Learning in Macroeconomics: some recent developments and experimental implications

George W. Evans

University of Oregon and University of St. Andrews

Macro theory and experimental workshop, June 9-10, 2014

Introduction

– Macroeconomic models are usually based on **optimizing** agents in dynamic, stochastic setting and can be summarized by a **dynamic system,** e.g.

$$\begin{aligned} y_t &= Q(y_t^e, w_t) \\ y_t &= Q(y_{t-1}, y_{t+1}^e, w_t) \\ \text{or } y_t &= Q(y_{t-1}, \left\{y_{t+1}^e\right\}_{j=0}^{\infty}, w_t) \end{aligned}$$

$y_t =$ vector of economic variables at time t (unemployment, inflation, investment, etc.), $y_{t+1}^e =$ expectations of these variables, $w_t =$ exogenous random factors at t. Nonstochastic models also of interest.

– The presence of **expectations** y_t^e or y_{t+1}^e, and the assumption that agents can solve dynamic programing problems, makes macroeconomics inherently different from natural science.

– The standard assumption of **rational expectations** (RE) assumes too much **knowledge** & **coordination** for economic agents. We need a **realistic** model of **rationality** What form should this take?

– My general answer is given by the **Cognitive Consistency Principle** (CCP): economic agents should be about as smart as (good) economists, e.g.

model agents like **economic theorists** – the **eductive** approach, or

model them as **econometricians** – the **adaptive** approach.

– We also need to reflect on the optimization assumption. In dynamic stochastic settings the CCP and introspection suggests relaxing the assumption.

– Agents may fall short of the CCP standard but CCP is a good benchmark.

A Muth/Lucas-type Model

Consider a simple univariate reduced form:

$$p_t = \mu + \alpha E^*_{t-1} p_t + \delta' w_{t-1} + \eta_t, \text{ with } \alpha \neq 1. \qquad \text{(RF)}$$

$E^*_{t-1} p_t$ denotes expectations of p_t formed at $t-1$, w_{t-1} is a vector of exogenous observables and η_t is an unobserved iid shock.

Muth cobweb example. Demand and supply equations:

$$\begin{aligned} d_t &= m_I - m_p p_t + v_{1t} \\ s_t &= r_I + r_p E^*_{t-1} p_t + r'_w w_{t-1} + v_{2t}, \end{aligned}$$

$s_t = d_t$, yields (RF) where $\alpha = -r_p/m_p < 0$ if $r_p, m_p > 0$.

Lucas-type monetary model. AS + AD + monetary feedback:

$$\begin{aligned} q_t &= \bar{q} + \lambda(p_t - E^*_{t-1} p_t) + \zeta_t, \\ m_t + v_t &= p_t + q_t \text{ and } m_t = \bar{m} + u_t + \rho' w_{t-1} \end{aligned}$$

leads to yields (RF) with $0 < \alpha = \lambda/(1+\lambda) < 1$.

Adaptive, Least-Squares Learning

The model $p_t = \mu + \alpha E_{t-1} p_t + \delta' w_{t-1} + \eta_t$ has the **unique REE**

$$\begin{aligned} p_t &= \bar{a} + \bar{b}' w_{t-1} + \eta_t, \text{ where} \\ \bar{a} &= (1-\alpha)^{-1}\mu \text{ and } \bar{b} = (1-\alpha)^{-1}\delta. \end{aligned}$$

Special case: If only white noise shocks or the model is nonstochastic then $\delta = 0$. In this case $\bar{b} = 0$ and the REE is $p_t = \bar{a} + \eta_t$, with $E_{t-1} p_t = \bar{a}$.

Under **LS learning**, agents have the beliefs or perceived law of motion (PLM)

$$p_t = a + b w_{t-1} + \eta_t,$$

but a, b are unknown. At the end of time $t-1$ they estimate a, b by LS (Least Squares) using data through $t-1$. Then they use the estimated coefficients to make forecasts $E^*_{t-1} p_t$.

– End of $t-1$: w_{t-1} and p_{t-1} observed. Agents **update estimates** of a, b to a_{t-1}, b_{t-1} and make forecasts

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

– **Temporary equilibrium at t:** (i) p_t is determined as

$$p_t = \mu + \alpha E_{t-1}^{*}p_t + \delta' w_{t-1} + \eta_t$$

and w_t is realized. (ii) agents update estimates to a_t, b_t and forecast

$$E_t^{*}p_{t+1} = a_t + b_t'w_t.$$

The dynamic system under LS learning is written recursively (RLS) as

$$\begin{aligned}
E_{t-1}^{*}p_t &= \phi_{t-1}'z_{t-1} \text{ where } \phi_{t-1}' = (a_{t-1}, b_{t-1}') \text{ and } z_{t-1}' = (1, w_{t-1}) \\
p_t &= \mu + \alpha E_{t-1}^{*}p_t + \delta' w_{t-1} + \eta_t, \\
\phi_t &= \phi_{t-1} + t^{-1}R_t^{-1}z_{t-1}(p_t - \phi_{t-1}'z_{t-1}) \\
R_t &= R_{t-1} + t^{-1}(z_{t-1}z_{t-1}' - R_{t-1}).
\end{aligned}$$

Question: Will $(a_t, b_t) \longrightarrow (\bar{a}, \bar{b})$ as $t \longrightarrow \infty$?

