

The RPEs of RBCs and other DSGEs*

David Evans
University of Oregon

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

Bruce McGough
University of Oregon

November 4th, 2021

Abstract

In a broad class of non-linear representative agent models, represented by a system of difference equations, we replace rational expectations with linear forecast models conditioning on a predetermined set of regressors. Within this framework, a *restricted perceptions equilibrium* (RPE) corresponds to a forecast rule that is optimal among that class of models. Local uniqueness of a rational expectations equilibrium (REE) near the non-stochastic steady state is shown to guarantee existence and uniqueness of an RPE local to that steady state. This RPE is E-stable provided the REE is E-stable. A benchmark RBC model with government spending shocks is used to illustrate the theoretical results.

JEL Classifications: D83; D84; E24

Key Words: Real business cycle model, adaptive learning, E-stability, restricted perceptions

*We thank participants at the summer 2021 conference on *Expectations in dynamic macroeconomic models* held at the Czech National Bank in Prague, and participants in University of Oregon's *Macro group*.

1 Introduction

Dynamic stochastic general equilibrium (DSGE) models comprise the scaffolding of modern macroeconomics, and here the imperialism of the rational expectations hypothesis manifests quite clearly: these models, *unless otherwise specified*, take rational expectations equilibrium (REE) as their solution concept. In particular, agents are assumed to make forecasts optimally using the economy's endogenous conditional distributions, and to make decisions optimally conditional on their forecasts; and then, magically, these decisions aggregate to form realizations of the endogenous variables consistent with the endogenous distributions that formed the basis for the agents' decisions.

The mental gymnastics required of economic agents in rational expectations (RE) models doesn't square with experience: even professional forecasters use estimated models, with parsimonious linear models often competing with, if not out-classing, state-of-the-art non-linear models. This same experience suggests a natural alternative: assume agents populating a DSGE economy use simple linear forecast rules to form expectations, but otherwise act in a manner analogous to their rational counterparts. In particular, these agents must make decisions taking into account constraints and prices induced by a non-linear world.

The agent's use of a given linear forecast model determines a data generating process (DGP) for the economy, and by projecting onto the span of the agent's regressors the optimal linear forecast model is determined. A natural solution concept for this economy is a *restricted perceptions equilibrium* (RPE), which requires that agents use forecast models that are optimal within a given class. Like an REE, an RPE should be viewed as a Nash equilibrium concept: the associated forecast model is optimal only if (almost) all other agents in the economy are also using it.

The "restricted perceptions" terminology indicates that the forecast model used by agents does not nest RE forecast rules, and this can arise naturally in various ways that are familiar to applied econometricians. These include omitted explanatory variables, parsimonious lag structures and misspecified functional forms. In many cases, the formal existence of an RPE can be established.¹ In this paper, the RPE arises because agents are using a linear

¹The term RPE was introduced in Evans and Honkapohja (2001) but the concept is older. Branch (2006) provides a nice survey. In the context of linear models, Marcet and Sargent (1989a), Sargent (1991), Evans, Honkapohja, and Sargent (1993) and Bullard, Evans, and Honkapohja (2008) study forecast rules that omit relevant lags or dynamics. In Branch and Evans (2006a), Branch and Evans (2007) and Adam (2007) forecast models omit relevant explanatory variables, while Evans and McGough (2020b), Evans and McGough (2020a) focus on misspecified functional form. Evans and McGough (2018a) consider the case in which exogenous variables are unobserved and use autoregressions or VARs as forecast models. Branch, McGough, and Zhu (2021) combine non-observability of exogenous shocks with the presence of observable sunspot processes to demonstrate the existence of stable RPE even in models that are determinate under RE. Hommes and Zhu (2014) introduce the concept of *behavioral learning equilibria*, which casts agents as using simple AR(1) forecast models in complex economic environments. RPE have also been central in a number of empirical DSGE models, e.g. Slobodyan and Wouters (2012).

forecasting model in a non-linear world. A major gap in the literature has been the absence of RPE existence results in nonlinear models that include lagged endogenous variables, e.g. real business cycle (RBC) models and DSGE models that incorporate capital accumulation, habit persistence and/or inflation inertia. A central contribution of this paper is to provide tools to address this omission, and to show how they can be applied to nonlinear DSGE models in which agents use linear forecasting rules.

The second contribution of this paper concerns learning dynamics within a DSGE/RPE framework. The criticism that REE, as an equilibrium concept, is too demanding of a model's agents is often countered by claims that REE are most appropriately viewed as emergent outcomes of learned behaviors, in the same way that the child learns to catch a ball not by studying physics but instead by practice. In macroeconomic environments, adaptive learning (AL), identified as recursive least-squares updating of forecast model parameters, has become a dominant mechanism by which agents are modeled as learning to make forecasts and take decisions. If an economy populated with learning agents converges in an appropriate sense to an REE then it is indeed natural to view that equilibrium as an emergent outcome of learned behavior;² and the associated literature has met with much success by providing conditions that guarantee stability under adaptive learning in linearized models. However, whether the REE of the nonlinear model inherits this stability is an open question, mostly because it is difficult to model how agents learn about an object – in this case the collection of endogenous conditional distributions of the economy's state process – when the modelers themselves have difficulty even approximating those objects.

In contrast with REE of nonlinear models, RPE, by construction, are amenable to stability analysis. Under AL the forecasting models used by agents are updated over time as new data become available, just as is done in linear models. If, over time, the forecast models of agents converge in an appropriate sense to the optimal forecast model, then the RPE is *stable*, and can be viewed as an emergent outcome of a natural learning process. In linear REE models, stochastic approximation results can be used to obtain the relevant stability conditions and to establish, under AL with suitable assumptions, convergence to an REE.³ The stability conditions can generally be obtained using the “expectational stability” or “E-stability” principle described in Evans and Honkapohja (2001). Importantly, the E-stability technique can be extended in a straightforward way to assess stability under AL of RPEs.

In this paper, for a broad class of DSGE models, we provide conditions guaranteeing both existence and E-stability of an RPE local to a determinate, non-stochastic steady state. Conditional on forecast model specification, uniqueness is also established. One version of our results can be informally stated as follows: if the non-stochastic steady state of a DSGE model is determinate then, provided the exogenous drivers have sufficiently small support, there is a unique RPE associated with the linear forecast model that conditions on the same variables as those in the minimal state variable (MSV) solution to the rational

²For more on this view, see Sargent (2008).

³See Bray and Savin (1986), Marcet and Sargent (1989b) and Evans and Honkapohja (2001).

model. Furthermore, the RPE is E-stable provided that the REE of the linearized model is E-stable.

The main results of the paper leverage a technical lemma that allows us to navigate the complications introduced by lagged endogenous variables. To get a feel for the result, consider the non-linear Markov process associated with the dynamics of an economy where agents have a fixed forecasting rule. The lemma states that if the linearized dynamics are asymptotically stationary then, provided the innovations have small support, the non-linear Markov process will remain close to the linearized dynamics and will converge weakly to a unique ergodic distribution. This limiting distribution is necessary to characterize the projections which determine the optimal forecast model of the dynamic system and the resulting function mapping agent's forecast models to their optimal counterparts – the T-map. The RPE is a fixed point of the T-map, and the properties of the asymptotic distribution of the linearized dynamic are enough to guarantee uniqueness of this fixed point and its E-stability.

After establishing these results for a general class of DSGE models, we use a benchmark *real business cycle* (RBC) model to examine RPE and compare them to REE. In our model, government spending, financed by budget-neutral lump-sum taxes, serves as the exogenous driver. We compare impulse response functions across three models: the linearized REE, the RPE and the non-linear REE, the latter computed numerically by solving the associated planner's problem. We find that, compared to the linearized REE, the RPE more closely matches the nonlinear REE. This is not surprising as the RPE incorporates all of the non-linearities present in the economic environment. However we emphasize that, in our view, the importance of the RPE is *not* to provide a useful computational approximation to the REE. Instead, we treat the RPE as the appropriate equilibrium concept under natural bounded rationality assumptions. This is, first, because agents are making decisions based on plausible forecasting models; and, second, because the RPE in the RBC model is stable under adaptive learning, i.e. it can arise as an emergent outcome based on these forecasting models.

With these results for the RBC model in hand we illustrate the additional implications of the learning dynamics themselves when there is a structural change. Specifically we consider an announced, but unanticipated, permanent rise in (mean) government spending. Here, sharp departures from REE are observed in the learning dynamics: even though the policy change is announced, agents must learn over time how to adjust their forecast model. This leads to non-monotonic responses of consumption and capital en route to the new RPE.

The paper is structured as follows. Section 2 first provides a general formal framework, consistent with standard representative-agent DSGE models. Under suitable assumptions we establish, for sufficiently small bounded support of the exogenous shocks, that a locally unique RPE exists. Section 3 examines the nonlinear system under adaptive learning, specifically under recursive least-squares updating of the coefficients of the linear forecast-

ing rule, and establishes E-stability of the RPE. Section 4 adopts a standard RBC model to illustrate the results, using a calibrated model to show learning dynamics for the RPE, to compare IRFs for the REE, the linearized REE and the RPE, and to study the impact of structural change on the RPE under learning. Section 5 concludes.

2 DSGEs

We begin our examination in an abstract dynamic environment, taking the form of a system of expectational difference equations commonly associated with representative-agent DSGE models. Systems of this type can be represented in the following form:⁴

$$0 = E_t F(y_{t+1}, y_t, z_t, \zeta_t) \quad (1)$$

$$z_{t+1} = h(y_t, z_t, \zeta_t) \quad (2)$$

$$\zeta_{t+1} = (I - \rho_\zeta)\bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1}. \quad (3)$$

Here $y_t \in \mathbb{R}^n$ is the (column) vector of contemporaneously determined endogenous (jump/choice) variables, $z_t \in \mathbb{R}^m$ is the (column) vector of predetermined endogenous (state) variables and $\zeta_t \in \mathbb{R}^l$ is the (column) vector of exogenous (state) variables. The exogenous state is assumed to follow a VAR process with coefficient matrix ρ_ζ , long-run mean $\bar{\zeta}$, and exogenous shock process ε_t that represents the sole source of uncertainty in the economy.⁵ Finally, σ is a perturbation parameter which controls the level of risk in the economy.

