

Finite Horizon Learning

William Branch
University of California, Irvine

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

Bruce McGough
Oregon State University

November 14, 2010

Abstract

Incorporating adaptive learning into macroeconomics requires assumptions about how agents incorporate their forecasts into their decision-making. We develop a theory of bounded rationality that we call *finite-horizon learning*. This approach generalizes the two existing benchmarks in the literature: Euler-equation learning, which assumes that consumption decisions are made to satisfy the one-step-ahead perceived Euler equation; and infinite-horizon learning, in which consumption today is determined optimally from an infinite-horizon optimization problem with given beliefs. In our approach, agents hold a finite forecasting/planning horizon. We find for the Ramsey model that the unique rational expectations equilibrium is E-stable at all horizons. However, transitional dynamics can differ significantly depending upon the horizon.

1 Introduction

The rational expectations (RE) hypothesis of the 1970's places individual optimization and expectation formation at the forefront of macroeconomic research. Although RE is the natural benchmark for expectation formation, it is at the same time a very

strong assumption, subject both to theoretical criticisms¹ and to plausible modifications that allow for a broader notion of bounded rationality.²

Today, dynamic stochastic general equilibrium (DSGE) models are the mainstay of macroeconomic modeling. To the extent that DSGE models embracing RE are unable to account adequately for the co-movements and time-series properties observed in the macroeconomic data, alternative mechanisms for expectation formation provide a plausible avenue for reconciliation; and, in the 30 years since the birth of the literature, adaptive learning has become rationality's benchmark replacement.³ While this literature originally focused on the conditions under which an equilibrium would be stable when rational expectations are replaced with an adaptive learning rule, increasingly there has been an emphasis on transitional or persistent learning dynamics that have the potential for generating new phenomena.

In the early literature, adaptive learning was applied either to ad-hoc models or to models with repeated, finite horizons such as the Muth model (Bray (1982), Bray and Savin (1986)) and the overlapping generations model of money (Woodford (1990)). However, micro-founded infinite horizon DSGE models provide a distinct challenge. The first attempts at modeling adaptive learning in infinite horizon DSGE models employed what we call "reduced-form learning," in which RE are replaced in the equilibrium conditions with a boundedly rational expectations operator and the stability of the equilibrium is then studied (see, e.g., Evans and Honkapohja (2001) and Bullard and Mitra (2002)).

While this was a natural first step in the study of equilibrium stability in a DSGE model, the ad-hoc nature of reduced-form learning is disconnected from the underlying micro-foundations of modern macroeconomic models. To address this concern, and to better understand the link between agents' choices and their forecasts in the context of an infinite horizon model, Honkapohja, Mitra, and Evans (2002) and Evans and Honkapohja (2006) provide a model of bounded rationality, which they called "Euler-equation learning," in which individual agents are assumed to make forecasts both of the relevant prices and of their own behavior, and then make decisions based on these forecasts to satisfy their perceived Euler equation. The Euler equation itself is taken as a behavioral primitive, capturing individual decision making. Evans and Honkapohja (2006) show, in a New Keynesian model, that Euler-equation learning is equivalent to reduced form learning.

The literature has proposed other learning mechanisms as alternatives to Euler-equation learning. Infinite-horizon learning, developed in Marcet and Sargent (1989),

¹See, for example, Guesnerie (2005).

²See Sargent (1993) for a survey of possible approaches.

³For a recent discussion see Evans and Honkapohja (2010)

and emphasized by Preston (2005), posits that agents make decisions to meet their Euler equations at all forward iterates and, where appropriate, also imposes their expected lifetime budget constraint.⁴ Shadow price learning, developed in Evans and McGough (2010), assumes that agents make choices conditional on the perceived value of additional future state variables. These alternative learning mechanisms are discussed in more detail in Section 3.2.

Euler-equation learning identifies agents as 2-period planners: they make decisions today based on their forecasts of tomorrow. Under rationality, this type of behavior is optimal: forecasts of tomorrow contain all the information needed to make the best possible decision today. If agents are boundedly rational, however, it is less clear that a 2-period planning horizon is optimal, or even adequate: perhaps a longer planning horizon is appropriate. The infinite-horizon approach takes this position to the extreme by positing an infinite planning horizon. By incorporating the lifetime budget constraint into the choice process, the agent is making decisions to satisfy his (perceived) Euler equation at all iterations and his transversality condition; in fact, infinite-horizon learning can be interpreted as assuming that private agents each period fully solve their dynamic programming problem, given their beliefs. While this has appeal in that it is consistent with the micro-foundations of the associated model, it has a number of drawbacks:

1. agents are required to make forecasts at all horizons, even though most forecasters in fact have a finite horizon;
2. agents are assumed to have sufficient sophistication to solve their infinite-horizon dynamic programming problem;
3. agents' behavior is predicated upon the assumption that their beliefs are *correct*.

This last point, in particular, is a strong assumption.⁵ In an adaptive learning model, agents' beliefs are updated by an estimation procedure – for example, recursive least squares – and therefore in any given period they will not correctly capture the joint distribution of the model's endogenous variables. If an agent knows his beliefs are wrong and likely to change in the future, that is, if an agent recognizes that his parameter estimates will evolve over time, it is no longer obvious that the agent's optimal decision is determined by the fully optimal solution to his dynamic programming problem *given his current beliefs*. While this point holds both for short-horizon

⁴See also Sargent (1993), pp. 122 - 125.

⁵The approach is typically justified by appealing to an anticipated utility framework. See Kreps (1998) and Cogley and Sargent (2008).

and long-horizon learning, it is most telling in the case of infinite-horizon learning, in which considerable weight is placed on distant forecasts, using a forecasting model that may become greatly revised. This reasoning suggests that agents may do best with finite-horizon models that look further ahead than one period, but do not attempt to forecast beyond some suitable finite horizon.