Theorem (Bray & Savin (1986), Marcet & Sargent (1989)). Convergence to RE, i.e. $(a_t, b_t') \rightarrow (\bar{a}, \bar{b}')$ a.s. if $\alpha < 1$. If $\alpha > 1$ convergence with prob. 0.

Thus the REE is stable under LS learning both for Muth model ($\alpha < 0$) and Lucas model ($0 < \alpha < 1$), but is not stable if $\alpha > 1$.

In general models, stochastic approximation theorems are used to prove convergence results. However the **expectational stability** (E-stability) principle, below, gives the stability condition.

Special case: If $\delta = 0$ and agents have the PLM $p_t = a + \eta_t$, then they only regress on an intercept, i.e. $\phi_t = a_t$. The system then is

$$\begin{aligned} p_t &= \mu + \alpha E^*_{t-1} p_t + \eta_t \\ E^*_{t-1} p_t &= a_{t-1} \\ a_t &= a_{t-1} + t^{-1}(p_t - a_{t-1}) \end{aligned}$$

and $a_t \rightarrow \bar{a}$ a.s. if $\alpha < 1$ (a_t does not converge if $\alpha > 1$).

Remark: Speed of convergence depends on α and is slow if $\frac{1}{2} < \alpha < 1$.

E-STABILITY

There is a simple way to obtain the stability condition. Start with PLM

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

and suppose (a, b) were fixed at some (possibly non-REE) value. Then

$$E^*_{t-1} p_t = a + b' w_{t-1},$$

which would lead to the Actual Law of Motion (ALM)

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

The implied ALM gives the mapping T: PLM $\longrightarrow$ ALM:

$$T(a, b) = (\mu + \alpha a, \delta + \alpha b).$$

The REE $\bar{a}, \bar{b}$ is a fixed point of T.

Expectational-stability (“E-stability) is defined by the ODE

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

where τ is notional time. $\bar{a}, \bar{b}$ is **E-stable** if it is stable under this ODE. Here T is linear and the REE is E-stable when $\alpha < 1$.

Intuition: under LS learning the parameters a_t, b_t are slowly adjusted, on average, in the direction of the corresponding ALM parameter.

This technique can be used in nonlinear models, multivariate linear models $y_t = \alpha + \beta E_t^* y_{t+1} + \gamma y_{t-1} + \delta w_t + \varepsilon_t$, and if there are multiple equilibria.

For a wide range of models **E-stability** governs stability under LS learning, e.g. Evans & Honkapohja (2001). This is the **E-stability principle**.

In **NK models** some interest rate rules fail to deliver stability under learning.

Variation 1: constant-gain learning dynamics

– For **discounted LS** the "gain" t^{-1} is replaced by a constant $0 < \gamma < 1$, e.g. $\gamma = 0.04$. Often called "constant gain" (or "perpetual") learning.

– Especially plausible if agents are worried about structural change.

– With constant gain in the Muth/Lucas and $\alpha < 1$ convergence of (a_t, b_t) is to a stochastic process around $(\bar{a}, \bar{b})$.

– In the Cagan/asset-pricing model

$$
\begin{aligned}
p_t &= \mu + \alpha E_t^* p_{t+1} + \delta w_t \\
w_t &= \rho w_{t-1} + \varepsilon_t
\end{aligned}
$$

constant gain learning leads to excess volatility, correlated excess return, etc.

– Escape dynamics can also arise.

Special case: If $\delta = 0$, agents have the PLM $p_t = a + \eta_t$, and they use constant-gain learning with gain $0 < \gamma \leq 1$, then (in e.g. the cobweb model)

$$E_{t-1}^* p_t = a_{t-1} \text{ and } a_t = a_{t-1} + \gamma(p_t - a_{t-1}),$$

which is equivalent to

$$E_t^* p_{t+1} = E_{t-1}^* p_t + \gamma \left(p_t - E_{t-1}^* p_t \right).$$

This, of course, is simply "**adaptive expectations**" with AE parameter γ. Thus AE is a special case of LS learning with constant gain in which the only regressor is an intercept.

Variation 2: misspecified models

- Actual econometricians make specification errors. What happens if our agents make such errors, e.g. underparameterization of the list of regressors or underparameterization of dynamics?

- LS learning still converges if a modified E-stability condition is met, but convergence is to a **Restricted Perceptions Equilibrium** (RPE).

- In an RPE agents are using the best (minimum MSE) econometric model within the class they consider.

Variation 3: heterogeneous expectations

In practice, there is heterogeneity of expectations across agents. This arises, at least in transitional learning dynamics, if different agents have:
– different initial expectations (priors)
– different (possibly random) gains
– asynchronous updating of estimates

Heterogeneity also arises from dynamic predictor selection (Brock & Hommes)
– alternative heuristic forecasting models with discrete choice ('behavioral rationality,' Hommes)
– alternative econometric forecasting models, including misspecified models

General Implications of Adaptive Learning Theory

- Can assess **plausibility** of RE based on stability under LS learning

- Use local stability under learning as a **selection criterion** in models with **Multiple Equilibria**
  - Multiple steady states in nonlinear models
  - Cycles and sunspot equilibria (SSEs) in nonlinear models
  - Sunspot equilibria in linear models with indeterminate steady states

- **Persistent learning dynamics** arise with modified adaptive learning rules

- **Policy implications**: Policy should facilitate learning by private agents of the targeted REE.

