

Learning in Macroeconomic Models

Wouter J. Den Haan
London School of Economics

© Wouter J. Den Haan

December 14, 2014

Overview

- A bit of history of economic thought
- How expectations are formed can matter *in the long run*
 - Seignorage model
- Learning without feedback
- Learning with feedback
 - Simple adaptive learning
 - Least-squares learning
 - Bayesian versus least-squares learning
 - Decision theoretic foundation of Adam & Marcet

Overview continued

Topics

- Learning & PEA
- Learning & sunspots

Why are expectations important?

- Most economic problems have intertemporal consequences
 - \implies future matters
- Moreover, future is uncertain
- Characteristics/behavior other agents can also be uncertain
 - \implies expectations can also matter in one-period problems

History of economic thought

- adaptive expectations:

$$\hat{E}_t [x_{t+1}] = \hat{E}_{t-1} [x_t] + \omega (x_t - \hat{E}_{t-1} [x_t])$$

- very popular until the 70s

History of economic thought

problematic features of adaptive expectations:

- agents can be systematically wrong
- agents are completely passive:
 - $\hat{E}_t [x_{t+j}] , j \geq 1$ only changes (at best) when x_t changes
 - \implies Pigou cycles are not possible
 - \implies model predictions *underestimate* speed of adjustment (e.g. for disinflation policies)

History of economic thought

problematic features of adaptive expectations:

- adaptive expectations about $x_{t+1} \neq$ adaptive expectations about Δx_{t+1}
 - (e.g. price level versus inflation)
- why wouldn't (some) agents use existing models to form expectations?
- expectations matter but still no role for randomness (of future realizations)
 - so no reason for buffer stock savings
 - no role for (model) uncertainty either

History of economic thought

rational expectations became popular because:

- agents are no longer passive machines, but forward looking
 - i.e., agents *think* through what could be consequences of their own actions and those of others (in particular government)
- consistency between model predictions and of agents being described
- randomness of future events become important
 - e.g., $E_t [c_{t+1}^{-\gamma}] \neq (E_t [c_{t+1}])^{-\gamma}$

History of economic thought

problematic features of rational expectations

- agents have to know *complete* model
 - make correct predictions about all possible realizations
 - on *and* off the equilibrium path
- costs of forming expectations are ignored
- how agents get rational expectations is not explained

History of economic thought

problematic features of rational expectations

- makes analysis more complex
 - behavior this period depends on behavior tomorrow for all possible realizations
 - \implies we have to solve for policy *functions*, not just simulate the economy

Expectations matter

- Simple example to show that how expectations are formed can matter in the long run
 - See Adam, Evans, & Honkapohja (2006) for a more elaborate analysis

Model

- Overlapping generations
- Agents live for 2 periods
- Agents save by holding money
- No random shocks

Model

$$\max_{c_{1,t}, c_{2,t}} \ln c_{1,t} + \ln c_{2,t}$$

s.t.

$$c_{2,t} \leq 1 + \frac{P_t}{P_{t+1}^e} (2 - c_{1,t})$$

no randomness \implies we can work with expected value of variables instead of expected utility

Agent's behavior

First-order condition:

$$\frac{1}{c_{1,t}} = \frac{P_t}{P_{t+1}^e} \frac{1}{c_{2,t}} = \frac{1}{\pi_{t+1}^e} \frac{1}{c_{2,t}}$$

Solution for consumption:

$$c_{1,t} = 1 + \pi_{t+1}^e / 2$$

Solution for real money balance (=savings):

$$m_t = 2 - c_{1,t} = 1 - \pi_{t+1}^e / 2$$

Money supply

$$\overline{M}_t^s = \overline{M}$$

Equilibrium

Equilibrium in period t implies

$$\bar{M} = M_t$$

$$\bar{M} = P_t (1 - \pi_{t+1}^e / 2)$$

$$P_t = \frac{\bar{M}}{1 - \pi_{t+1}^e / 2}$$

Equilibrium

Combining with equilibrium in period $t - 1$ gives

$$\pi_t = \frac{P_t}{P_{t-1}} = \frac{1 - \pi_t^e/2}{1 - \pi_{t+1}^e/2}$$

Thus: π_t^e & $\pi_{t+1}^e \implies$ money demand \implies actual inflation π_t

Rational expectations solution

Optimizing behavior & equilibrium:

$$\frac{P_t}{P_{t-1}} = T(\pi_t^e, \pi_{t+1}^e)$$

Rational expectations equilibrium (REE):

$$\pi_t = \pi_t^e$$

$$\implies$$

$$\pi_t = T(\pi_t, \pi_{t+1})$$

$$\implies$$

$$\pi_{t+1} = 3 - \frac{2}{\pi_t}$$

$$\pi_{t+1} = R(\pi_t)$$

Multiple steady states

- There are two solutions to

$$\pi = 3 - \frac{2}{\pi}$$


⇒ there are two steady states

- $\pi = 1$ (no inflation) and perfect consumption smoothing
- $\pi = 2$ (high inflation), money has no value & no consumption smoothing at all

Unique solution

- Initial value for π_t not given, but given an initial condition the time path is fully determined
- π_t converging to 2 means m_t converging to zero and P_t converging to ∞

Rational expectations and stability



Rational expectations and stability

π_1 : value in period 1

$\pi_1 < 1$: divergence

$\pi_1 = 1$: economy stays at low-inflation steady state

$\pi_1 > 1$: convergence to high-inflation steady state


Alternative expectations

- Suppose that

$$\pi_{t+1}^e = \frac{1}{2}\pi_{t-1} + \frac{1}{2}\pi_t^e$$

- still the same two steady states, but
 - $\pi = 1$ is stable
 - $\pi = 2$ is not stable

Adaptive expectations and stability



Learning without feedback

Setup:

- 1 Agents know the complete model, except they do **not** know *dgp* exogenous processes
- 2 Agents use observations to update beliefs
- 3 Exogenous processes do not depend on beliefs
⇒ no feedback from learning to behavior of variable being forecasted

Learning without feedback & convergence

- If agents can learn the dgp of the exogenous processes, then you typically converge to REE
- They may not learn the correct dgp if
 - Agents use limited amount of data
 - Agents use misspecified time series process

Learning without feedback - Example

- Consider the following asset pricing model

$$P_t = E_t [\beta (P_{t+1} + D_{t+1})]$$

- If

$$\lim_{j \rightarrow \infty} \beta^{t+j} D_{t+j} = 0$$

then

$$P_t = E_t \left[\sum_{j=1}^{\infty} \beta^j D_{t+j} \right]$$

Learning without feedback - Example

- Suppose that

$$D_t = \rho D_{t-1} + \varepsilon_t, \varepsilon_t \sim N(0, \sigma^2) \quad (1)$$

- REE:

$$P_t = \frac{D_t}{1 - \beta\rho}$$

(note that P_t could be negative so P_t is like a deviation from steady state level)

Learning without feedback - Example

- Suppose that agents do not know value of ρ
- Approach here:
 - If period t belief = $\hat{\rho}_t$, then

$$P_t = \frac{D_t}{1 - \beta \hat{\rho}_t}$$

- Agents ignore that their beliefs may change,
 - i.e., $\hat{E}_t [P_{t+j}] = E_t \left[\frac{D_{t+j}}{1 - \beta \hat{\rho}_{t+j}} \right]$ is assumed to equal $\frac{1}{1 - \beta \hat{\rho}_t} E_t [D_{t+j}]$

Learning without feedback - Example

How to learn about ρ ?