This paper generalizes the existing approaches to decision making to environments in which agents form expectations adaptively. We bridge the gap between Euler-equation learning and infinite-horizon learning, by developing a theory of finite-horizon learning. We ground our analysis in a simple dynamic general equilibrium model, the Ramsey model, and our approach is to allow agents to make decisions based on a planning horizon of a given finite length N . Euler equation learning is particularly easy to generalize: we iterate the Euler equation forward N periods and assume agents make consumption decisions today based on forecasts of consumption N periods in the future, and on forecasts of the evolution of interest rates during those N periods. We call this implementation of learning “N-step Euler equation learning.”

For reasons discussed below, N-step Euler equation learning does not reduce to infinite-horizon learning in the limit as the horizon approaches infinity. In fact, a distinct learning mechanism is required to provide a finite horizon analog to infinite horizon learning. We accomplish this by incorporating the Euler equation, iterated forward n periods for $1 < n \leq N$, into the budget constraint, which itself is discounted and summed N times. Through this construction, decisions are conditional on savings yesterday, the future evolution of interest rates and wages, and on expected future savings. We call the resulting learning mechanism “N-step optimal learning” because it leads to decisions which would be optimal given an N period problem *conditional on expected future savings*.

We show that for both learning mechanisms and all horizon lengths, the Ramsey model’s unique rational expectations equilibrium is stable under learning. There are, however, important differences along a transition path. By examining the expected paths of agents’ beliefs, we find that both learning mechanisms impart oscillatory dynamics. However, for longer planning horizons, these oscillations become negligible.

2 The Ramsey model and reduced form learning

The Ramsey model provides a simple, tractable laboratory for our exploration of finite horizon learning (FHL). In this section, we review the model and analyze the stability of its unique rational expectations equilibrium under reduced-form learning.

2.1 The Ramsey model

We consider a standard version of the Ramsey model. There are many identical households with CRRA preferences who supply their unit endowment of labor inelastically and face a consumption/savings decision. The representative household's problem is given by

$$\begin{aligned} \max_{\{c_t, k_t\}_{t \geq 0}} \quad & E \sum_{t \geq 0} \beta^t u(c_t) \\ \text{s.t.} \quad & s_t = w_t + (1 + r_t)s_{t-1} - c_t + \pi_t \end{aligned}$$

where s_{t-1} is the savings (in the form of capital) held by the household at the beginning of time t , c_t is the time t consumption level, r_t is the real return on savings, w_t is the real wage, and π_t is profit from the household's portfolio of firm shares. Here s_{-1} is given, and $s_t \geq 0$ and $0 \leq c_t \leq w_t + (1 + r_t)s_{t-1}$ are additional constraints.

The associated Euler equation is given by

$$u'(c_t) = \beta E_t(1 + r_{t+1})u'(c_{t+1}),$$

which we may linearize as

$$c_t = E_t c_{t+1} + a E_t r_{t+1} \tag{1}$$

where $a = -r\beta/\sigma$, and σ is the relative risk aversion. Also, all variables are now written in proportional deviation from steady state form.

There are many identical firms, each having access to a Cobb-Douglas production function $F = k^\alpha n^{1-\alpha}$ in capital and labor.⁶ Firms rent capital and hire labor in competitive factor markets, sell in a competitive goods market, and face no adjustment costs. This simple modeling of firm behavior, together with the assumptions on the production function, implies that factor prices are equal to the associated marginal products and firms' profits are zero. Incorporating these implications into the flow budget constraint and using market clearing to identify s_t with k_{t+1} provides the capital accumulation equation. Imposing equilibrium interest rates into the household's Euler equation results in the following reduced form system of expectational difference equations:

$$k_{t+1} = \delta_1 c_t + \delta_2 k_t \tag{2}$$

$$c_t = E_t c_{t+1} + b k_{t+1}. \tag{3}$$

⁶Typically there would be a stochastic productivity component in the production function. Without a loss of generality, the analysis in this paper assumes a non-stochastic economy.

The coefficients are $\delta_1 = -c/k$, $\delta_2 = (1 + F_k - \delta)$ and $b = akF_{kk}/r$, where variables without time subscripts are steady-state levels and all derivatives are evaluated at the steady state. Notice that because capital is predetermined there is no expectations operator in front of k_{t+1} . The system (2), (3) is generically determinate (under the usual assumptions on utility and technology), and the unique REE may be written $c_t = \hat{A}k_t$ and $k_t = (\delta_1\hat{A} + \delta_2)k_{t-1}$.

2.2 Reduced form learning

In order to form rational expectations, agents in the economy must know the actual distributions of all variables, which depend in part on their own behavior and beliefs. Instead of adopting this framework, Evans and Honkapohja (2001) assume that agents behave as econometricians: given a forecasting model whose specification is consistent with the equilibrium of interest, agents form conditional expectations and update their perceived coefficients as new data become available. Specifically, throughout the remainder of the paper, we attribute to agents a *perceived law of motion* for consumption:

$$c_t = H + Ak_t. \quad (4)$$

Since only the Euler equation depends explicitly on expectations, it seems reasonable to assume that agents know the coefficients for the capital accumulation equation (2) and the manner in which real interest rates are related to the capital stock. We could have agents estimate these coefficients, but since there is no feedback involved in this estimation, stability results would not be affected.

In the present case of reduced-form learning we will not be precise about the “actions” taken given the forecasts and whether these are consistent with economic equilibrium. Therein lies the fundamental difference between RF-learning and agent-based learning mechanisms.