Some Experimental Implications, Examples and Discussion

- Although adaptive learning is usually examined in models with exogenous random shocks, experiments often leave these out.

- The most testable proposition seems then to be the stability criterion.

- In an OG model, Marimon, Spear and Sunder (JET1993) provided some experimental evidence of E-stable 2-cycles and nearby 2-SSEs.

- Evans, Honkapohja and Marimon (MD, 2001), "Convergence in Inflation Models with Heterogeneous Learning Rules," considered an OG hyperinflation model with additional fiscal constraint.

- The usual (dynamic) seigniorage model

$$\pi_t = \frac{a - b\pi_t^e}{a - b\pi_{t+1}^e - d}$$

has two steady states (with $d > 0$ small) and none if $d > 0$ is large.

- We add random shocks and a fiscal constraint: $d_t \leq \lambda \times GDP$. There is then one E-stable steady state for d large. For d small there are three steady states, with the high and low steady states E-stable.

- Experimental results generally in line with the theory. Note tightening of the fiscal constraint in regime 4 led to it not binding.

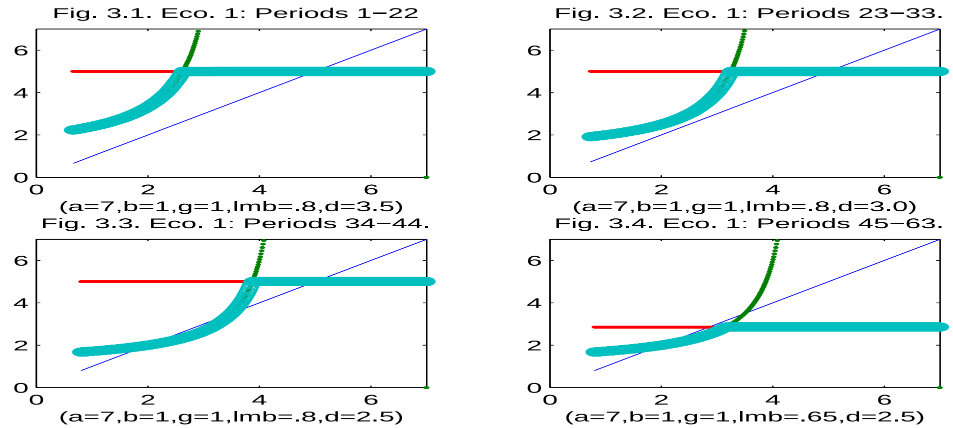
Fig. 3.1. Eco. 1: Periods 1−22
(a=7,b=1,g=1,lmb=.8,d=3.5)
Fig. 3.2. Eco. 1: Periods 23−33.
(a=7,b=1,g=1,lmb=.8,d=3.0)
Fig. 3.3. Eco. 1: Periods 34−44.
(a=7,b=1,g=1,lmb=.8,d=2.5)
Fig. 3.4. Eco. 1: Periods 45−63.
(a=7,b=1,g=1,lmb=.65,d=2.5)

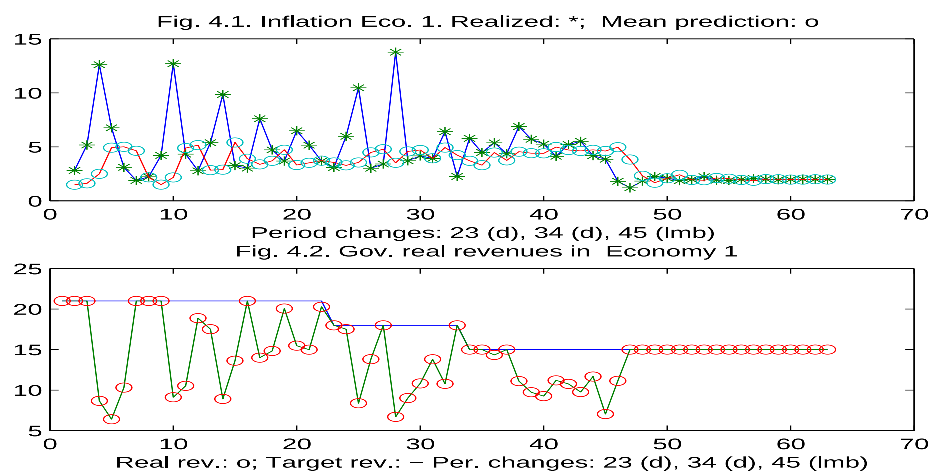
Fig. 4.1. Inflation Eco. 1. Realized: *; Mean prediction: o
Period changes: 23 (d), 34 (d), 45 (lmb)
Fig. 4.2. Gov. real revenues in Economy 1
Real rev.: o; Target rev.: − Per. changes: 23 (d), 34 (d), 45 (lmb)

- Arifovic, Evans and Kostyshyna, "Are Sunspots Learnable? An Experimental Investigation in a Simple Macro Model" (2014) consider a repeated static economy with a positive production externality and three steady states. The key equation is

$$n_t^i = \left(\psi(\overline{N_t}^{e,i})/(2w)\right)^2$$

where $\psi(\cdot)$ captures the production externality.

- Two steady state are E-stable (locally stable under learning). So are 2-state Markov SSEs that fluctuate between these two steady states.

- Experimental results indicate that coordination on any of these locally stable solutions can arise.