It may be helpful to have a simple model in mind as an example. In the Ramsey model y can be taken as the vector of consumption (c), real interest rates (r) and real wages (w), z can be taken as the capital stock (k) and ζ as the log TFP shock. The function h is the law of motion for the endogenous state variable and, in the Ramsey model, can be constructed from the resource constraint. Finally, the function F includes the remaining equilibrium conditions of the model: the representative household's Euler equation and the no-arbitrage conditions of the firm.

In our general set-up, (1)-(3), the model's primitives are F , h , ρ_ζ , and the stochastic properties of ε_t . On these we make the following assumptions:

Assumptions A.

A.1 The innovations ε_t are zero mean, i.i.d., with compact support \mathcal{E} and measure μ . The matrix ρ_ζ has spectral radius less than one.

⁴DSGE models with a finite number of agent types can also be put in this form, though in the presence of incomplete asset markets the more common approach to incorporating heterogeneity assumes a continuum of agent-types, and the associated dynamics are more involved.

⁵The VAR assumption is for notational simplicity: all our results extend to general Markov processes that are locally stable.

A.2 The functions F and h are twice differentiable and there is a locally isolated non-stochastic steady state (\bar{y}, \bar{z}) .

A.3 The steady state is locally determinate.⁶

Given initial conditions z_0 and ζ_0 , a *rational expectations equilibrium* (REE) is a uniformly bounded (a.e.) triple of processes (y_t, z_t, ζ_t) satisfying (1)-(3). Assumption A.3 implies that, provided σ is sufficiently small and z_0 is near \bar{z} , there is a unique REE that remains close to the steady state. A recursive representation of this REE is a function f such that

$$y_t = f(z_t, \zeta_t). \quad (4)$$

The equations (2)-(4) then characterize the model's equilibrium dynamics when the agents are rational.

A common approach to studying the rational expectations equilibrium is to construct a linear approximation. The equilibrium dynamics, (2) and (4), may be linearized to obtain

$$\hat{y}_t = \bar{y} + f_z(\hat{z}_t - \bar{z}) + f_\zeta(\zeta_t - \bar{\zeta}) \quad (5)$$

$$\hat{z}_{t+1} = \bar{z} + h_y(\hat{y}_t - \bar{y}) + h_z(z_t - \bar{z}) + h_\zeta(\zeta_t - \bar{\zeta}), \quad (6)$$

where the hats are used to indicate the solution to the linearized model, and derivatives are evaluated at the steady state. The matrices h_y and h_ζ can be computed directly by differentiating equation (2) while f_z and f_ζ can be found using the implicit function theorem. After substituting for $(\hat{y}_t - \bar{y})$ in (6) using (5), one obtains the following law of motion for the states:

$$\begin{pmatrix} \hat{z}_{t+1} - \bar{z} \\ \zeta_{t+1} - \bar{\zeta} \end{pmatrix} = \begin{pmatrix} h_y f_z + h_z & h_y f_\zeta + h_\zeta \\ 0 & \rho_\zeta \end{pmatrix} \begin{pmatrix} \hat{z}_t - \bar{z} \\ \zeta_t - \bar{\zeta} \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} \sigma \varepsilon_{t+1}. \quad (7)$$

Our assumptions guarantee that $h_y f_z + h_z$ has spectral radius less than one, which implies that the linearized dynamics have a unique stable ergodic distribution.

A final observation will nicely motivate the expectations of our boundedly rational agents in the next section. Iterating equation (5) forward one period and rearranging terms gives a natural linear rule for forecasting y_{t+1} given current states, namely,

$$y_{t+1} \approx \hat{\psi}_0 + \hat{\psi}_z z_{t+1} + \hat{\psi}_\zeta \left((I - \rho_\zeta) \bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1} \right), \quad (8)$$

where the coefficients are given by $\hat{\psi}_0 = \bar{y} - f_z \bar{z} - f_\zeta \bar{\zeta}$, $\hat{\psi}_z = f_z$ and $\hat{\psi}_\zeta = f_\zeta$. The coefficients $\hat{\psi}$ can be viewed as approximating, to first order, the forecasting model used by rational agents when forming expectations of future choices. Even in a non-linear world this forecasting function can be used to easily evaluate the conditional expectation of any non-linear function $G(y_{t+1})$:

$$E_t^{\hat{\psi}} G(y_{t+1}) = \int_{\mathcal{E}} G(\hat{\psi}_0 + \hat{\psi}_z z_{t+1} + \hat{\psi}_\zeta (I - \rho_\zeta) \bar{\zeta} + \hat{\psi}_\zeta \rho_\zeta \zeta_t + \hat{\psi}_\zeta \sigma \varepsilon_{t+1}) \mu(d\varepsilon_{t+1}).$$

⁶This amounts to studying a linearized system associated with equations (1)-(3) and verifying the Blanchard and Kahn (1980) conditions. See Fernández-Villaverde, Rubio-Ramírez, and Schorfheide (2016) for an overview.

Note that the endogenous state vector z is predetermined: $E_t^{\hat{\Psi}} z_{t+1} = z_{t+1}$. Since z_{t+1} depends on y_t , the values of y_t and forecasts of (functions of) y_{t+1} are simultaneously determined. Returning to the Ramsey example, this implies that agents may condition on their savings decisions today when making forecasts about, say, the marginal utility of income tomorrow.

2.1 RPEs

In a rational expectations equilibrium, agents' forecasts are optimal: given their information, agents' expectations are formed using the forecast model that minimizes mean-square error. In a restricted perceptions equilibrium, agents' forecasts are constrained optimal: given their information, agents' expectations are formed using the forecast model that minimizes mean-square error, subject to the constraint that the forecast model must lie within a pre-selected class, and predicated on the assumption that all other agents behave similarly.

We assume that agents observe both state variables, z_{t+1} and ζ_t , and use a forecasting model motivated by the linearized rational expectations equilibrium. Specifically, we assume that agents forecast future values of y_{t+1} using the *perceived law of motion* (PLM)

$$y_{t+1} = \psi_0 + \psi_z z_{t+1} + \psi_\zeta \zeta_{t+1} = \psi_0 + \psi_z z_{t+1} + \psi_\zeta \left((I - \rho_\zeta) \bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1} \right). \quad (9)$$

The matrix $\psi = (\psi_0, \psi_z, \psi_\zeta)$ represents the beliefs of the economy's representative agent.

For a given set of beliefs, ψ , and states, z_t and ζ_t , the time t endogenous outcomes, y_t , obtain from a temporary equilibrium that solves

$$\underbrace{\int_{\mathcal{E}} F(\psi_0 + \psi_z h(y_t, z_t, \zeta_t) + \psi_\zeta \left((I - \rho_\zeta) \bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1} \right), y_t, z_t, \zeta_t) \mu(d\varepsilon_{t+1})}_{\equiv \hat{F}(y_t, z_t, \zeta_t, \psi, \sigma)} = 0. \quad (10)$$

This equation, which is obtained by inserting the PLM (9) into the reduced-form equation (1), implicitly gives an equilibrium policy rule $y_t = \hat{f}(z_t, \zeta_t, \psi, \sigma)$ for an economy populated by agents with beliefs ψ . While it is not possible in general to guarantee that a temporary equilibrium exists or is unique, we can guarantee that this is the case when beliefs are close enough to $\hat{\psi}$ and σ is small enough.

For fixed beliefs, the economy then evolves according to

$$y_t = \hat{f}(z_t, \zeta_t, \psi, \sigma) \quad (11)$$

$$z_{t+1} = h(y_t, z_t, \zeta_t) \quad (12)$$

$$\zeta_{t+1} = (I - \rho_\zeta) \bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1}. \quad (13)$$

Note that we are explicitly assuming that agents know z_{t+1} and the law of motion for ζ_t . We make these assumptions for simplicity. We could assume that agents do not observe

z_{t+1} and must, instead, construct forecasts for both z_{t+1} and ζ_{t+1} . All the results would go through in this environment. We discuss this further in Section 2.2.

One immediate problem faced in this analysis is that for given values of ψ and σ there is no reason to expect that the dynamics (11)-(13) are stable in any natural sense. However, the linearization of this system with beliefs $\hat{\psi}$ has a unique ergodic distribution. We would expect this property to hold for ψ near $\hat{\psi}$ and for sufficiently small σ . We formalize this concept in the following lemma.

Lemma 1 (Local Ergodicity). *If ψ is near $\hat{\psi}$, z_0 is near \bar{z} , and σ is small enough, then process (y_t, z_t, ζ_t) defined by the system (11)-(13) has a unique ergodic distribution $\mu(\psi, \sigma)$, and*

$$(y_t, z_t, \zeta_t) \xrightarrow{\mathcal{D}} (y(\psi, \sigma), z(\psi, \sigma), \zeta) \sim \mu(\psi, \sigma). \quad (14)$$

All proofs are in the Appendix. This lemma is key to our analysis. It guarantees the existence of an ergodic distribution of state variables for beliefs close enough to $\hat{\psi}$ and for small enough σ . Intuitively, the stability of the linearized REE implies that the linearization of the dynamics (11)-(13) has a unique ergodic distribution when beliefs are consistent with the linearized REE, i.e. $\psi = \hat{\psi}$. For small enough σ it's possible to bound the errors of the linear dynamics relative to the true non-linear process enough to show that the state dynamics are locally an average contraction, which guarantees a unique ergodic distribution. The remaining work is to show that we can lift this property to an open set of beliefs containing $\hat{\psi}$.