We take as given the reduced form equation

$$c_t = E_t^* c_{t+1} + bE_t^* k_{t+1}, \quad (5)$$

which has been modified to incorporate bounded rationality: E_t^* is taken to be a boundedly rational expectations operator based on the agents’ forecast model. Conditional on the perceived law of motion (4), expectations are

$$E_t^* c_{t+1} = H + AE_t^* k_{t+1}.$$

It remains to identify agents’ forecasts of the future capital stock, and we do so by assuming agents know the coefficients in the capital accumulation equation (2), δ_i , so

that

$$E_t^* k_{t+1} = \delta_1 c_t + \delta_2 k_t.$$

Plugging in expectations to the reduced-form equation (5) leads to the following actual law of motion (ALM) for consumption:

$$c_t = \frac{H}{1 - (A + b)\delta_1} + \frac{(A + b)\delta_2}{1 - (A + b)\delta_1} k_t. \quad (6)$$

Much like when employing the method of undetermined coefficients, there is a mapping from the perceived coefficients in (4) to the actual coefficients in (6) that are implied by the PLM. Notice, in particular, that the PLM for consumption consists of a constant and a coefficient on the current capital stock. In the actual law of motion, which depend on these beliefs, actual consumption depends on a constant and the current capital stock. Referring to the mapping from the perceived law of motion to the actual law of motion as the “T-map”, it is immediate that the ALM identifies the T-map as

$$\begin{aligned} A &\rightarrow \frac{(A + b)\delta_2}{1 - (A + b)\delta_1} \\ H &\rightarrow \frac{H}{1 - (A + b)\delta_1}. \end{aligned}$$

Notice that the unique REE $(0, \hat{A})$ is a fixed point to the T-map. The T-map plays a prominent role in expectational stability analysis as we see next.

Expectational stability analysis asks whether reasonable learning rules based on perceived laws of motion like (4) will converge to a rational expectations equilibrium. It turns out that a straightforward and intuitive condition governs whether an equilibrium is E-stable. Let $\Theta = (H, A)'$ summarize the household’s beliefs. Since the REE is a fixed point of the T-map it is also a resting point of the following ordinary differential equation

$$\dot{\Theta} = T(\Theta) - \Theta. \quad (7)$$

The right hand side of the ode is the difference between the actual coefficients and the perceived coefficients. According to the ode, a reasonable learning rule should adjust perceived coefficients toward actual coefficients, with the resting point being an REE. The E-stability Principle states that if an REE corresponds to a Lyapunov stable rest point of the E-stability differential equation then it is locally stable under least squares learning. An REE will be E-stable when the T-map contracts to the unique REE. Thus stability under learning may be assessed by analyzing the stability properties of (7). Below, we compute the T-map for each learning environment and assess the E-stability properties.

While the reduced form learning mechanism is simple and appealing, it is vague on the interaction between forecasts and the implied agent behavior. The argument for this mechanism is that agents form forecasts and then “act accordingly,” and that the implications of their actions are well-captured by the reduced form equation (5). This may be greeted with some suspicion because, while (5) is developed from the agent’s Euler equation, it already has equilibrium prices imposed. More sophisticated DSGE models, such as RBC or New Keynesian models, have reduced form equations that are considerably more complicated, thus making interpretation of reduced form learning that much more difficult.

3 Euler equation learning and alternatives

To place learning in DSGE models on a more firm footing, Evans and Honkapohja (2006) introduce Euler equation learning. Evans and Honkapohja take the Euler equation (1) as the behavioral primitive, and take care to distinguish between individual quantities and aggregate variables. As it will serve as a platform to launch our investigations of finite horizon learning, we review Euler equation learning in detail; then we provide some discussion of other learning mechanisms.

3.1 Euler equation learning

Under Euler equation learning, the Euler equation is taken as the primitive equation capturing agent behavior. Intuitively, agents make consumption decisions today to equate marginal loss with expected marginal benefit. For each agent i , there is an Euler equation:

$$c_t^i = E_t^i c_{t+1}^i + a E_t^i r_{t+1}, \quad (8)$$

where E^i is agent i ’s (possibly) boundedly rational expectations operator. We emphasize the behavioral assumption identifying Euler equation learning as follows:

Euler equation learning behavioral assumption. *The Euler equation learning assumption identifying consumption behavior in terms of future forecasts is given by (8).*

Agent i forms forecasts of r_{t+1} and c_{t+1}^i , and then uses these forecasts to determine demand for the current period’s consumption goods. Since

$$r_t = (F_{kk}k/r) k_t \equiv Bk_t,$$

that is, since there is no feedback in the determination of the dependence of r on k , we assume agents know $r_t = Bk_t$, and thus forecast future interest rates by forecasting

future aggregate capital. To forecast future consumption, we assume that agents adopt a perceived law of motion which conditions on current interest rates and current wealth. For simplicity, and to promote comparison to real time learning, we exploit the homogeneity of the model and assume agents recognize that past wealth has been equal to aggregate capital, and that they forecast future wealth accordingly. Together, these assumptions provide the following forecasting model:

$$c_t^i = H^i + A^i k_t.$$

As for reduced form learning, we assume that agents know the values of δ_i . Therefore,

$$\begin{aligned} E_t^i r_{t+1} &= B\delta_1 c_t + B\delta_2 k_t \\ E_t^i c_{t+1}^i &= H^i + A^i(\delta_1 c_t + \delta_2 k_t). \end{aligned}$$

Given these forecasts, we may use (8) to identify agent i 's consumption decision:

$$c_t^i = H^i + (A^i + aB)\delta_1 c_t + (A^i + aB)\delta_2 k_t.$$

Imposing homogeneity, so that $c^i = c$, $H^i = H$, and $A^i = A$, allows us to compute the equilibrium dynamics given beliefs:

$$c_t = \frac{H}{1 - (A + aB)\delta_1} + \frac{(A + aB)\delta_2}{1 - (A + aB)\delta_1} k_t.$$

These dynamics comprise the ALM for the economy and thus identify the Euler equation learning model's T-map:

$$A \rightarrow \frac{(A + b)\delta_2}{1 - (A + aB)\delta_1} \tag{9}$$

$$H \rightarrow \frac{H}{1 - (A + aB)\delta_1}, \tag{10}$$

which may then be used to analyze stability under learning. Since $aB = b$, we note that Euler equation learning provides the same T-map as reduced form learning. In this way, Evans and Honkapohja are able to justify and provide a foundation for reduced-form learning.