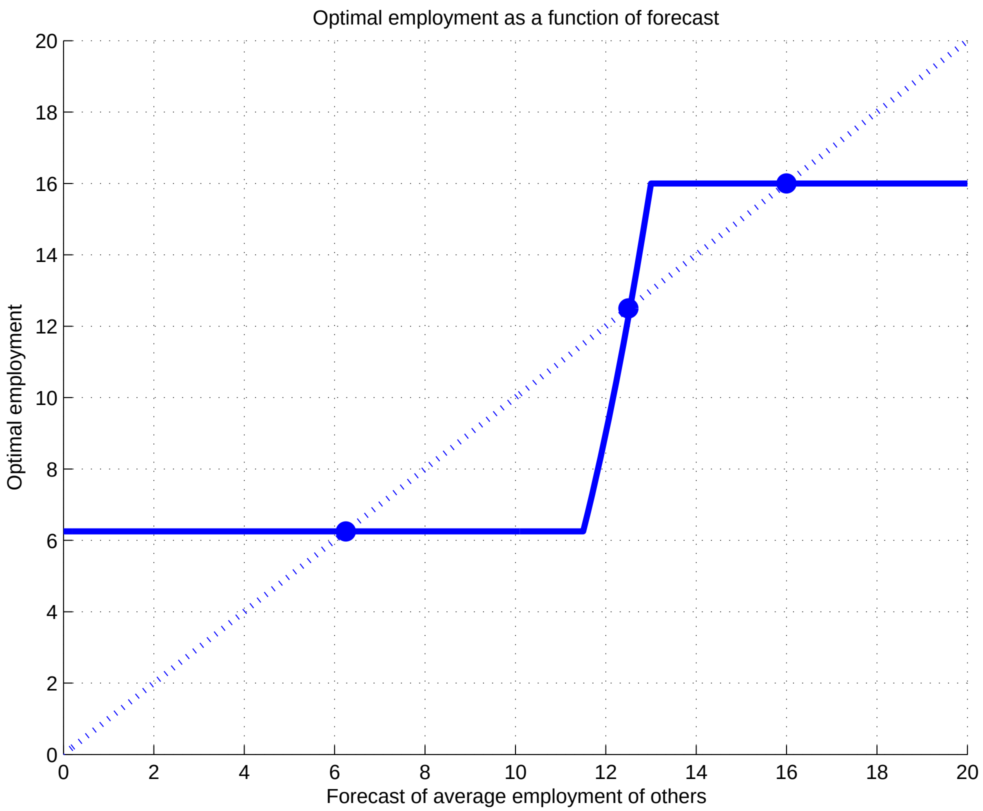

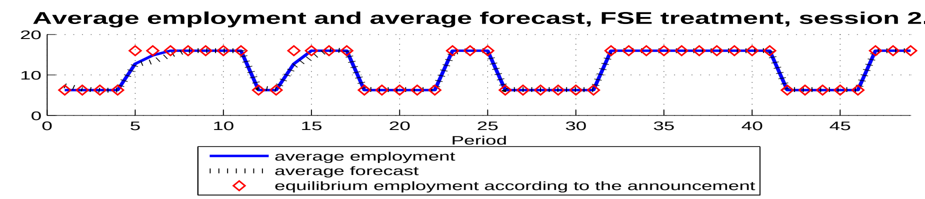
Average employment and average forecast, FSE treatment, session 2.
20
10
0
0
5
10
15
20
25
30
35
40
45
Period
average employment
average forecast
equilibrium employment according to the announcement

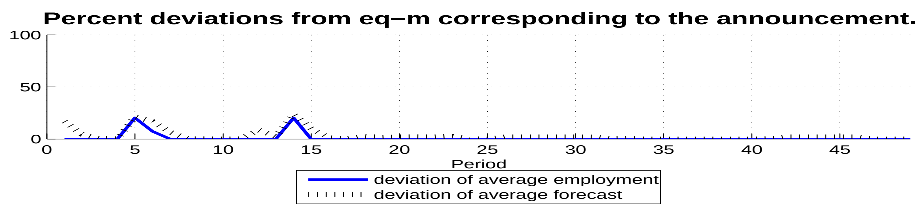
Percent deviations from eq−m corresponding to the announcement.
100
50
0
0
5
10
15
20
25
30
35
40
45
Period
deviation of average employment
deviation of average forecast

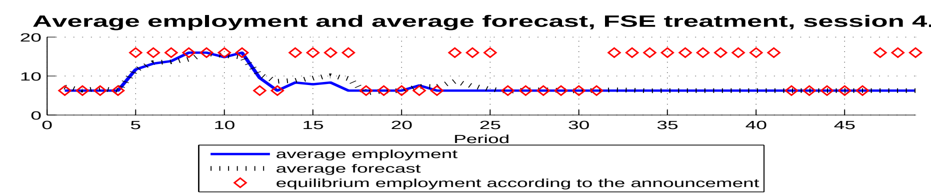
Average employment and average forecast, FSE treatment, session 4.
20
10
0
0
5
10
15
20
25
30
35
40
45
Period
average employment
average forecast
equilibrium employment according to the announcement

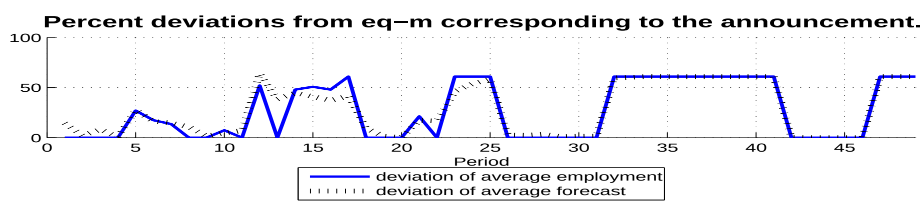
Percent deviations from eq−m corresponding to the announcement.
100
50
0
0
5
10
15
20
25
30
35
40
45
Period
deviation of average employment
deviation of average forecast

- Experiments reported in Hommes (JEDC, 2011) emphasize the importance of positive vs negative feedback.