We use Lemma 1 to construct an RPE by defining a function $\psi \rightarrow T(\psi, \sigma)$ that maps beliefs onto the coefficients obtained by projecting $y(\psi, \sigma)$ onto $z(\psi, \sigma)$, $\zeta(\psi, \sigma)$, and a constant. Thus the T-map takes beliefs to the coefficients of the linear forecast model that is optimal given the ergodic distribution implied by those beliefs. A *restricted perceptions equilibrium* is a fixed point $\psi^*(\sigma)$ of this T-map. Two points merit emphasis. First, RPE beliefs are constrained optimal at the agent-level: if the economy is in an RPE and a given agent is required to pick, once and for all, a linear forecast model that conditions on current states and exogenous shocks, then $\psi^*(\sigma)$ is their best choice. Second, like an REE, an RPE is appropriately interpreted as a Nash concept: RPE beliefs are optimal for a given agent only if (almost) all other agents hold RPE beliefs. The first of our two main results may now be stated.

Theorem 1. *Given assumptions A, if $\sigma > 0$ is sufficiently small then an RPE $\psi^*(\sigma)$ exists. Furthermore:*

1. *The RPE is locally unique: for given σ there is an open neighborhood U of $\hat{\psi}$ such that $\psi^*(\sigma)$ is the unique fixed point of $T(\cdot, \sigma)$ in U .*
2. *$\psi^*(\sigma) \rightarrow \hat{\psi}$ as $\sigma \rightarrow 0$.*

This theorem provides two important conclusions. First, an RPE exists and is locally

unique for small shocks. This is as much as could be expected from a result that conditions on the local determinacy of the non-stochastic steady state. Second, an RPE meets the linearized REE in the small-shock limit, which is exactly as expected: the optimal linear forecast model becomes unconstrained optimal as the economy becomes linear. This last point is subtle and it is the reason we are using the concept of a RPE rather than a REE. The agent's forecasting model is necessarily misspecified: they are using a linear forecasting model in a non-linear world. Their decisions take into account the curvature of their utility functions and that prices that clear markets have a non-linear dependence on the aggregate state.

2.2 Alternative RPEs

The RPEs of Theorem 1 are specific to the functional form of the forecast model used by agents when forming expectations: see equation (9). This PLM was chosen for its alignment with the linearized REE given in (8). However, it may be quite natural, within a given modeling environment, to make different assumptions about the information available to agents and about the set of regressors they use to make forecasts; and the nature, and possibly the number, of the resulting RPEs will depend on the informational assumptions made and the set of regressors used. Within our framework, three natural variations arise, and we discuss these here in turn.

Above it was assumed that agents observe, and condition forecasts, on z_{t+1} . Alternatively, it may be natural to assume agents don't observe z_{t+1} , and instead use a PLM to forecast it. For example, the PLM

$$z_{t+1} = a_0 + a_z z_t + a_y y_t + a_\zeta \zeta_t \tag{15}$$

can be used to replace h in (10), yielding the temporary equilibrium condition

$$\hat{F}(y_t, z_t, \zeta_t, \psi, a, \sigma) = 0.$$

The analysis proceeds just as above, though now the RPE will have two endogenous beliefs vectors: $\psi^*(\sigma)$ and $a^*(\sigma)$. Of course establishing existence of RPE in this case requires extending Theorem 1, but this extension would be a mechanical exercise as h is invariant to beliefs.

It may be natural to assume that agents do not observe, or otherwise do not use, some of the state variables when making forecasts. Furthermore, agents may choose to condition forecasts on contemporaneously determined variables, especially when the exogenous states are not observed. For example, in the Ramsey model it may be natural to assume that agents do not observe the TFP shock, and instead forecast wages and interest rates using a vector autoregression. RPE associated with underspecified forecast models have been studied extensively within the context of linearized models. To capture this type of

behavior within our framework, we use the following PLM:

$$y_{t+1} = \psi_0 + \psi_y I_y y_t + \psi_z I_z z_{t+1} + \psi_\zeta I_\zeta \zeta_{t+1},$$

where I_\star is a selector matrix identifying which of the variables in \star are used as regressors. The analysis proceeds as before. Note that PLMs of this type may be used in conjunction with PLMs of the form (15).

Finally, it may also be natural for agents to use a PLM that allows for interaction terms or other nonlinearities. This poses no conceptual difficulty: the only requirement for our analysis is that the PLM be linear in coefficients, so that projection can be used to establish conditions needed for consistency of beliefs. In fact, RPE associated with forecast models that condition on higher order terms can globally approximate the REE to arbitrary precision.

3 Learning your RPEs

Equilibrium objects requiring accumulated information and endogenous coordination among agents are most naturally viewed as emergent outcomes of natural learning processes. In macroeconomics, adaptive learning provides the needed natural process: agents are assumed to update their forecast rules as new data become available; if, over time, their forecast rules become consistent with the equilibrium object, then it is said to be stable under adaptive learning, and thus can be viewed as an emergent outcome of a natural learning process.

RPE, by construction, are amenable to this type of stability analysis. For the case at hand, we proceed as follows: let ψ_t be the representative agent's beliefs used in period t to form forecasts and take decisions. To update their beliefs as new data become available, the agent uses recursive least squares to regress realized outcomes y_t on their information set $x_t^\top = (1, z_t^\top, \zeta_t^\top)$. The recursion can be written

$$\begin{aligned} R_{t+1} &= R_t + \gamma_{t+1} (x_t x_t^\top - R_t) \\ \psi_{t+1}^\top &= \psi_t^\top + \gamma_{t+1} R_{t+1}^{-1} x_t (y_t - \psi_t x_t)^\top. \end{aligned} \tag{16}$$

Here γ_t is the ‘‘gain sequence’’ and R_t is the period t estimate of the regressors' second moments.

It is well-known that for the decreasing-gain sequence $\gamma_t = 1/t$, equation (16) reproduces the sequence of standard least-squares estimates ψ_t , e.g. see Ch. 2 of Evans and Honkapohja (2001). In stochastic systems with AL, analytical results showing convergence to REE or RPE parameters require decreasing gains in which $\lim_{t \rightarrow \infty} \gamma_t \rightarrow 0$ at a suitable rate like t^{-1} . This assumption is appropriate when the economic environment is stationary and agents perceive it to be stationary.

In applied work the alternative “constant-gain” assumption is often used, in which $\gamma_t = \gamma$, where $0 < \gamma < 1$, with γ typically taken to be “small.” Under constant-gain recursive least squares (RLS), agents in effect discount older data at geometric rate $1 - \gamma$. For the small gain limit, formal convergence results are available for constant-gain RLS in stationary systems. These results show that beliefs converge to a distribution around the REE (or RPE) beliefs. However, in practice stochastic simulations are frequently employed to study constant-gain learning dynamics.

While decreasing gain is appropriate in stationary environments, realistic economies are subject to structural change, and consequently applied models typically assume that agents use constant-gain RLS. Prominent early examples include Sargent (1999), Branch and Evans (2006b), Bullard and Eusepi (2005), Milani (2007) and Orphanides and Williams (2007). As emphasized in Evans, Evans, and McGough (2021), three distinct reasons can be given for the use of constant gain learning:

- i *Structural change and model misspecification.* See, for example, Williams (2019).
- ii *Perception of potential structural change.* This can be viewed as a form of robust decision making, see Evans, Honkapohja, and Williams (2010).
- iii *Availability bias.* This refers to a tendency, observed in the behavioral economics literature, notably Kahneman (2011), to weight more heavily experiences that readily come to mind. In the AL literature this is interpreted as a tendency to weight current and recent experience more heavily than past experiences. This has also been described as a “fading memory” effect, e.g. Malmandier and Nagel (2016), or a “recency bias”, Cole and Milani (2021).

In line with these rationales and with the applied literature, our simulations use constant-gain learning.

Equation (16) can be combined with equations (11)-(13) to obtain a recursive system characterizing the dynamics of our RBC model populated with learning agents. The full system, reproduced here for completeness, takes as the state $(z_t, \zeta_t, R_t, \psi_t)$, and is given by

$$\begin{aligned}
x_t &= (1, z_t, \zeta_t) \\
y_t &= \hat{f}(z_t, \zeta_t, \psi_t, \sigma) \\
z_{t+1} &= h(y_t, z_t, \zeta_t) \\
\zeta_{t+1} &= (I - \rho_\zeta) \bar{\zeta} + \rho_\zeta \zeta_t + \sigma \varepsilon_{t+1} \\
R_{t+1} &= R_t + \gamma_{t+1} (x_t x_t^\top - R_t) \\
\psi_{t+1}^\top &= \psi_t^\top + \gamma_{t+1} R_{t+1}^{-1} x_t (y_t - \psi_t x_t)^\top.
\end{aligned} \tag{17}$$

Given an initial state, the system (17) determines the time path of the our economy populated by homogenous adaptive learners. If, for appropriate initial conditions, $\psi_t \rightarrow \psi^*(\sigma)$ almost surely then we say that the RPE $\psi^*(\sigma)$ is locally stable under adaptive learning.

3.1 RPEs and E-stability

Assessment of stability under adaptive learning is commonly conducted using the *E-stability Principle*, which invites practitioners to forgo the challenges of the discrete, stochastic system (17), and instead consider the continuous, deterministic system

$$\dot{\psi} = T(\psi, \sigma) - \psi. \quad (18)$$

An RPE $\psi^*(\sigma)$, as captured by a fixed point of the T-map, is *E-stable* if it is a Lyapunov stable equilibrium of (18). The E-stability Principle says that E-stable equilibria are locally stable under appropriate adaptive learning algorithms.

The dynamics of beliefs under adaptive learning algorithms like (17) are governed by the forecast error, in this case $\hat{f}(z_t, \zeta_t, \psi_t, \sigma) - \psi_t x_t$. The system (18) is analogous: the RHS acts as a forecast error, measuring the discrepancy between perception ψ and implication $T(\psi, \sigma)$. If $\psi^*(\sigma)$ is E-stable then “moving” according to this discrepancy pushes beliefs towards the RPE. The E-stability principle encapsulates the intuitive connection between these two dynamic systems.