3.2 Other implementations of learning

The coupling of agent level decision making and boundedly rational forecasting has been considered by a variety of other authors, and in this section we discuss two alternate implementations of learning in infinite horizon models: shadow price learning and infinite horizon learning.

The infinite horizon learning mechanism was first developed by Marcet and Sargent (1989) and has received renewed attention in Preston (2005) and Eusepi and Preston (2010). Under infinite horizon learning, agents make decisions so as to meet their Euler equations at all forward iterates and their expected lifetime budget constraint. Notably, this requires that agents account, a-priori, for their transversality condition; in this way, agents are making optimal decisions given their beliefs, which are captured by their forecasting model. Preston has found that under some circumstances the stability conditions implied by infinite horizon learning are different (and more restrictive) than the conditions implied by Euler equation learning. A nice comparison of Euler equation learning and infinite horizon learning is provided by Evans, Honkapohja, and Mitra (2009). In the next section, we establish infinite horizon learning as a limiting case of one of our finite horizon learning implementations.

Evans and McGough (2010) take a different approach to coupling decision theory and learning agents: they model agents as 2-period planners who choose controls today based on their perceived value of the state tomorrow. Evans and McGough call this simple behavioral procedure “shadow price learning,” and establish a general result showing that under shadow price learning, agents will eventually learn to make optimal decisions. They further show that, under certain circumstances, shadow price learning reduces to Euler equation learning.

4 Stability under finite horizon learning

By modeling the representative household as an Euler equation learner, we impose that decisions be made based on one-period-ahead forecasts. This assumption is in sharp contrast to the benchmark behavior of the rational agent and to the imposed behavior in infinite horizon learning: each is to required form forecasts at all horizons. Existing models involve only the two extreme cases – one-period horizon and infinite horizon – which are at odds with the casual observation that most forecasters have a finite forecasting horizon. This section presents our generalization of adaptive learning to environments with finite planning/forecasting horizons.

We construct two finite-horizon learning mechanisms: “N-step Euler equation learning”, which generalizes Euler equation learning to an N-period planning horizon; and, “N-step optimal learning” where agents solve an N-period optimization problem with boundedly rational forecasts. We note that N-step optimal learning has infinite horizon learning as a limiting case.

4.1 N-step Euler equation learning

We modify Euler equation learning to allow for more far-sighted individuals by iterating (8) forward N periods:

$$c_t^i = E_t^i c_{t+N}^i + a E_t^i \sum_{s=1}^N r_{t+s}. \quad (11)$$

We interpret this equation as capturing individuals who are concerned about long run consumption levels and short run price fluctuations, and we call this learning mechanism “N-step Euler equation learning”.

N-step Euler equation learning behavioral assumption. *The N-step Euler equation learning assumption identifying consumption behavior in terms of future forecasts is given by (11).*

To forecast, for example, k_{t+n+1} , agent i must forecast c_{t+n} – an issue we did not encounter when investigating Euler equation learning. One option would be to provide agent i with a forecasting model for aggregate consumption. For simplicity and comparability, we make the alternative assumption that agent i thinks he is “average” and so his best forecast of c_{t+n} is c_{t+n}^i .

It remains to specify how agents form forecasts of c_{t+n}^i . As above we provide agent i with a forecasting model that is linear in aggregate capital: $c_t^i = H^i + A^i k_t$. These assumptions yield the following forecasts:

$$\begin{aligned} E_t^i c_{t+n}^i &= H^i + A^i E_t^i k_{t+n} \\ E_t^i k_{t+n} &= \delta_1 E_t^i c_{t+n-1}^i + \delta_2 E_t^i k_{t+n-1} = H^i S_n(A^i) + (\delta_1 A^i + \delta_2)^{n-1} (\delta_1 c_t + \delta_2 k_t) \\ E_t^i r_{t+n} &= B E_t^i k_{t+n}, \end{aligned}$$

where

$$S_n(A^i) = \delta_1 \sum_{m=0}^{n-2} (\delta_1 A^i + \delta_2)^m.$$

These forecasts may be combined with the behavioral equation (11) to determine agent i 's consumption decision: for appropriate functions \hat{C}_j we have

$$c_t^i = \hat{C}_0(A^i) H^i + \hat{C}_1(A^i) c_t + \hat{C}_2(A^i) k_t. \quad (12)$$

Equation (12) determines the behavior of agent i given our implementation of N-step learning. Note that agent i 's behavior depends on his beliefs and on aggregate realizations.

Imposing homogeneity provides the equilibrium dynamics dictated by (12). Let $\psi(A) = \delta_1 A + \delta_2$, and set

$$\gamma_1(A, N) = 1 + A\delta_1 \left(\frac{1 - \psi(A)^{N-1}}{1 - \psi(A)} \right) + \frac{aB\delta_1}{1 - \psi(A)} \left(N - 1 - \psi(A) \left(\frac{1 - \psi(A)^{N-1}}{1 - \psi(A)} \right) \right)$$

$$\gamma_2(A, N) = A\psi(A)^{N-1} + aB \left(\frac{1 - \psi(A)^N}{1 - \psi(A)} \right).$$

Then the T-map is given by

$$A \rightarrow \frac{\delta_2 \gamma_2(A, N)}{1 - \delta_1 \gamma_2(A, N)} \quad (13)$$

$$H \rightarrow \frac{\gamma_1(A, N)H}{1 - \delta_1 \gamma_2(A, N)}. \quad (14)$$

This map may be used to assess stability under N-step Euler equation learning.