  – negative feedback in cobweb model with $-1 < \alpha < 0$ leads to quick convergence to REE
  – negative feedback in cobweb model with $\alpha \ll -1$ may fail to converge to REE

  – positive feedback in asset pricing/Cagan model with $\alpha < 1$ near 1 can lead to coordination on different patterns.

- The case $\alpha < 1$ near 1 under adaptive learning with multiple misspecified models can lead to regime-switching (Branch and Evans RED 2007). See also Bullard, Evans and Honkapohja (2008, 2010), Branch & Evans (2011) and Evans, Honkapohja, Sargent and Williams (2013).

- A major issue in experiments is whether subjects are asked to **make decisions** or just to **make forecasts**. Many macro experiments are designed as "learning to forecast" experiments: the experimenter makes the optimal decision for the agent, given their forecast.

- Another major issue is how much **detailed structural information** to give to subjects. Do they know the full structure of the model (and given time to analyze it), or are they given only key qualitative features.

- A major unanswered question for experiments: what about forecasting when there are **exogenous shocks**? Experiments tend to eliminate them. But most macroeconomists believe exogenous random shock are important.

Recent Methodological Issues: The Planning Horizon

- In the Muth-Lucas model outcomes depend on one step ahead forecasts. Also true for Cagan and OG models in which $p_t = Q(p^e_{t+1}, w_t)$.

- However, in standard macro models (e.g. RBC or NK models) agents have long (infinite) lives and usually have corresponding planning horizons.

- Approaches when agents have long lives:
  (i) "internal rationality" – Bayesian updating + full optimization
  (ii) "shadow-price" learning or "Euler-equation" adaptive learning
  (iii) "infinite-horizon" adaptive learning.

- Approaches based on finite planning horizons are also possible.

Bayesian Decision-Making and Internal Rationality

- The standard stochastic dynamic programing problem takes the form

$$
\begin{aligned}
V(x_0) &= \max E_0 \sum_{t \geq 0} \beta^t r(x_t, u_t) \\
\text{subject to } x_{t+1} &= g(x_t, u_t, \varepsilon_{t+1}),\ x_0 \text{ given.}
\end{aligned}
$$

- Fully optimal decision-making assumes g known and that agents can solve for the optimal rule $u_t = h(x_t)$ for controls $u_t \in \Gamma(x_t) \subset \mathbb{R}^m$ as a function of the state $x_t \in \mathbb{R}^n$, e.g. by solving Bellman's equation.

- Cogley and Sargent (IER 2008) show how to solve a finite-horizon permanent income problem using Bayesian decision-making if instead g depends in part on an exogenous 2-state Markov process with unknown parameters.

- Adam and Marcet (JET 2011) develop an "internal rationality" approach in an asset-pricing set-up with risk-neutral agents.

- Agents solve their decision problem using conditional distributions obtained from a well-defined system of subjective probability beliefs about variables exogenous to their decisions. These beliefs need not be consistent with the external truth.

- Under internal rationality, agents in general are sophisticated Bayesian decision-makers and are able to solve difficult dynamic programming problems.

Shadow-Price Learning

Evans and McGough (2014) “Learning to Optimize” shows that a boundedly rational agent can learn to optimize using perceived first-order conditions combined with LS learning (a “short-horizon” decision rule).

- This is a “bounded optimality” approach.

- Consider the standard linear-quadratic (LQ) dynamic programing problem for a single agent: given initial state x_0,

$$\begin{aligned} V(x_0) = \max \quad & -E_0 \sum \beta^t \left(x_t' R x_t + u_t' Q u_t + 2 x_t' W u_t \right) \\ s.t. \quad & x_{t+1} = A x_t + B u_t + C \varepsilon_{t+1}, \end{aligned}$$

  A simple example is an LQ Robinson Crusoe economy. Hansen and Sargent (2014) give many examples of LQ economies.

- Under standard conditions the optimal decision rule is $u_t = F^* x_t$, where F^* depends on the solution P to the Riccati equation.

- Solving the Riccati equation is generally only possible numerically. Obtaining F^* requires advanced knowledge and computational skills.

- We replace RE and full optimality with (i) adaptive learning and (ii) bounded optimality, based on (iii) the Lagrangian approach

$$\begin{aligned}\mathcal{L} &= E_0 \sum \beta^t \{-x_t' R x_t - u_t' Q u_t - 2x_t' W u_t + \\ &\quad \lambda_t' (A x_{t-1} + B u_{t-1} + C \varepsilon_t - x_t)\}.\end{aligned}$$

- Agents replace (A, B) with (A_t, B_t), estimated and updated by recursive LS (RLS), and replace λ_t with perceived shadow prices λ_t^*.