- Least squares learning using $\{D_t\}_{t=1}^T$ & correct dgp
- Least squares learning using $\{D_t\}_{t=1}^T$ & incorrect dgp
- Least squares learning using $\{D_t\}_{t=T-\bar{T}}^T$ & correct dgp
- Least squares learning using $\{D_t\}_{t=T-\bar{T}}^T$ & incorrect dgp
- Bayesian updating (also called rational learning)
- Lots of other possibilities

Convergence again

- Suppose that the true dgp is given by

$$\begin{aligned}D_t &= \rho_t D_{t-1} + \varepsilon_t \\ \rho_t &\in \left\{ \rho_{\text{low}}, \rho_{\text{high}} \right\} \\ \rho_{t+1} &= \begin{cases} \rho_{\text{high}} & \text{w.p. } p(\rho_t) \\ \rho_{\text{low}} & \text{w.p. } 1 - p(\rho_t) \end{cases}\end{aligned}$$

- Suppose that agents think the true dgp is given by

$$D_t = \rho D_{t-1} + \varepsilon_t$$

- \implies Agents will never learn
(see homework for importance of sample used to estimate ρ)

Recursive least-squares

- time-series model:

$$y_t = x_t' \gamma + u_t$$

- least-squares estimator

$$\hat{\gamma}_T = R_T^{-1} \frac{X_T' Y_T}{T}$$

where

$$X_T' = \begin{bmatrix} x_1 & x_2 & \cdots & x_T \end{bmatrix}$$

$$Y_T' = \begin{bmatrix} y_1 & y_2 & \cdots & y_T \end{bmatrix}$$

$$R_T = X_T' X_T / T$$

Recursive least-squares

$$R_T = R_{T-1} + \frac{(x_T x_T' - R_{T-1})}{T}$$

$$\hat{\gamma}_T = \hat{\gamma}_{T-1} + \frac{R_T^{-1} x_T (y_T - x_T' \hat{\gamma}_{T-1})}{T}$$

Proof for R

$$\begin{aligned}
 \frac{X'_T X_T}{T} & \stackrel{?}{=} \frac{X'_{T-1} X_{T-1}}{(T-1)} + \frac{x_T x'_T}{T} - \frac{X'_{T-1} X_{T-1}}{T(T-1)} \\
 \left(\frac{T-1}{T}\right) X'_T X_T & \stackrel{?}{=} X'_{T-1} X_{T-1} + \frac{T-1}{T} x_T x'_T - \frac{X'_{T-1} X_{T-1}}{T} \\
 X'_T X_T - \frac{X'_T X_T}{T} & \stackrel{?}{=} X'_{T-1} X_{T-1} + x_T x'_T - \frac{x_T x'_T}{T} - \frac{X'_{T-1} X_{T-1}}{T} \\
 X'_{T-1} X_{T-1} + x_T x'_T & \stackrel{?}{=} X'_{T-1} X_{T-1} + x_T x'_T \\
 - \frac{X'_{T-1} X_{T-1} + x_T x'_T}{T} & \stackrel{?}{=} - \frac{x_T x'_T + X'_{T-1} X_{T-1}}{T}
 \end{aligned}$$

Proof for gamma

$$\begin{aligned}
 (X_T' X_T)^{-1} & \stackrel{?}{=} (X_T' X_T)^{-1} \left((X_{T-1}' X_{T-1})^{-1} X_{T-1}' Y_{T-1} + \right. \\
 \times X_T' Y_T & \left. \begin{matrix} x_T y_T \\ -x_T x_T' (X_{T-1}' X_{T-1})^{-1} X_{T-1}' Y_{T-1} \end{matrix} \right) \\
 X_T' Y_T & \stackrel{?}{=} (X_{T-1}' X_{T-1} + x_T x_T') (X_{T-1}' X_{T-1})^{-1} X_{T-1}' Y_{T-1} \\
 & \quad + \begin{pmatrix} x_T y_T \\ -x_T x_T' (X_{T-1}' X_{T-1})^{-1} X_{T-1}' Y_{T-1} \end{pmatrix} \\
 X_T' Y_T & \stackrel{?}{=} \begin{pmatrix} (I + x_T x_T' (X_{T-1}' X_{T-1})^{-1}) X_{T-1}' Y_{T-1} \\ x_T y_T \\ -x_T x_T' (X_{T-1}' X_{T-1})^{-1} X_{T-1}' Y_{T-1} \end{pmatrix}
 \end{aligned}$$