The formal connection between the E-stability principle and systems like (17) is quite deep and technical, and relies on the theory of stochastic recursive algorithms, as developed by Ljung (1977) and first used by Marcet and Sargent (1989b) in the context of adaptive learning. The E-stability Principle provides a simple side-step to the technicalities: sufficient conditions for Lyapunov stability are obtained from the eigenvalues of $DT(\psi^*, \sigma)$: see Evans and Honkapohja (2001) for many details.

E-stability of the RPE follows as collateral serendipity from the proof of Theorem 1, where the spectral radius of DT is shown to be smaller than one.⁷ In particular, we have

Corollary 1. *If F satisfies assumptions A and if $\sigma > 0$ is sufficiently small then the RPE, $\psi^*(\sigma)$, is E-stable.*

Theorem 1 demonstrates existence and uniqueness of RPE. This corollary shows that an RPE is the expected, emergent outcome of a natural learning process.

3.2 Agent-level learning and the shadow-price approach

Our implementation of boundedly rational expectations formation has been necessarily mechanical: we did not use the details of the underlying economic model to inform the behavioral assumptions we made (e.g. which variables were forecasted and what regressors were used) because we did not specify an underlying model. This was intentional: we

⁷To establish stability under least-squares learning might be challenging – the state dynamics are not conditionally linear so the more general Markov versions of the SRA theorems are needed.

wanted to provide generic existence, uniqueness and stability results based on determinacy and E-stability, without referring to the particulars of a given economy.

The mechanical approach taken here is sometimes referred to as *reduced-form learning*: bounded rationality is implemented only after aggregation, market clearing, and often considerable algebraic manipulation has taken place. Due to its relative tractability, this is by far the most popular approach for the implementation of adaptive learning in a given modeling environment, particularly for addressing questions of applied interest.

The alternative is to adopt bounded rationality, in terms of both forecasting and *decision making*, as a behavioral primitive, necessarily implemented prior to aggregation and market clearing. This general approach – *agent-level learning* – has been examined in a variety of settings, and based on a variety of different models of boundedly rational decision making. Evans and Honkapohja (2006) use *Euler equation learning* to justify reduced form learning in a new Keynesian environment; Preston (2005) emphasizes *the long-horizon approach* in which households use a form of anticipated utility, together with their life-time budget constraints, to take decisions; and Adam and Marcet (2011) put forth the concept of *internal rationality*, in which agents act fully rationally based on their beliefs, but these beliefs may not align with endogenous reality.

While the papers just referenced, as well as others like them, consider agent-level learning implementations within the context of specific DSGE environments, Evans and McGough (2018b) develop a theory of boundedly rational decision making that can be applied to any environment in which a fully rational agent would be modeled as solving a dynamic program. Their implementation – *the shadow price approach* – takes as primitive that agents make decisions using perceived shadow prices to assess the trade-off between today and the future; and they update their perceptions over time as they realize the consequences of past decisions. Evans and McGough (2021) provide a road map for the application of this approach in general environments, as well as a comparison of it to the approaches listed above.

Evans and McGough (2018b) demonstrate that the shadow price approach achieves asymptotically optimal decision making in linear-quadratic environments. Using the new technologies developed in this paper, we are currently working to extend the theoretical results of Evans and McGough (2018b) to fully general environments.

4 The RPEs of RBCs

We adopt a benchmark *real business cycle* (RBC) model as a laboratory for a specific examination of the RPE associated with a DSGE model. Using numerical methods, we analyze the dynamics of the economy in an RPE and under the assumption that agents use a constant gain learning algorithm to learn the RPE. We compare these dynamics to those

associated with the model's REE. The dynamics under RE are obtained and presented in two ways: using a local approximation to the REE by linearizing the model around the non-stochastic steady state; and using a global approximation to the REE by applying standard techniques to approximate the solution the associated planner's problem.

4.1 Environment and calibration

Time is discrete. There is a continuum of identical households (agents) who receive flow utility from the consumption of goods and leisure. Via competitive factor markets, agents supply capital and labor to firms, the former in-elastically and the latter with Frisch elasticity χ^{-1} . The representative household's problem is

$$\begin{aligned} \max E_0 \sum_{t \geq 0} \beta^t (u(c_t) - v(n_t)) \\ s_t = r_t s_{t-1} + w_t n_t - c_t - \tau_t, \end{aligned} \quad (19)$$

together with the restriction that agents cannot run Ponzi games (NPG). Here c is consumption, s is savings in the form of (real) capital, and n is the quantity of labor supplied. Wages w and gross returns r are written in terms of consumption goods. τ represents lump sum taxes used to finance an exogenous government expenditure process. The function u measures flow utility from goods consumption, and is taken to be CRRA with risk parameter $\theta > 0$. The function v measures the disutility of labor and has the form $v(n) = \xi (1 + \chi)^{-1} n^{1+\chi}$, with $\chi > 0$ capturing the inverse Frisch elasticity, and $\xi > 0$ a scaling parameter used to tune steady-state labor hours.

Under rationality, the representative agent enters period t holding savings s_{t-1} and facing prices r_t and w_t , and taxes τ_t . Their period t consumption, savings, and labor-supply decisions must satisfy their budget constraint (19) and

$$u'(c_t) = \beta E_t \lambda_{t+1} \quad (20)$$

$$v'(n_t) = u'(c_t) w_t, \quad (21)$$

$$\lambda_t = r_t u'(c_t), \quad (22)$$

where λ_t is their period t shadow price of savings. Their decisions must also satisfy the usual transversality condition (TVC). Equation (20) is the standard Euler condition measuring the agent's consumption/savings trade-off, and equation (21) measures their labor/leisure trade-off.

There is a continuum of identical firms, each owning technology f that is CRTS in capital k and labor n : $f(k_t, n_t) = k_t^\alpha n_t^{1-\alpha}$. There are no inter-temporal frictions and capital depreciates at rate $\delta \in (0, 1)$. Competitive factor markets ensure

$$r_t = 1 + f_k(k_t, n_t) - \delta \text{ and } w_t = f_n(k_t, n_t). \quad (23)$$

Finally, each period, the government must finance an exogenous expenditure process, g_t . We assume that g_t follows a stationary AR(1) process in logs:

$$\log g_t = (1 - \rho_g) \log \bar{g} + \rho_g \log g_{t-1} + \sigma \varepsilon_t, \quad (24)$$

with $\rho_g \in (0, 1)$ and ε_t begin zero-mean, iid, and having compact support. The government runs a balanced budget each period, $\tau_t = g_t$. Market clearing, therefore, implies $k_t = s_{t-1}$, and so we may identify these variables to characterize the REE.

By modern standards, this is a particularly simple version of a DSGE model. We could have embellished this model by adding multiple shocks; pricing or financial frictions; and even heterogeneous agents. All of our theorems would still apply to any of these extensions. However, our main interest is understanding the behavior of a model populated by small agents who are using linear forecast rules in a non-linear world. To isolate this behavior, we have opted to study one of the simplest non-linear models who's dynamics are well understood.

To map this environment into our general framework, we specify our model in logs and set

$$y_t = (\log c_t, \log n_t, \log \lambda_t, \log r_t, \log w_t) \text{ and } z_t = \log k_t \text{ and } \zeta_t = \log g_t \quad (25)$$

Then the functions F and h in (1) and (2) are given by

$$F(y_{t+1}, y_t, z_t, \zeta_t) = \begin{pmatrix} u'(c_t) - \beta \lambda_{t+1} \\ u'(c_t) w_t - v'(n_t) \\ \lambda_t - r_t u'(c_t) \\ r_t - 1 - f_k(k_t, n_t) + \delta \\ w_t - f_n(k_t, n_t) \end{pmatrix} \quad (26)$$

$$h(y_t, k_t, \zeta_t) = f(k_t, n_t) + (1 - \delta)k_t - c_t - g_t. \quad (27)$$

For standard calibrations of the RBC model, assumptions A hold, so that the steady state is determinate. In the unique REE y_t is a function of the states $\log k_t$ and $\log g_t$. Following Section 2.1, we assume that boundedly rational agents forecast future values of y with a PLM which is linear in these states and given by (9). As $\log \lambda_{t+1}$ is the only variable that requires forecasting, the linear PLM can be written simply as

$$\log \lambda_{t+1} = \psi_0 + \psi_k \log k_{t+1} + \psi_g \log g_{t+1}, \quad (28)$$

with expectations formed according to

$$E_t^\psi \lambda_{t+1} = \exp(\psi_0 + \psi_k \log k_{t+1} + \psi_g ((1 - \rho_g) \log \bar{g} + \rho_g \log g_t)) \cdot E \exp(\sigma \psi_g \varepsilon_{t+1}).$$

This PLM can be combined with equations (26) and (27) to obtain the temporary equilibrium function \hat{F} , given in (10). Finally, we can write the REE functional dependence of

the shadow price on the states as: $\log \lambda_t = \Lambda(\log k_t, \log g_t)$. Linearizing Λ about the steady state yields the beliefs $\hat{\psi}$ associated with the linearized REE.

Our calibration follows that of Evans, Li, and McGough (2021). The model is calibrated to a quarterly frequency with β set to 0.99 to give an annualized steady-state real interest rate of 4%. α is set to 0.36 to target capital's share of output, while the depreciation rate δ is set to 0.025 to match a capital to output ratio of 10.26 (see Den Haan, Judd, and Juillard (2010)). Flow utility from consumption is assumed logarithmic ($\theta = 1$), and we take the Frisch elasticity, χ^{-1} , to be 1. The parameter ξ is calibrated to 9.72 to target the average supply of hours by the households of 1/3. Finally, the parameters of the government expenditures process are set to $\rho = 0.965$ and $\sigma = 0.0136$ to match the autocorrelation and standard deviation of log-linearly detrended real government expenditures, with $\bar{g} = 0.185$ to give a steady-state spending to gdp ratio of 15%. For convenience, the following table records the calibrated values.