- The perceived FOCs are

$$\begin{aligned} u_t &= -Q^{-1}W'x_t + (\beta/2)Q^{-1}B_t'\hat{E}_t\lambda_{t+1}^* \\ \lambda_t^* &= -2Rx_t - 2Wu_t + \beta A_t'\hat{E}_t\lambda_{t+1}^*. \end{aligned}$$

  Both are standard 2-period marginal equations, e.g. the first weighs the marginal cost of a unit of u_t against its expected discounted benefit in producing more state x_{t+1}.

- Finally we specify how agents make forecasts $\hat{E}_t\lambda_{t+1}^*$. Under optimal decisions $\lambda_t = H^*x_t$ so we assume agents use RLS to estimate

$$\begin{aligned} \lambda_t^* &= Hx_t + \mu_t, \text{ which gives} \\ \hat{E}_t\lambda_{t+1}^* &= H_t(A_tx_t + B_tu_t). \end{aligned}$$

  Given x_t, A_t, B_t the above equations determine $u_t, \hat{E}_t\lambda_{t+1}^*$ and λ_t^*

- Together with RLS equations for parameter updating, this fully describes SP-learning as a recursive system.

- **Theorem** *Under standard assumptions, and assuming a suitable projection facility, then under SP-learning* (H_t, A_t, B_t) *converges to* (H^*, A, B) *almost surely.*

- This is a striking result: decisions converge asymptotically to the fully rational optimal solution. By estimating shadow prices, we have converted an infinite-horizon problem into a two-period optimization problem.

- Although we prove asymptotic optimality for the LQ set-up, SP learning can be applied to the general dynamic programing set-up.

- In special cases SP-learning can be shown to be equivalent to "Euler-equation learning." Value-function learning is an alternative implementation of our results.

- SP learning can be embedded in general equilibrium models. The paper illustrates this for the stochastic Ramsey model.

- An advantage of SP-learning: agents need only solve 2-period optimization problems using one-step ahead forecasts of states and shadow prices. SP-learning is simple, general and intuitive.

- Experimental tests of SP-learning would be desirable.

Infinite-Horizon (IH) Learning

- This approach has been stressed by Preston, IJCB (2005), JME (2006), and used in by Eusepi and Preston, AEJMac (2010), AER (2010). See also Marcet and Sargent (1989) and Sargent (1993).

- In IH ('optimal') learning at each t agents make fully optimal decisions given long-horizon forecasts for variables outside their control. Their TVC and any IBC is explicitly incorporated.

- Agents do not take into account that their forecast rules under learning change over time (the "anticipated utility" approach of Kreps (1998)).

- If the structure is not stationary, e.g. due to an expected future change in macro policy, this can be incorporated.

Ramsey Model example

To illustrate consider a discrete-time non-stochastic Ramsey model

$$\max \; E_t^*\{\sum_{s=t}^{\infty} \beta^{s-t} \frac{c_s^{1-\sigma}}{1-\sigma}\} \text{ s.t. } a_{s+1} = w_s + r_s a_s - c_s - \tau_s, \text{ for all } s \geq t,$$

where $a_s = k_s + b_s$ and r_s is the real rate of return factor. Assume initially a balanced budget with constant $g_t = g$ and no debt.

The Euler equation is $c_t^{-\sigma} = \beta E_t^*(r_{t+1} c_{t+1}^{-\sigma})$. In this set-up **SP-learning** is "equivalent" to **Euler-equation** (EE) **learning**. With point expectations

$$c_t = \beta^{-\frac{1}{\sigma}} \left(r_{t+1}^e(t) \right)^{-\frac{1}{\sigma}} c_{t+1}^e(t).$$

Under EE-learning agents use adaptive learning to form forecasts of $r_{t+1}^e(t)$ and of **their own consumption** next period, $c_{t+1}^e(t)$.

IH-learning

Under IH learning, agents fully optimize using their IBC and TVC. The Euler equation with point expectations implies $c^e_{t+j}(t) = c_t \beta^{\frac{j}{\sigma}} (\prod_{i=1}^{j} r^e_{t+i}(t))^{\frac{1}{\sigma}}$. Substituting into the IBC (with $\sigma = 1$) gives

$$c_t = (1-\beta) \times \left(r_t k_t + w_t - \tau_t + \sum_{j=1}^{\infty} (\prod_{i=1}^{j} r^e_{t+i}(t))^{-1} (w^e_{t+j}(t) - \tau^e_{t+j}(t)) \right)$$

Forecasts $w^e_{t+j}, r^e_{t+1}(t)$ are obtained from simple adaptive learning rules. Given these and k_t this equation gives temporary equilibrium c_t and hence k_{t+1}.

Together with learning this gives the path of the economy. The REE in the Ramsey model is stable under IH learning.

Anticipated Fiscal policy

Eusepi & Preston emphasize that knowledge of the monetary policy rule helps stabilize the economy.

Evans, Honkapohja & Mitra JME (2009) use IH-learning to incorporate anticipated fiscal changes.