4.2 N-step optimal learning

Under N-step Euler equation learning, savings behavior is passive in that it is determined by the budget constraint after the consumption decision is made; and because of this assumption, individual wealth enters into the agent's decision only if it influences the agent's forecasts of either future consumption or future interest rates. An alternative formulation of agent behavior, which we call "N-step optimal learning", takes wealth – both current and expected future values – as central by incorporating the budget constraint into consumption decisions.

To develop N-step optimal learning, set

$$R_t^n = \prod_{k=1}^n (1 + r_{t+k})^{-1},$$

with $R_0^t = 1$. Iterate agent i 's flow budget constraint forward N -periods to get

$$\sum_{n=0}^N R_n^t c_{t+n}^i = \sum_{n=0}^N R_n^t w_{t+n} + (1 + r_t) s_{t-1}^i - R_N^t s_{t+N}^i. \quad (15)$$

Log-linearize (15) and assume agent i makes decisions so that it binds in expectation. Thus agent i 's behavior must satisfy

$$c_t^i + c \sum_{n=1}^N \beta^n E_t^i c_{t+n}^i = \zeta_1 s_{t-1}^i + \zeta_2(N) E_t^i s_{t+N}^i + \sum_{n=0}^N \zeta_2(N, n) E_t^i r_{t+n} + w \sum_{n=0}^N \beta^n E_t^i w_{t+n},$$

for appropriate functions ζ_i .

We now use agent i 's Euler equation iterated forward appropriately to eliminate explicit dependence of consumption today on expected future consumption; this yields the following behavioral equation:⁷

$$c_t^i = \phi_1(N)s_{t-1}^i + \phi_2(N)E_t^i s_{t+N}^i + \sum_{n=1}^N \phi_3(N, n)E_t^i r_{t+n} + \phi_4(N) \sum_{n=1}^N \beta^n E_t^i w_{t+n}. \quad (16)$$

Here

$$\begin{aligned} \beta(N, n) &= \frac{\beta^n}{1 - \beta} (1 - \beta^{N-n+1}) \\ \phi_1(N) &= \frac{s\beta^{-1}}{c\beta(N, 0)} \\ \phi_2(N) &= -\frac{\beta^N s}{c\beta(N, 0)} \\ \phi_3(N) &= \frac{1}{c\beta(N, 0)} \left(\left(cr\beta \left(1 - \frac{1}{\sigma} \right) - wr\beta \right) \beta(N, n) + rs\beta^{N+1} \right) \\ \phi_4(N) &= \frac{w}{c\beta(N, 0)}. \end{aligned}$$

N-step optimal learning behavioral assumption. *The N-step optimal learning assumption identifying consumption behavior in terms of future forecasts and current savings is given by (16).*

To close the model, we must specify how these forecasts are formed. To remain consistent with, and comparable to N-step Euler equation learning, we assume agent i forecasts his future savings as being equal to aggregate capital holdings: $E_t^i s_{t+N}^i = E_t^i k_{t+N+1}$; modeled this way, only a PLM for aggregate consumption is required. As above, assume agent i forecasts k_{t+n} using the known aggregate capital accumulation equation. This requires forecasts of aggregate consumption c_{t+n} . Because agent i is no longer forecasting individual consumption, we provide him with a forecasting model for aggregate consumption: $c_t = H^i + A^i k_t$. Finally, we assume agents know $w_t = \alpha k_t$.

Imposing homogeneity provides the equilibrium dynamics, which yields the fol-

⁷Because the production function has constant returns to scale, the explicit dependence of consumption on current real wage and real interest rates washes out.

lowing T-map:

$$A \rightarrow R(A, N) \left(\phi_1(N) + \delta_2 \theta_N(A) \psi(A)^{N-1} + \delta_2 \sum_{n=1}^{N-1} \eta(N, n) \psi(A)^{n-1} \right) \quad (17)$$

$$H \rightarrow R(A, N) \left(\delta_1 \phi_2(N) + \theta_N(A) S_N(A) + \sum_{n=1}^{N-1} \eta(N, n) S_n(A) \right) \quad (18)$$

where

$$\begin{aligned} \eta(N, n) &= B\phi_3(N, n) + \alpha\phi_4(N)\beta^n \\ \theta_n(A) &= \phi_2(n)\psi(A) + \eta(N, N) \\ R(A, N) &= \left(1 - \delta_1 \theta_N(A) \psi(A)^{N-1} - \delta_1 \sum_{n=1}^{N-1} \eta(N, n) \psi(A)^{n-1} \right)^{-1}. \end{aligned}$$

This map may be used to assess stability under N-step optimal learning.

4.3 Discussion

Two observations concerning N-step Euler equation learning are immediate. First, N-step Euler equation learning is a generalization of Euler equation learning mechanism, as developed by Evans and Honkapohja: indeed, by setting the horizon $N = 1$, the behavioral assumption (11) reduces to (8) and the T-maps (13) and (14) reduce to (9) and (10), respectively. On the other hand, it is not possible to provide an interpretation of N-step Euler equation learning at an infinite horizon. To see this, note that at the rational expectations equilibrium, captured by \hat{A} , the time path for capital is given by $k_{t+1} = \psi(\hat{A})k_t$, and must converge to the steady state; thus, $|\psi(\hat{A})| < 1$. Since $a > 0, B > 0$ and $\delta_1 < 0$, it follows that $\gamma_1(\hat{A}, N) \rightarrow -\infty$ as the planner horizon gets large. This does not overturn stability; however it does prevent identifying an “infinite horizon” version of Euler equation learning, and suggests that the planning horizon may strongly influence the path taken by beliefs – and hence the economy – along a path that converges to the rational expectations equilibrium obtains.