Table 1: Calibration

parameter	β	θ	χ	ξ	α	δ	ρ	σ_ε	\bar{g}
value	0.99	1.0	1.0	9.72	0.36	0.025	0.965	0.0136	0.185

4.2 RPE Beliefs and Stability

The RPE can be computed using a sample analog of the T-map: for fixed ψ , the economy can be simulated to approximate the ergodic distribution of $(\log \lambda_t, \log k_t, \log g_t)$; we can then sample from this distribution and regress $\log \lambda$ on $\log k$, $\log g$ and a constant to obtain the sample analog of the T-map. The fixed point can then be found using a non-linear solver.

We find that the RPE values of beliefs in the non-linear model are very close but not exactly identical to $\hat{\psi}$, i.e. the REE beliefs in the linearized model. This can be observed by looking at the red and blue dashed lines in Figure 1 which represent the linear and RPE beliefs respectively. That the beliefs are close to each other reflects that our RBC model is very close to linear in logs; and indeed, for this same reason, the linearized model provides a very good approximation to non-linear model along certain measures. However, this should not be misconstrued as meaning that the models' dynamics themselves are isomorphic; indeed, in the RPE the Euler equation predicates that an agent's consumption decision is non-linear in shadow-price forecasts.

We now consider stability of RPE under real-time adaptive learning. As discussed Section 3, E-stability of the RPE strongly suggests that it will be stable under adaptive learning. Figure 1 confirms this. The figure was created using 100,000 simulations with gain set at $\gamma = 0.05$. The initial value of the second moments matrix R was taken as common across simulations and set at the linear REE value. Initial beliefs were drawn randomly within

$\pm 50\%$ of the linear REE values. The thicker black curve is the cross-sectional mean of beliefs and the outer curves identify the lowest and highest deciles.

Two observations are warranted. First, the simulations capture the stability of the RPE under learning. Even though the initial beliefs of the economy start away from the RPE they eventually coalesce around the RPE. Second, the simulations take the gain to be constant. For this reason, beliefs do not converge to RPE values but rather to an ergodic distribution that is centered very near them.⁸ If the gain decreased at an appropriate rate, for example $\gamma_t \propto t^{-1}$, then we would expect beliefs to converge almost surely to the RPE.⁹ All taken together, we find results consistent with our theoretical statements: namely, that an RPE exists, that it lies close to the linearized REE, and that it is stable under learning.

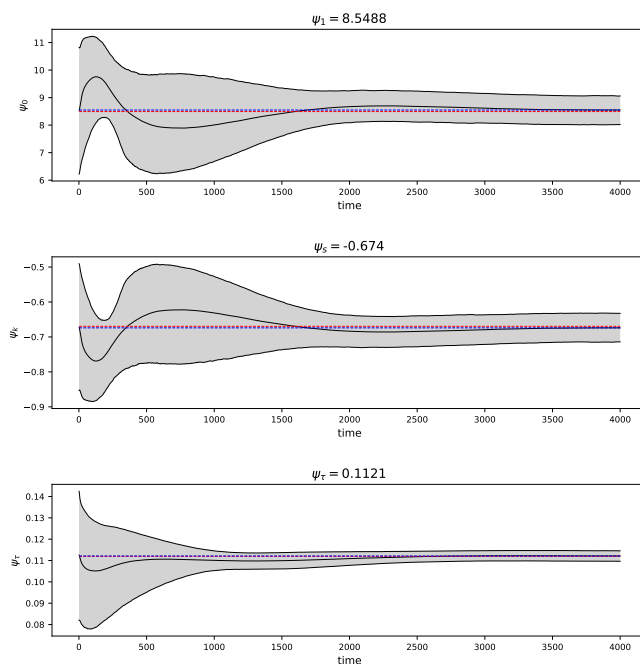


Figure 1: Beliefs trajectories, black curves gives mean beliefs, dashed (red) lines are linearized REE beliefs, dotted (blue) lines correspond to RPE beliefs, and the shaded band identifies lower and upper deciles. The upper label of each panel gives the numerical value of the associated RPE coefficient.

⁸As the gain approaches zero the mean of the ergodic distribution of beliefs approaches the RPE.

⁹To ensure almost most sure convergence it is necessarily to augment the learning algorithm with a projection facility: see Evans and Honkapohja (2001) for details.

4.3 RPE dynamics

Impulse responses to a one-standard-deviation government expenditure shock also highlight both the similarities and differences across models: see the left-hand panels of Figure 2.¹⁰ The red curve is the linear REE time-path, with black and blue the RPE and non-linear REE paths respectively. All variables are in percent deviations from their initial state values. The non-linear REE and RPE are almost indistinguishable, demonstrating, even with a linear forecasting rule, that the model with boundedly rational agents does a good job of capturing the non-linearities of the model.

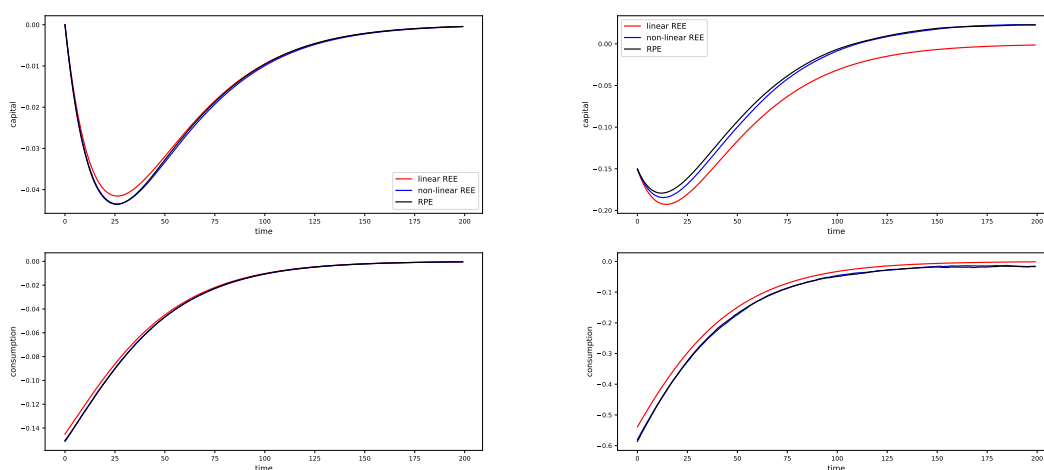


Figure 2: IRF to a government expenditure shock with $\sigma = 0.0136$ (left) and transition dynamics (right). Red curve corresponds to linear REE. The blue curve represents the response of the REE while the black line plots the response of the RPE. All variables are percent deviations from non-stochastic steady state.

To illustrate this point further, we study the transition dynamics from a point farther away from the non-stochastic steady state, yet still visited frequently by the dynamic model. To achieve this, we take the ergodic distribution of state variables generated by the REE and compute their average values conditional on government expenditures being above their mean. We initialize all three models at these states and plot expected paths of the capital stock and consumption in the right-hand panels of figure 2. All values are in percent deviation from the non-stochastic steady state. As initial state values are further away from the non-stochastic steady state, the linearized REE does a poorer job approximating the

¹⁰The IRFs for the the non-linear REE and RPE paths are computed by averaging across multiple simulations, and, for each simulation, netting out the dynamics resulting from the same sequence of innovations, but without the initial impulse shock.

outcomes of the non-linear REE while the RPE is much closer. Interestingly, the non-linear REE and RPE both converge to the same long run average values, which are different from the non-stochastic steady state. Even though the forecasting rule in the RPE is linear, agents are still exposed to the non-linearities of the model and feature the same precautionary savings motives present in the non-linear REE. This leads to both models having long-run capital levels that are larger than those of the non-stochastic steady state, which is the long run average of the linear REE model.

The take-away so far is that RPE are easily computed, and they are stable under adaptive learning, as is strongly suggested by our E-stability finding. In this way an RPE can be viewed as an emergent outcome of a natural learning process. In addition, the behavior of the economy in an RPE and with learning agents departs only very modestly from the non-linear REE, and captures aspects of the model's inherent non-linearity missed by the linearized model.

4.4 Structural change

The findings of Section 4.3 reflect that the environment is stationary. In the presence of structural change, the behavior of models with learning agents can differ quite significantly, and in our view, realistically, from the predictions of rational expectations. To examine this, we consider a simple comparative dynamics experiment: assume that the economy has existed for an extended period of time and has reached, at time 0, an ergodic distribution of state variables.¹¹ At time 0, the government unexpectedly and permanently raises government expenditures by 5%. This is achieved by raising current government expenditures, g_0 , by 5% and permanently raising the long government expenditures to a new level $\bar{g}' = 1.05\bar{g}$. Figure 3 plots the expected paths of capital and consumption as percentage deviations from their time 0 values. The blue curve is the non-linear REE and the black curve is the dynamics under learning.

The outcomes of the REE are as expected. The permanent increase in spending is implemented as a lump-sum tax, which reduces consumption and leisure via the income effect. The increase in labor raises the return to capital, thus increasing savings. This raises income over time, with consumption and leisure rising to their new, lower steady state values.

The time paths in the learning economy are strikingly different. On impact, just like the REE environment, the rise in taxes imparts an income effect leading to a fall in consumption and leisure. The differences lie in the expectations about the future. Expectations of the learning agents are formed according to

$$E_t^\psi \lambda_{t+1} = \exp(\psi_{t,0} + \psi_{t,k} \log k_{t+1} + \psi_{t,g} ((1 - \rho_g) \log \bar{g}' + \rho_g \log g_t)) \cdot E \exp(\sigma \psi_{t,g} \varepsilon_{t+1}). \quad (29)$$

where ψ_t represents the agent's beliefs at time t . When forming expectations in (29),

¹¹For the REE, $(\log k_0, \log g_0)$ is drawn from the ergodic distribution of the REE model. For the learning economy, $(\log k_0, \log g_0, \psi_0)$ is drawn from the ergodic distribution of the constant-gain learning model.