A hallmark of RE is that announced future policies have an impact now. This also happens with IH learning. EHM (2009) show the impact in the Ramsey model of an announced future permanent increase in government spending.

$$g_t = g_0 = \tau_0 \text{ when } t < T_p \text{ and } g_t = g_1 = \tau_1 \text{ for } t \geq T_p.$$

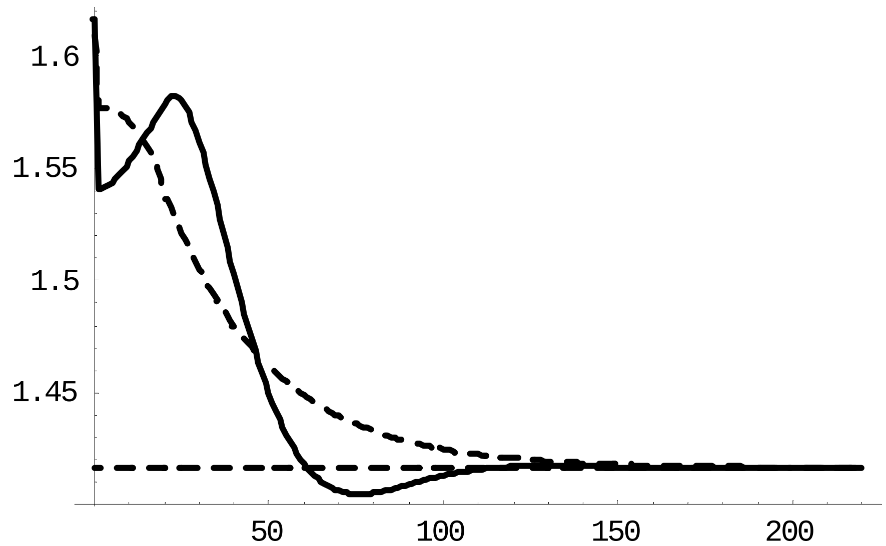

c_t dynamics under learning (solid curve) and perfect foresight (dashed curve). Straight dashed line is new steady state for c. $T_p = 20$.

Immediate impact due to the understanding by agents that future taxes will be higher. Learning dynamics differ from RE because agents do not know GE effects and use adaptive learning to forecast w_{t+i} and r_{t+i}.

If same policy change were repeated many times, agents could eventually learn RE, but policy changes typically have unique features.

RBC models: business cycle fluctuations

Eusepi-Preston, AER (2010) look at equilibrium fluctuations in a an RBC model with IH-learning. As above c_t depends on forecasted wages and rental rates for capital. Forecasts are made using a linear model, with coefficients updated over time using constant-gain LS learning.

The basic findings for a calibrated learning model (compared to RE):
(i) Same output volatility with smaller technology shocks.
(ii) More volatility in investment and hours.
(iii) Captures persistence in investment and hours.

Adding IH-learning to the RBC model improves fit to the data. Key mechanism: partially self-fulfilling shifts in expectations arising from technology shocks that generate temporary but persistent movements in estimated coefficients.

Other applications of IH-learning

Eusepi & Preston have NK applications showing the importance of communication of monetary policy and the role of debt.

Gasteiger and Zhang JEDC (2013) have looked at anticipated policy in models with labor-tax distortions.

Mitra, Evans and Honkapohja JEDC (2013) and other work consider announced fiscal policy changes and spending multipliers in an RBC model.

Benhabib, Evans and Honkapohja JEDC (in press), "Liquidity Traps and Expectation Dynamics: Fiscal Stimulus or Fiscal Austerity?" use IH learning in an NK model with multiple steady states due to the ZLB.

EHM have experiments planned to contrast RE with IH learning.

Benhabib, Evans and Honkapohja (JEDC, in press)