N-step optimal learning is like N-step Euler equation learning in that the behavioral primitive governing N-step optimal learning asserts that agents make decisions today based on forecasts of future prices, and on their own future behavior – consumption in case of Euler equation learning and savings in case of optimal learning;

however, unlike N-step Euler equation learning, N-step optimal learning conditions also on current savings. Second, as suggested by the nomenclature, under N-step optimal learning, the agent is behaving *optimally* conditional on current wealth *and conditional on expected future wealth*, that is, she is behaving as if she is solving an N-period problem with terminal wealth taken as given. Finally, it can be shown that provided beliefs imply $|\psi(A)| < 1$, the T-map given by (17), (18) above converges to the T-map obtained under infinite horizon learning: in this way, N-step optimal learning may be viewed as the finite horizon version of infinite horizon learning.

4.4 Stability under finite horizon learning.

To conduct stability analysis of the Ramsey model’s unique REE under finite horizon learning, we appeal to Evans and Honkapohja’s E-stability principle, and thus examine the Lyapunov stability of the systems of differential equations of the form (7), corresponding either to equations (17) and (18) or to equations (13) and (14). While there is nothing difficult in principle about this type of stability analysis – simply compute $DT - I$ and see if the real parts of the eigenvalues are negative – the dependence of DT on the planning horizon and on the model’s deep parameters is quite complicated, and prevents analytic results. Instead, we rely on numerical analysis, and we obtain the following result:

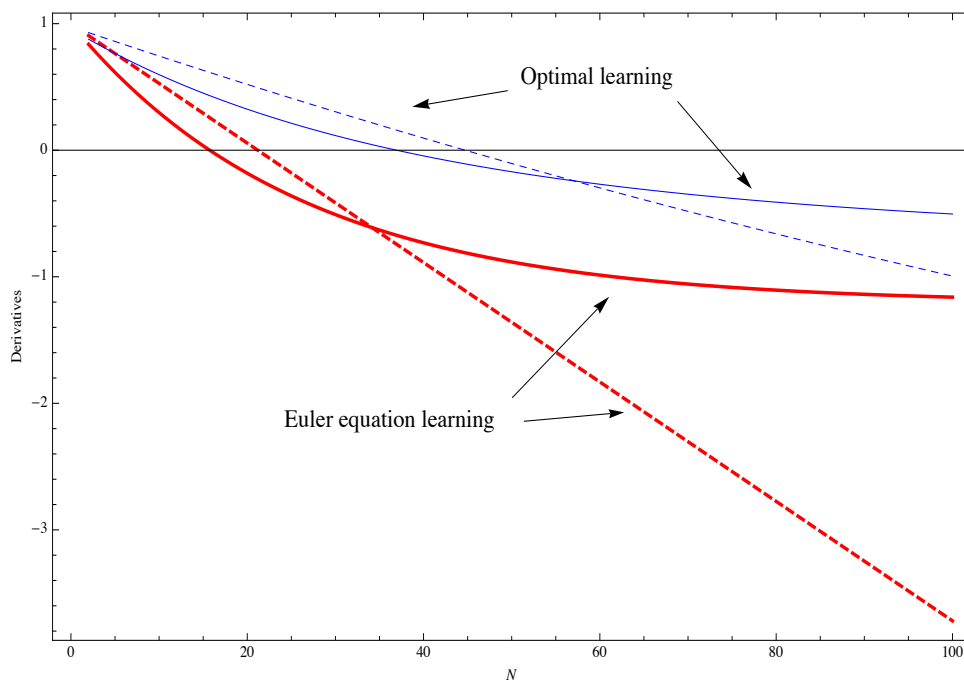
Result. *For all parameter constellations examined, for all planning horizons N , and for both learning mechanisms, the unique REE is E-stable.*

Our numerical result indicates that planning horizon and learning mechanism are irrelevant asymptotically, but they are not pairwise equivalent. While this will be explored in more detail in the next section, we can get a taste for the potential differences here by plotting the derivatives of the T-maps evaluated at the REE. Note that both T-map systems decouple so that the derivatives with respect to A and H may be evaluated separately. For the numerical analysis presented here and throughout the paper we use the standard calibration

$$\alpha = 1/3, \quad \beta = .99, \quad \delta = .025, \quad \sigma = 1.$$

Figure 1 plots DT_A and DT_H for both N-step Euler equation learning and N-step optimal learning, for $N \in \{2, \dots, 100\}$. The solid curves indicate the values of DT_A and the dashed curves indicate values of DT_H . E-stability requires that these eigenvalues have real parts less than one. Notice that while stability obtains for all values of N , the magnitude of the derivatives vary across both horizon length and implementation type. While the connection is not completely understood nor particularly precise, there are formal results and numerical evidence to suggest that, small values of DT

Figure 1: T-map derivatives for N-step Euler and optimal learning: DT_A is solid and DT_H is dashed.



imply faster convergence. In this way, Figure 1 suggests that, Euler equation learning is faster than optimal learning, and longer horizons provide more rapid convergence to the REE.

5 Transition dynamics of finite horizon learning

The behavioral assumption of N-step Euler equation learning implies strong negative feedback for large N: this is evidenced by the failure of N-step Euler equation learning to exist in the limit (as $N \rightarrow \infty$) and by the exploding behavior of DT_H in Figure 1. Intuitively, an agent forecasting above average aggregate consumption for the next N periods (corresponding to the belief $H > 0$) will subsequently forecast low aggregate capital stocks and high real interest rates for these periods as well; high real interest rate forecasts raise the relative price of consumption today and the agent responds by lowering c_t^i . A long planning horizon exacerbates this effect.