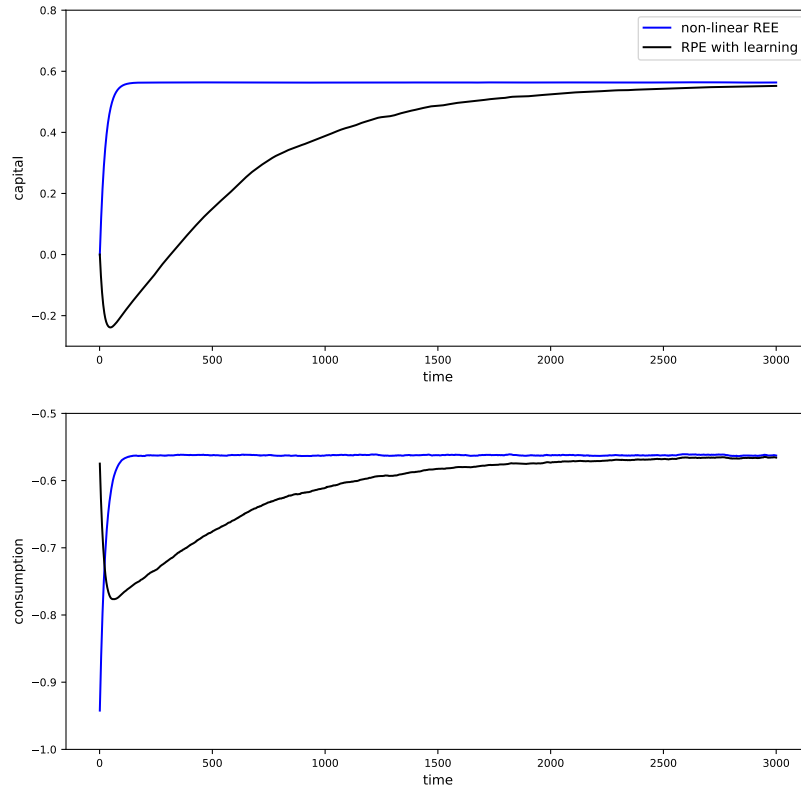


Figure 3: Expected path of consumption and capital to a permanent government spending increase (left) and dynamics of beliefs (right). Blue curve is non-linear REE, red curve is linear REE, and black curve is learning RPE. Black curves give mean beliefs, and the shaded band identifies lower and upper deciles.

agents fully recognize that the increase in government expenditures is permanent and thus accurately forecast $\log(g_{t+1})$. They, however, have never experienced a permanent change in government expenditure and, thus, must rely on their experience with transitory changes in government expenditure to forecast the effect of the permanent change on their shadow price of wealth. As a result, the 5% increase in expenditures yields an expected increase in the log shadow price of $0.05\psi_{t,g}$. The positive coefficient, $\psi_{t,g}$, implies that agents update their expectations of the shadow price of savings to be higher due to the increase in government expenditure, but not to the same degree as the fully rational agent. Thus, the fall in consumption in the learning model is smaller than that of the rational expectations economy and, as a result, the capital stock initially declines. This occurs even though agents have had experience with government expenditure shocks that were highly persistent. Over

time, however, the learning agents notice that their realized shadow price of savings is higher than expected and adjust their beliefs accordingly. In the long run, the capital stock converges to a new average value close to that of the non-linear REE.

5 Conclusions

In DSGE models the decision rules used by agents are inherently both forward-looking and nonlinear. Under rational expectations the forecasts of agents are assumed to be optimal, based on correct conditional distributions. However, those conditional distributions are endogenously determined by the future decisions of agents. Hence, the REE must be viewed as an equilibrium, and this begs the question of how the equilibrium is attained, a concern that is acute in nonlinear settings. This paper has shown a natural path through these issues, based on plausible bounded rationality considerations.

We assume agents estimate linear forecast models and use them to form the required conditional expectations needed for subjectively optimal decision-making. Estimated coefficients of the forecast models are updated over time using least-squares learning. Within an abstract framework that can represent most representative-agent DSGE models, we show under standard assumptions that there exists a restricted perceptions equilibrium – which can be viewed as a counterpart to the models associated REE – in which agents’ forecast rules are optimal within the linear class under consideration. Furthermore, the RPE is stable under least-square learning, i.e. forecast rules converge over time to the RPE, which addresses the issue of attainability of the RPE.

We used the RBC model to demonstrate the tractability of our approach, as well as to study the novel transitional learning dynamics following structural change. Applications to New Keynesian models and to DSGE models that include empirically relevant frictions are clearly feasible, using the shadow-price approach developed in Evans and McGough (2018b), and this approach can be extended to allow for agent heterogeneity. Finally, the theoretical framework developed in this paper can be applied to decision problems of the kind commonly modeled using dynamic programming, thus allowing for the development of a theory of bounded-rational decision making that can be used to model agent-level learning in DSGE environments. These applications and extensions are the subject of current research.

Appendix

We begin with a technical lemma that provides conditions for existence of stable ergodic distributions associated with homogeneous recursive dynamic systems. First, some notation. Let $\mathcal{S} \subset \mathbb{R}^n$ be open, $\mathcal{E} \subset \mathbb{R}^m$ be compact, convex, with non-empty interior containing the origin, and $\Psi \subset \mathbb{R}^k$ be open. Let $\phi : \mathcal{S} \oplus \mathcal{E} \oplus \Psi \oplus \mathbb{R} \rightarrow \mathcal{S}$ be C^k on the interior for $k \geq 2$. The dynamical system of interest is

$$y_t = \phi(y_{t-1}, \sigma \varepsilon_t, \psi, \sigma), \quad (30)$$

where $\sigma \in (0, 1)$, $y_{-1} \in \mathcal{S}$, and $\varepsilon_t \in \mathcal{E}$ is iid, zero mean.

The technical lemma establishes conditions sufficient to guarantee that (30) has a unique, stable ergodic distribution, and to do this we will apply results from Loskot and Rudnicki (1995). A definition is needed. Fixing ψ and σ , note that $\phi(\cdot, \sigma \times \cdot, \psi, \sigma) : \mathcal{S} \times \mathcal{E} \rightarrow \mathcal{S}$ is measurable.

Definition 1. *The map $\phi(\cdot, \sigma \times \cdot, \psi, \sigma)$ is an average contraction on \mathcal{S} if there is a measurable function $L : \mathcal{E} \rightarrow \mathbb{R}$ such that*

1. $\|\phi(y, \sigma \varepsilon, \psi, \sigma) - \phi(y', \sigma \varepsilon, \psi, \sigma)\| \leq L(\varepsilon) \|y - y'\|$ for all $y, y' \in \mathcal{S}$
2. $E(L(\varepsilon)) < 1$

The following assumptions provide the needed restrictions on ϕ .

Assumptions B.

- B.1 There exists $\hat{\psi} \in \Psi$ and $\bar{y}(\hat{\psi}) \in \mathcal{S}$ such that $\bar{y}(\hat{\psi}) = \phi(\bar{y}(\hat{\psi}), 0, \hat{\psi})$.
- B.2 The matrix $D_y \phi(\bar{y}(\hat{\psi}), 0, \hat{\psi})$ is invertible, with spectral radius smaller than one.

The following result is the desired technical lemma.

Lemma (Technical Lemma). *Suppose ϕ satisfies assumptions B. Then there exists $\bar{\sigma} > 0$, $r_\psi > 0$ and $r_y > 0$ such that*

1. *If $\sigma \in (0, \bar{\sigma})$ and $\psi \in B(\hat{\psi}, r_\psi)$ then the system (30) has a unique ergodic distribution $\mu(\psi, \sigma)$ with support in $Cl(B(\bar{y}(\hat{\psi}), r_y))$.*
2. *If $y_{-1} \in Cl(B(\bar{y}(\hat{\psi}), r_y))$ then $y_t \in Cl(B(\bar{y}(\hat{\psi}), r_y))$ and $y_t \xrightarrow{\mathcal{D}} y \sim \mu(\psi, \sigma)$.*

Proof. The work of the proof is establishing that $\phi(\cdot, \sigma \times \cdot, \psi, \sigma)$ is an average contraction for appropriately restricted parameters and on an appropriate state space. To this end, for given $r_y > 0$, define $L(\cdot, \sigma, \psi, r_y) : \mathcal{E} \rightarrow \mathbb{R}$ as follows:

$$L(\varepsilon, \sigma, \psi, r_y) = \sup \left\{ \|D_y \phi(y, \sigma \varepsilon, \psi, \sigma)\| : y \in Cl(B(\bar{y}(\hat{\psi}), 2r_y)) \right\}. \quad (31)$$

Noting that $\varepsilon \rightarrow L(\varepsilon, \sigma, \psi, r_y)$ is measurable, our goal is to find $\bar{\sigma} > 0$, $r_\psi > 0$ and $r_y > 0$ so that $\sigma \in (0, \bar{\sigma})$ and $\psi \in B(\hat{\psi}, r_\psi)$ imply $\phi(\cdot, \sigma \times \cdot, \psi)$ and $L(\cdot, \sigma, \psi, r_y)$ satisfy items 1 and 2 of definition 1 on $Cl(B(\bar{y}(\hat{\psi}), r_y))$.

For given $r_y > 0$, a version of the vector-valued mean value theorem implies

$$\|\phi(y', \sigma \varepsilon, \psi, \sigma) - \phi(y, \sigma \varepsilon, \psi, \sigma)\| \leq L(\varepsilon, \sigma, \psi, r_y) \|y' - y\| \text{ for } y', y \in Cl(B(\bar{y}(\hat{\psi}), r_y)),$$

which establishes item 1. Since the spectral radius of $D_y \phi(\bar{y}(\hat{\psi}), 0, \hat{\psi})$ is less than one, we may choose a matrix norm so that $\|D_y \phi(\bar{y}(\hat{\psi}), 0, \hat{\psi})\| < 1$. Since $D_y \phi$ is continuous and \mathcal{E} is compact, we can choose $r_y > 0$, $\tilde{r}_\psi > 0$ and $\bar{\sigma} > 0$ so that $\psi \in B(\hat{\psi}, \tilde{r}_\psi)$ and $\sigma \in (0, \bar{\sigma})$ implies $L(\varepsilon) \leq \Delta < 1$ for all $\varepsilon \in \mathcal{E}$. This establishes item 2.

