–33.

White, H. (1994): Estimation, Inference and Specification Analysis, Econometric Society Monograph. Cambridge University Press, Cambridge UK.

Williams, N. (2019): "Escape Dynamics in Learning Modelsa," Review of Economic Studies, 86(2), 882–912.

Woodford, M. (2013): "Macroeconomic Analysis without the Rational Expectations Hypothesis," in Annual Review of Economics.