The same thought experiment leads to a different intuition for N-step optimal learning. By incorporating the budget constraint into the optimal decision, our use of

log utility washes the income/substitution effect of expected interest rate movements: this may be seen in the expression for $\phi_3(N)$. Interest rates still affect consumption through a wealth effect. Thus an expected decrease in future capital stock, leading to an increase in expected interest rates, reduces the present value of future wage earnings, and thus puts downward pressure on consumption. This effect is compounded by the decrease in expected future wage resulting from the expected decrease in future capital stock. However, both of these effects are mitigated by the expectation that future savings falls: a reduction in expected future savings leads to an increase in consumption today; while this may seem counter-intuitive, remember that the N-step optimal learner is, in effect, solving an N-period planning problem, taking expected future savings as given; a reduction in expected future savings relaxes the agent's constraint and so allows for increased consumption today.

While both learning implementations imply negative feedback for large N , the magnitude of the implied feedback is smaller for N-step optimal learning. Also, given a particular learning implementation – either N-step optimal learning or N-step Euler equation learning – the feedback varies dramatically across planning horizon. These observations suggest that the transition dynamics – the time path of beliefs as convergence to the REE obtains – should vary across both planning horizon and learning implementation. To investigate these possibilities we analyze the different learning algorithms' "mean dynamics".

5.1 Mean dynamics

Let $\Theta = (H, A)'$ capture a representative agent's beliefs. The mean dynamics are given by

$$\begin{aligned}\dot{\Theta} &= S^{-1}M(\Theta)(T(\Theta) - \Theta) \\ \dot{S} &= M(\Theta) - S,\end{aligned}$$

where S is the sample second moment matrix of the regressors $(1, k_t)$ and M is the corresponding population second moment matrix assuming fixed beliefs Θ . Intuitively, the mean dynamics provide an approximation to the expected time-path of beliefs given that agents are using recursive least squares to estimate their forecasting model: for more details, see Evans and Honkapohja (2001).⁸

⁸Because our model is non-stochastic, the matrix M , which captures the asymptotic second moment of the regressors under fixed beliefs in the recursive least squares updating algorithm, must be modified. We follow the ridge regression literature and perturb $M(\Theta)$ by adding εI . For the graphs in this paper, $\varepsilon = .05$. The qualitative features of the analysis are not affected by small changes in ε .

Figure 2: Time-path for beliefs under Euler equation learning: A is solid curve and H is dashed curve.





Figure 2 plots the time-path for beliefs implied by the mean dynamics under N -step Euler equation learning, and given the initial condition $A = .357$, $H = .99$.⁹ The time-paths for A are solid curves and those for H are dashed. The REE value for A is approximately $.6$. We note that the beliefs on capital, A , oscillates as it approaches its REE value; also, while convergence is indicated for all planning horizons, convergence is much faster for larger N .

In Figure 3, we plot the time-path for beliefs implied by the mean dynamics under N -step optimal learning, and for the same initial conditions. As with N -step Euler equation learning, a longer planning horizon results in faster convergence. Also, the oscillatory nature of the time-paths under N -step optimal equation learning are quite similar to N -step Euler equation learning; however, under N -step Euler equation learning, these oscillations largely disappear, whereas they remain for optimal learners.

⁹The matrix S must also be given an initial condition. In a stochastic model, the natural initial condition for this matrix is the regressor's covariance; however, since our model is non-stochastic, our initial condition is necessarily ad-hoc. While the time-paths do depend quantitatively on the initial condition chosen, we found that the qualitative results to be quite robust.

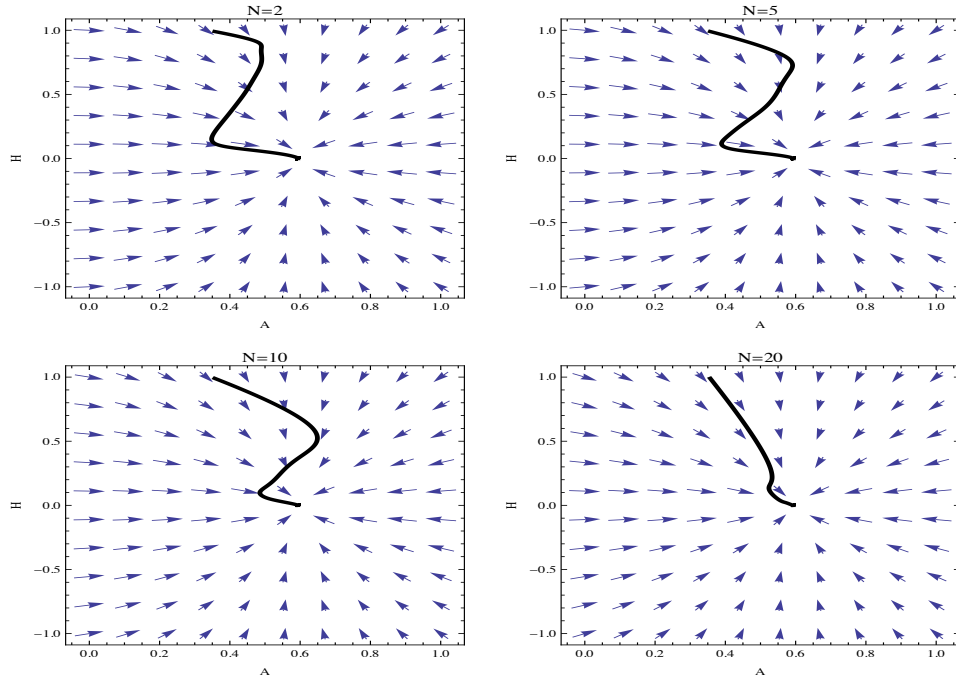
Figure 3: Time-path for beliefs under optimal equation learning: A is solid curve and H is dashed curve.