It remains to demonstrate that for appropriately chosen $\bar{\sigma} \in (0, \bar{\sigma})$ and $r_\psi \in (0, \tilde{r}_\psi)$, the dynamics remain in the compact set $Cl(B(\bar{y}(\hat{\psi}), r_y))$ provided $\psi \in B(\hat{\psi}, r_\psi)$ and $\sigma \in (0, \bar{\sigma})$. Define

$$\Xi(\sigma, \psi) = \sup \{ \|\phi(y, \sigma \varepsilon, \psi, \sigma) - \phi(y, \sigma \varepsilon, \hat{\psi}, \sigma)\| : y \in Cl(B(\bar{y}(\hat{\psi}), r_y)), \varepsilon \in \mathcal{E} \}$$

Noting that Ξ is continuous and that $\Xi(\sigma, \hat{\psi}) = 0$, we may choose $r_\psi \in (0, \tilde{r}_\psi)$ and $\bar{\sigma} \in (0, \bar{\sigma})$ so that $\Xi(\sigma, \psi) < 1/3(1 - \Delta)r_y$. Finally, define

$$\Upsilon(\sigma) = \sup \{ \|\phi(\bar{y}(\hat{\psi}), \sigma \varepsilon, \hat{\psi}, \sigma) - \bar{y}(\hat{\psi})\| : \varepsilon \in \mathcal{E} \}.$$

Then Υ is continuous and $\Upsilon(0) = 0$, so we may further restrict $\bar{\sigma}$ to ensure $\Upsilon(\sigma) < 1/3(1 - \Delta)r_y$ for all $\sigma \in (0, \bar{\sigma})$. It follows that $y \in Cl(B(\bar{y}(\hat{\psi}), r_y)) \implies$

$$\begin{aligned} \|\phi(y, \sigma \varepsilon, \psi, \sigma) - \bar{y}(\hat{\psi})\| &\leq \|\phi(y, \sigma \varepsilon, \psi, \sigma) - \phi(y, \sigma \varepsilon, \hat{\psi}, \sigma)\| + \|\phi(y, \sigma \varepsilon, \hat{\psi}, \sigma) - \bar{y}(\hat{\psi})\| \\ &\leq \|\phi(y, \sigma \varepsilon, \psi, \sigma) - \phi(y, \sigma \varepsilon, \hat{\psi}, \sigma)\| + \|\phi(y, \sigma \varepsilon, \hat{\psi}, \sigma) - \phi(\bar{y}(\hat{\psi}), \sigma \varepsilon, \hat{\psi}, \sigma)\| \\ &\quad + \|\phi(\bar{y}(\hat{\psi}), \sigma \varepsilon, \hat{\psi}, \sigma) - \bar{y}(\hat{\psi})\| \\ &\leq \Xi(\sigma, \psi, r_y) + \Upsilon(\sigma) + L(\varepsilon, \sigma, \hat{\psi}) \|y - \bar{y}(\hat{\psi})\| \\ &\leq \frac{2(1 - \Delta)r_y}{3} + \Delta \cdot r_y < r_y, \end{aligned}$$

which demonstrates that the dynamics remain in $Cl(B(\bar{y}(\hat{\psi}), r_y))$. This completes the argument that $\phi(\cdot, \sigma \times \cdot, \psi)$ is an average contraction on $Cl(B(\bar{y}(\hat{\psi}), r_y))$ provided $\bar{\sigma} \in (0, \bar{\sigma})$ and $\psi \in (0, r_\psi)$. Proposition 1 in the Appendix of Loskot and Rudnicki (1995) then completes the proof. ■

Proof of Lemma 1 and Theorem 1. Combining equations (11) and (12) yields the following dynamics for the state variable

$$z_{t+1} = h(\hat{f}(z_t, \zeta_t, \psi, \sigma), z_t, \zeta_t) \equiv H(z_t, \zeta_t, \psi, \sigma).$$

Along with equation (13), H can be used to determine a dynamic system for $\tilde{\xi}_t \equiv (z_t, \zeta_t)$, which we denote by

$$\tilde{\xi}_t = \phi(\tilde{\xi}_{t-1}, \sigma \varepsilon_t, \psi, \sigma). \quad (32)$$

We claim ϕ satisfies assumptions A. First, notice that $\bar{z} = H(\bar{z}, \bar{\zeta}, \hat{\psi}, 0)$; it follows that ϕ satisfies A.1. Next, notice that the eigenvalues of $D_{\tilde{\xi}} \phi$ are given by the eigenvalues

of $D_z H$ and ρ_ζ , so it suffices to show that $D_z H(\bar{z}, \bar{\zeta}, 0)$ has spectral radius less than one. To this end, recall that $\hat{\psi}$ is constructed such that the linearized model with beliefs $\hat{\psi}$ is identical to the linearized rational expectations equilibrium. A direct consequence of this construction is that $D_k H(\bar{z}, \bar{\zeta}, \hat{\psi}, 0)$ must match the behavior of the endogenous state variable in the linearized RE equilibrium, $h_y f_z + h_z$, and, hence, local stability of the RE equilibrium implies $D_k H(\bar{z}, \bar{\zeta}, \hat{\psi}, 0)$ has spectral radius less than 1. We conclude that ϕ satisfies assumptions A, so that $\hat{\xi}_t$ converges to a stationary distribution for appropriately restricted parameters, and concludes the proof of Lemma 1. For the remainder of the proof we assume ψ and σ are chosen appropriately.

Finally, it will be helpful to expand the state to $\xi_t \equiv (y_t, z_t, \zeta_t)$. Since y_t is determined by z_t and ζ_t , we have that $\xi_t \xrightarrow{\mathcal{D}} \xi(\psi, \sigma)$. Abusing notation somewhat, we will denote by ϕ the transition dynamic for ξ_t . For the remainder of the proof it will prove convenient to define the random variables $\xi_1(\psi, \sigma)$ and $\xi_2(\psi, \sigma)$ such that $y_t \xrightarrow{\mathcal{D}} \xi_1(\psi, \sigma)$ and $(z_t, \zeta_t) \xrightarrow{\mathcal{D}} \xi_2(\psi, \sigma)$ with $(\xi_1(\psi, \sigma), \xi_2(\psi, \sigma)) = \xi(\psi, \sigma)$.

Now let $x(\psi, \sigma) = (1, \xi_2(\psi, \sigma))$, and define the T-map as

$$T(\psi, \sigma)^\top = (E[x(\psi, \sigma) \otimes x(\psi, \sigma)])^{-1} E[x(\psi, \sigma) \xi_1(\psi, \sigma)^\top],$$

where the expectation is taken with respect to the distribution $\xi(\psi, \sigma)$. To analyze the T-map we need to approximate $\xi(\psi, \sigma)$. To this end, define the process $\hat{\xi}_t(\psi)$ recursively as

$$\hat{\xi}_t(\psi) = D_\xi \phi \cdot \hat{\xi}_{t-1}(\psi) + D_\varepsilon \phi \cdot \varepsilon_t, \quad (33)$$

where the derivatives $D_* \phi$ are understood to be evaluated at $(\bar{\xi}(\psi, 0), 0, \psi, 0)$. Since the matrix $D_\xi \phi$ is stable there is an ergodic distribution $\hat{\xi}(\psi)$ such that $\hat{\xi}_t(\psi) \xrightarrow{\mathcal{D}} \hat{\xi}(\psi)$. Let $\hat{x}(\psi) = (1, \hat{\xi}_2(\psi))$, and let $\hat{T}(\psi)$ be the coefficients of the projection of $\hat{\xi}_1(\psi)$ onto $\hat{x}(\psi)$, i.e.

$$\hat{T}(\psi)^\top = (E[\hat{x}(\psi) \otimes \hat{x}(\psi)])^{-1} E[\hat{x}(\psi) \hat{\xi}_1(\psi)^\top],$$

where the expectation is taken with respect to the distribution $\hat{\xi}(\psi)$. We analyze the T-map by connecting it to \hat{T} .