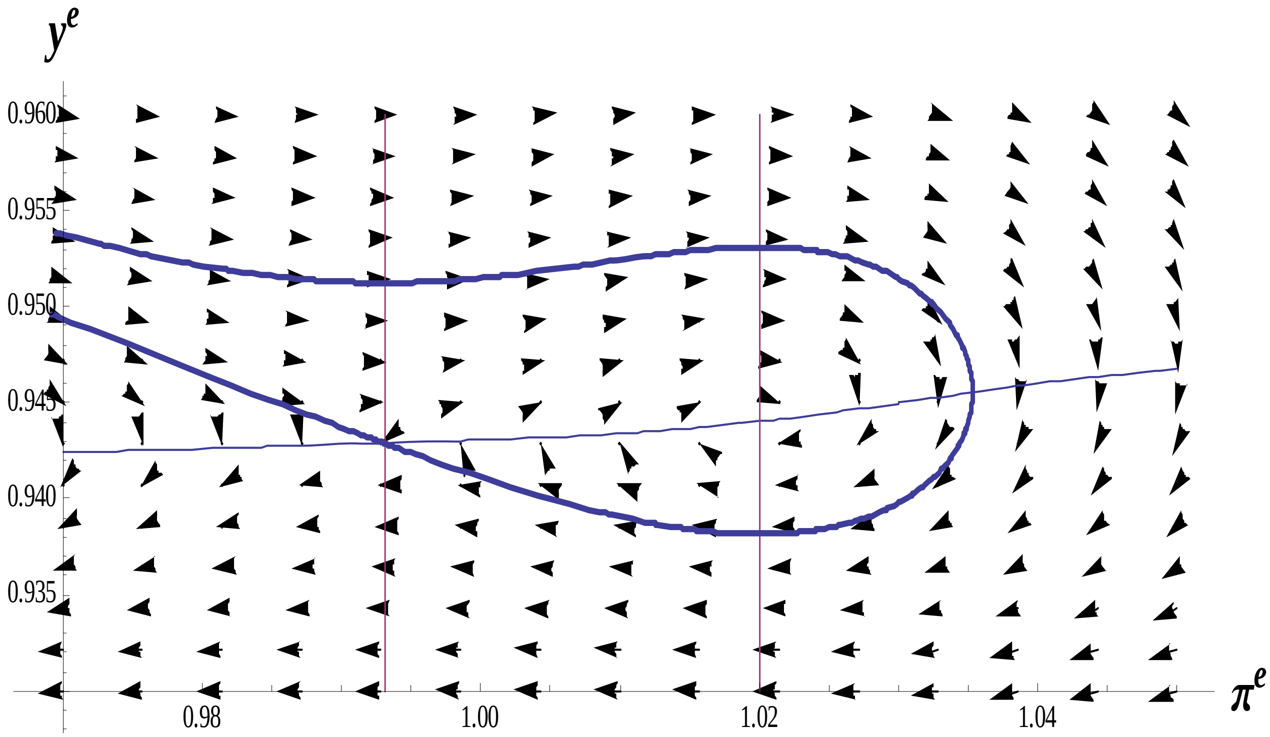

Global learning dynamics – the Ricardian case.

EDUCTIVE LEARNING

LS learning is adaptive: it occurs in real time in response to data.

Eductive learning occurs in mental time using common knowledge (CK) and reasoning. See e.g. Guesnerie (1992, 2002), Evans and Guesnerie (1993, 2005).

Consider the non-stochastic cobweb model with firms $\omega \in [0,1]$:

$$\begin{aligned} d_t &= m_I - m_p p_t \text{ with } s_t(\omega) = r_I + r_p p_t^e(\omega) \\ s_t &= \int_0^1 s_t(\omega) d\omega, \text{ and } d_t = s_t. \\ &\longrightarrow p_t = \mu + \alpha \int_0^1 p_t^e(\omega) d\omega, \end{aligned}$$

where $\alpha = -m_p^{-1} r_p < 0$ if $m_p, r_p > 0$. The REE is $p_t = \bar{a} \equiv (1-\alpha)^{-1}\mu$.

The cobweb model is $p_t = \mu + \alpha \int_0^1 p_t^e(\omega) d\omega$.

The **eductive argument**: Suppose it is CK that agents are rational and know the structure, and it is also CK that $p_t \in V(\bar{p})$. If $|\alpha| < 1$ agents can deduce that $p_t = \bar{p}$: $p_t^e(\omega) \in V(\bar{p})$ implies $p_t \in |\alpha| V(\bar{p})$ and this is CK. If $|\alpha| < 1$ this can be iterated. Hence $p_t = \bar{p}$ is CK and we say $p_t = \bar{p}$ is **eductively stable**. If $|\alpha| > 1$ the argument fails: $p_t = \bar{p}$ is not eductively stable.

– If $-1 < \alpha < 1$ the economy is both adaptively and eductively stable; if $\alpha < -1$ it is adaptively stable but not eductively stable.

– The stricter condition $|\alpha| < 1$ can be obtained from iterating the T-map

$$a_{N+1} = T(a_N), \text{ for } N = 0, 1, 2, \ldots, \text{ where here } T(a) = \mu + \alpha a.$$

In general "Iterative E-stability" is a necessary condition for eductive stability.

Eductive vs adaptive learning

- Eductive stability models agents as theorists and assumes CK of known structure and (hyper)rationality of agents.
  – If an REE is eductively stable, learning can be instantaneous.
  – In practice, allowing for some lack of CK, one might still expect more rapid coordination on the REE than under adaptive learning.

- What if an economy is not eductively stable?
  – Then we would expect coordination to be difficult.
  – Guesnerie argues that then there exist plausible adaptive learning rules that fail to converge.
  – It seems, at the very least, one would expect less rapid coordination on the REE than under adaptive learning.

- Experimental papers relevant to eductive vs. adaptive learning.
  – Hommes (JEDC, 2011) finds lack of coordination in the cobweb model when $\alpha < -1$.
  – Bao and Duffy (2013) focus on the eductive/adaptive issue. They find convergence when $\alpha = -2$, but convergence is slow. When $-1 < \alpha < 0$ convergence is faster than one would expect under adaptive learning.
  – Of course also relevant: "guess the average" experiments, Nagel (1995).

- Perhaps experimental results reflect a mixture of eductive and adaptive learning. It is likely that structural knowledge, setting, information, etc., may be crucial.

- Evans, Guesnerie and McGough (2014) show the RBC model not strongly eductively stable, though it is stable under many adaptive learning rules. Thus the eductive/adaptive issue will be important going forward.

Conclusions and Issues for Experimental Macroeconomics

- The CCP suggests the need for learning and bounded rationality.

- Least-squares and adaptive learning in macro is a flexible and well-developed general framework. Behavioral and eductive models are also available.

- Even in simple set-ups, various specific settings should be examined in experiments, controlling for knowledge of structure, information variables, forecasting tools, multiple equilibria, etc.

- One long-standing question in experiments concerns the relative importance of bounded rationality in forecasting vs optimization.

- An issue not much addressed in experiments is models with serially correlated observable shocks.

- Do adaptive subjects choose between simple behavioral rules or use statistical forecasting rules?

- An important issue in macro learning is the decision horizon:
  (i) How do boundedly rational agents solve dynamic programing problems?
  (ii) How do they incorporate knowledge of future structural changes like announced changes in policy?

- Do agents behave adaptively, eductively, or use a mixture?

- Experiments should be able to shed light on all these issues.