5.2 Phase plots

The possibility of intriguing oscillatory dynamics is not evident from the T-maps or from the E-stability differential system; and, in fact, the oscillations are caused not by the T-maps themselves, but rather by the interaction of the beliefs with the covariance matrix S . To expose this dichotomy more effectively, consider Figure 4 where we plot, in phase space, the time-path of beliefs under Euler equation learning: see solid curves in the various panels. However, we plot this time-path against the vector field capturing the E-stability differential equation (7). The vector field indicates that the REE is a sink; however, the mean dynamics impart a path for beliefs that, at times, moves away from the REE values, against the direction dictated by the E-stability vector field. The figure indicates the potential importance of using mean dynamics rather than the E-stability differential system to study transition paths.

Figure 4: Time-path for beliefs in phase space under Euler equation learning: vector field given by E-stability differential equation.



5.3 Discussion

The mean dynamics capture the expected transition to the model's rational expectations equilibrium, and the evidence presented in Figures 2 – 4 above indicates that the transition depends on both planning horizon and learning mechanism: short horizon learning models indicate slower convergence and the potential for oscillations in beliefs; these oscillations persist under optimal learning as the planning horizon increases, but under Euler equation learning, the strength of the feedback at long planning horizon dominates and washes out the oscillations.

The distinctive transitional behavior indicated by planning horizon and learning mechanism suggests empirical implications. Coupling finite horizon learning with constant-gain recursive updating, and then embedding these mechanisms in more realistic DSGE models – for example, real business cycle models or New-Keynesian models – may improve fit, better capture internal propagation, and allow for reduced reliance on exogenous shocks with unrealistic, or at least unmodeled, time-series properties.

6 Conclusion

To the extent that DSGE and finance models that embrace the rational expectations hypothesis are unable to account for co-movements in the data, alternative expectation formation mechanisms are clearly a natural focus; and, in the 30 years since birth of the literature, adaptive learning has become rationality's benchmark replacement. Originally, adaptive learning and the corresponding stability analysis was applied either to ad-hoc models or models with repeated, finite horizons; however, micro-founded infinite horizon DSGE models provided a distinct challenge. On the one hand, Euler-equation learning has been offered as a simple behavioral rule providing a boundedly rational justification for examining adaptive learning within one-step ahead reduced-form systems. On the other hand, the principal alternative proposal has been to assume that agents solve their infinite-horizon dynamic optimization problem each period, using current estimates of the forecasting model to form expectations infinitely far into the future. In contrast, introspection and common sense suggests that boundedly rational decision-making is usually based on a finite horizon, the length of which depends on many factors.

This paper has explored a generalization of Euler-equation learning that extends the planning horizon to any finite number of periods. We have also formulated a new type of mechanism – optimal learning – designed explicitly to provide a finite-planning horizon analogue to infinite-horizon learning. The asymptotic stability implications of finite-horizon learning within the Ramsey model are simple to summarize: all roads lead to rationality. This is good news to those researchers hoping to justify the rational expectations hypothesis and to those researchers who have relied on Euler-equation learning, or reduced-form learning, to conduct their stability analysis.

Equally important to the stability analysis, though, are the results on transitional dynamics – results which, under constant-gain learning, would carry over to persistent learning dynamics. Our results indicate that agents' choices are strongly affected by planning horizon. If these results hold in more realistic models then researchers interested in embedding learning agents into fitted DSGE models should consider the planning horizon as a key parameter that needs to be estimated.

References

- BRAY, M. (1982): “Learning, Estimation, and the Stability of Rational Expectations Equilibria,” *Journal of Economic Theory*, 26, 318–339.
- BRAY, M., AND N. SAVIN (1986): “Rational Expectations Equilibria, Learning, and Model Specification,” *Econometrica*, 54, 1129–1160.
- BULLARD, J., AND K. MITRA (2002): “Learning About Monetary Policy Rules,” *Journal of Monetary Economics*, 49, 1105–1129.
- COGLEY, T., AND T. J. SARGENT (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49, 185–221.
- EUSEPI, S., AND B. PRESTON (2010): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, forthcoming.
- 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.
- (2010): “Learning as a Rational Foundation for Macroeconomics and Finance,” mimeo.
- EVANS, G. W., S. HONKAPOHJA, AND K. MITRA (2009): “Anticipated Fiscal Policy and Learning,” *Journal of Monetary Economics*, 56, 930–953.
- EVANS, G. W., AND B. MCGOUGH (2010): “Learning to Optimize,” mimeo.
- GUESNERIE, R. (2005): *Assessing Rational Expectations 2: Eductive Stability in Economics*. MIT Press, Cambridge, Mass.
- HONKAPOHJA, S., K. MITRA, AND G. W. EVANS (2002): “Notes on Agents’ Behavioral Rules Under Adaptive Learning and Recent Studies of Monetary Policy,” manuscript, www.valt.helsinki.fi/RAKA/seppo.html, and <http://darkwing.uoregon.edu/~gevans>.
- JACOBS, D., E. KALAI, AND M. KAMIEN (eds.) (1998): *Frontiers of Research in Economic Theory*. Cambridge University Press, Cambridge.

- KREPS, D. M. (1998): “Anticipated Utility and Dynamic Choice,” in Jacobs, Kalai, and Kamien (1998), pp. 242–274.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,” *Journal of Economic Theory*, 48, 337–368.
- PRESTON, B. (2005): “Learning about Monetary Policy Rules when Long-Horizon Expectations Matter,” *International Journal of Central Banking*, 1, 81–126.
- SARGENT, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford.
- WOODFORD, M. (1990): “Learning to Believe in Sunspots,” *Econometrica*, 58, 277–307.