We begin by expanding ϕ about $\xi = \bar{\xi}(\psi)$ and $\sigma = 0$, which gives us

$$\frac{\xi_t(\psi, \sigma) - \bar{\xi}(\psi)}{\sigma} \xrightarrow{\mathcal{D}} \hat{\xi}(\psi) + \Delta(\psi, \sigma), \text{ with } \Delta(\psi, \sigma) = \mathcal{O}(\sigma). \quad (34)$$

We now use (34) establishing the following:

Claim. For appropriately chosen r_ψ ,

1. $T(\cdot, \sigma) : B(\hat{\psi}, r_\psi) \rightarrow \mathbb{R}^3$ is C^1
2. $\lim_{\sigma \rightarrow 0} T(\hat{\psi}, \sigma) = \hat{T}(\hat{\psi}) = \hat{\psi}$

$$3. \lim_{\sigma \rightarrow 0} DT_{\psi}(\hat{\psi}, \sigma) = D\hat{T}(\hat{\psi})$$

To establish this claim, fix σ and drop it as an argument; then, write equation (32) as $\xi_t(\psi) = \Phi_t(\xi_{-1}, \varepsilon^t, \psi)$, where $\varepsilon^t = (\varepsilon_0, \dots, \varepsilon_t)$, and Φ_t is C^2 in ψ . We may compute

$$T_t(\psi)^\top = \left(E_{\mathcal{F}_t} \begin{pmatrix} 1 & \xi_{2t}(\psi)^\top \\ \xi_{2t}(\psi) & \xi_{2t}(\psi)\xi_{2t}(\psi)^\top \end{pmatrix} \right)^{-1} E_{\mathcal{F}_t} \begin{pmatrix} 1 \\ \xi_{2t}(\psi) \end{pmatrix} \xi_{1t}(\psi)^\top, \quad (35)$$

where \mathcal{F}_t is the joint distribution of $(\xi_{-1}, \varepsilon^t)$. Now notice that, for fixed ψ , the functions that comprise the components of the matrices in (35) are smooth in $\xi(\psi)$, so that $T_t(\psi) \rightarrow T(\psi)$. In fact, more is true. Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be C^k for $k \geq 1$ and let $f_t(\psi) = E_{\mathcal{F}_t} f(\xi_t(\psi))$. Then

$$D_{\psi_i} f_t(\psi) = \int D_{\xi} f(\xi_t) \cdot D_{\psi_i} \Phi_t(\xi_{-1}, \varepsilon^t, \psi) d\mathcal{F}_t(\varepsilon^t)$$

is C^{k-1} . It follows that $D_{\psi_i} T_t(\psi)$ is continuous. Combining this with the compactness of the state space, we can find r_ψ so that $D_{\psi_i} T_t$ is uniformly bounded on $Cl(B(\hat{\psi}, r_\psi))$.¹² Thus $\{T_t\}$ and $\{DT_t\}$ are uniformly equicontinuous on $Cl(B(\hat{\psi}, r_\psi))$. We may use the point-wise convergence of $T_t(\psi)$ to $T(\psi)$ and apply the Ascoli-Arzela theorem to conclude that T_t converges to T uniformly. We may again apply the Ascoli-Arzela theorem, now to $\{DT_t\}$ to conclude that a subsequence $\{DT_{\tau(t)}\}$ converges uniformly to some continuous function on $Cl(B(\hat{\psi}, r_\psi))$. But we may then restrict attention to this subsequence to get that $T_{\tau(t)} \rightarrow T$ uniformly and $\{DT_{\tau(t)}\}$ converges uniformly, whence T is differentiable and $DT = \lim_{t \rightarrow \infty} DT_{\tau(t)}$, which is continuous. This establishes item 1 of the claim. Items 2 and 3 are straightforward computations based on the RHS of (34).

The final piece of the puzzle concerns continuity of T in σ . We may use symmetry to define $T(\psi, \cdot)$ on $(-\bar{\sigma}, 0)$, and define $T(\psi, 0) \equiv \lim_{\sigma \rightarrow 0} T(\psi, \sigma)$. We claim that $T(\psi, \cdot): (-\bar{\sigma}, \bar{\sigma}) \rightarrow B(\hat{\psi}, \delta)$ is continuous. It suffices to consider $\sigma \neq 0$. Pick a closed interval I about σ that does not contain zero. Then $D_\sigma T_t$ is uniformly bounded on I , so that the Ascoli-Arzela theorem provides for uniform convergence of T_t to T on I .¹³ Thus T is continuous on the interior of I , which contains σ .

¹²That ϕ satisfies assumption A allows us to conclude, through lemma 1, that ξ_t is uniformly bounded. To obtain the same result for $D_\psi \xi_t$ we note that

$$D_\psi \xi_t = D_\psi \phi(\xi_{t-1}, \sigma \varepsilon_t, \psi, \sigma) + D_\xi \phi(\xi_{t-1}, \sigma \varepsilon_t, \psi, \sigma) D_\psi \xi_{t-1}.$$

Combined with the law of motion for ξ_t this equation constructs a function $\hat{\phi}$ operating on the stacked vector $\hat{\xi}_t = [\xi_t, D_\psi \xi_t]'$ such that

$$\hat{\xi}_t = \hat{\phi} \left(\hat{\xi}_{t-1}, \sigma \varepsilon_t, \psi, \sigma \right).$$

It is straightforward to verify that $\hat{\phi}$ satisfies assumption A and, hence, $D_\psi \xi_t$ is uniformly bounded.

¹³Uniform boundedness of $D_\sigma T_t$ is inherited from the uniform boundedness of $D_\sigma \xi_t$. As $D_\sigma \xi_t$ satisfies

$$D_\sigma \xi_t = D_\sigma \phi(\xi_{t-1}, \sigma \varepsilon_t, \psi, \sigma) + D_\xi \phi(\xi_{t-1}, \sigma \varepsilon_t, \psi, \sigma) D_\sigma \xi_{t-1}$$

we can use the same approach as footnote 12 to show uniform boundedness of $D_\sigma \xi_t$.

We are finally able to finish the proof. It is straightforward to show that $D\hat{T}(\hat{\psi})$ invertible. It follows from items 1 and 3 of the claim that, local to $\sigma = 0$, the map $\psi \rightarrow T(\psi, \sigma)$ is injective. By item 2 of the claim, $T(\hat{\psi}, 0) = \hat{\psi}$, and then the claim plus the continuity of T in σ allows for the application of a continuous version of implicit function theorem – see Theorem 5.1 of Blackadar (2015) – to show that for σ near zero there is a unique solution to $T(\psi, \sigma) = \psi$. ■

References

- ADAM, K. (2007): “Experimental Evidence on the Persistence of Output and Inflation,” *Economic Journal*, 117, 603–636.
- ADAM, K., AND A. MARCET (2011): “Internal Rationality, Imperfect Market Knowledge and Asset Prices,” *Journal of Economic Theory*, 146, 1224–1252.
- BLACKADAR, B. (2015): “A General Implicit/Inverse Function Theorem,” Discussion paper.
- BLANCHARD, O. J., AND C. M. KAHN (1980): “The solution of linear difference models under rational expectations,” *Econometrica: Journal of the Econometric Society*, pp. 1305–1311.
- BRANCH, W. (2006): “Restricted Perceptions Equilibria and Learning in Macroeconomics,” pp. 135–160. Cambridge University Press, Cambridge, UK.
- BRANCH, W., B. MCGOUGH, AND M. ZHU (2021): “Statistical Sunspots,” *Theoretical Economics*.
- 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.
- BRAY, M., AND N. SAVIN (1986): “Rational Expectations Equilibria, Learning, and Model Specification,” *Econometrica*, 54, 1129–1160.
- 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.
- BULLARD, J., G. W. EVANS, AND S. HONKAPOHJA (2008): “Monetary Policy, Judgment and Near-Rational Exuberance,” *American Economic Review*, 98, 1163–1177.
- COLE, S. J., AND F. MILANI (2021): “Heterogeneity in Individual Expectations, Sentiment, and Constant-Gain Learning,” *Journal of Economic Behavior and Organization*, 188, 627–650.
- DEN HAAN, W. J., K. L. JUDD, AND M. JUILLARD (2010): “Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty,” *Journal of Economic Dynamics and Control*, 34(1), 1–3.

- EVANS, D., G. W. EVANS, AND B. MCGOUGH (2021): “Learning When to Say No,” *Journal of Economic Theory*, 94, doi: <https://doi.org/10.1093/acrefore/9780190625979.013.508>.
- EVANS, D., J. LI, AND B. MCGOUGH (2021): “Local Rationality,” Discussion paper.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.
- 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., S. HONKAPOHJA, AND N. WILLIAMS (2010): “Generalized Stochastic Gradient Learning,” *International Economic Review*, 51, 237–262.
- EVANS, G. W., AND B. MCGOUGH (2018a): “Equilibrium Selection, Observability and Backward-Stable Solutions,” *Journal of Monetary Economics*, 98, 1–10.
- (2018b): “Learning to Optimize,” mimeo, University of Oregon.
- (2020a): “Equilibrium Stability in a Nonlinear Cobweb Model,” *Economics Letters*, 193, 109130. <https://doi.org/10.1016/j.econlet.2020.109130>.
- (2020b): “Stable Near-rational Sunspot Equilibria,” *Journal of Economic Theory*, 186, 104987. <https://doi.org/10.1016/j.jet.2019.104987>.
- (2021): “Agent-level Adaptive Learning,” *Oxford Research Encyclopedia of Economics and Finance*, <https://doi.org/10.1093/acrefore/9780190625979.013.620>.
- FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ, AND F. SCHORFHEIDE (2016): “Solution and estimation methods for DSGE models,” in *Handbook of macroeconomics*, vol. 2, pp. 527–724. Elsevier.
- HOMMES, C. H., AND M. ZHU (2014): “Behavioral Learning equilibria,” *Journal of Economic Theory*, 150, 778–814.
- KAHNEMAN, D. (2011): *Thinking Fast and Slow*. Farrar, Straus and Giroux, New York.
- LJUNG, L. (1977): “Analysis of Recursive Stochastic Algorithms,” *IEEE Transactions on Automatic Control*, 22, 551–575.
- LOSKOT, K., AND R. RUDNICKI (1995): “Limit Theorems for Stochastically Perturbed Dynamical Systems,” *Journal of Applied Probability*, 32(2), 459–469.

- MALMANDIER, U., AND S. NAGEL (2016): “Learning From Inflationary Experiences,” *Quarterly Journal of Economics*, 131, 53–87.
- 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.
- MILANI, F. (2007): “Expectations, Learning and Macroeconomic Persistence,” *Journal of Monetary Economics*, 54, 2065–2082.
- 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.
- SARGENT, T. J. (1991): “Equilibrium with Signal Extraction from Endogenous Variables,” *Journal of Economic Dynamics and Control*, 15, 245–273.
- (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- (2008): “Evolution and Intelligent Design,” *American Economic Review*, 98, 5–37.
- SLOBODYAN, S., AND R. WOUTERS (2012): “Estimating a Medium-Scale DSGE Model with Expectations Based on Small Forecasting Models,” *American Economic Journal: Macroeconomics*, 4, 65–101.
- WILLIAMS, N. (2019): “Escape Dynamics in Learning Models,” *Review of Economic Studies*, 86, 882–912.