Reasons to adopt recursive formulation

- makes proving analytical results easier
- less computer intensive,
 - but standard LS gives the same answer
- there are intuitive generalizations:

$$R_T = R_{T-1} + \omega(T) (x_T x_T' - R_{T-1})$$

$$\hat{\gamma}_T = \hat{\gamma}_{T-1} + \omega(T) R_T^{-1} x_T (y_T - x_T' \hat{\gamma}_{T-1})$$

$\omega(T)$ is the "gain"

Learning with feedback

- 1 Explanation of the idea
- 2 Simple adaptive learning
- 3 Least-squares learning
 - E-stability and convergence
- 4 Bayesian versus least-squares learning
- 5 Decision theoretic foundation of Adam & Marcet

Learning with feedback - basic setup

Model:

$$p_t = \rho \hat{E}_{t-1} [p_t] + \delta x_{t-1} + \varepsilon_t$$

RE solution:

$$\begin{aligned} p_t &= \frac{\delta}{1 - \rho} x_{t-1} + \varepsilon_t \\ &= a_{\text{re}} x_{t-1} + \varepsilon_t \end{aligned}$$

What is behind model

Model:

$$p_t = \rho \hat{E}_{t-1} [p_t] + \delta x_{t-1} + \varepsilon_t$$

Stories:

- Lucas aggregate supply model
- Muth market model

See Evans and Honkapohja (2009) for details

Learning with feedback - basic setup

Perceived law of motion (PLM) at $t - 1$:

$$p_t = \hat{a}_{t-1}x_{t-1} + \varepsilon_t \quad (2)$$

Actual law of motion (ALM):

$$p_t = \rho\hat{a}_{t-1}x_{t-1} + \delta x_{t-1} + \varepsilon_t = (\rho\hat{a}_{t-1} + \delta)x_{t-1} + \varepsilon_t \quad (3)$$

Updating beliefs I: Simple adaptive

$$\text{ALM: } p_t = (\rho \hat{a}_{t-1} + \delta) x_{t-1} + \varepsilon_t$$

Simple adaptive learning:

- $\hat{a}_t = \rho \hat{a}_{t-1} + \delta$
- could be rationalized if
 - agents observe x_{t-1} and ε_t
 - t is more like an iteration and in each iteration agents get to observe long time-series to update

Simple adaptive learning: Convergence

$$\hat{a}_t = \rho \hat{a}_{t-1} + \delta$$

or in general

$$\hat{a}_t - \hat{a}_{t-1} = T(\hat{a}_{t-1})$$

Simple adaptive learning: Convergence

Key questions:

- 1 Does \hat{a}_t converge?
- 2 If yes, does it converge to a ?

Answers: If $|\rho| < 1$, then the answer to both is yes.

Updating beliefs: LS learning

Suppose agents use least-squares learning

$$\begin{aligned}\hat{a}_t &= \hat{a}_{t-1} + \frac{R_t^{-1} x_{t-1} (p_t - x_{t-1} \hat{a}_{t-1})}{t} \\ &= \hat{a}_{t-1} + \frac{R_t^{-1} x_{t-1} ((\rho \hat{a}_{t-1} + \delta) x_{t-1} + \varepsilon_t - x_{t-1} \hat{a}_{t-1})}{t} \\ R_t &= R_{t-1} + \frac{(x_{t-1} x_{t-1} - R_{t-1})}{t}\end{aligned}$$

Updating beliefs: LS learning

$$\begin{aligned}\hat{a}_t &= \hat{a}_{t-1} + \frac{1}{t} R_t^{-1} x_{t-1} (p_t - x_{t-1} \hat{a}_{t-1}) \\ &= \hat{a}_{t-1} + \frac{1}{t} R_t^{-1} x_{t-1} ((\rho \hat{a}_{t-1} + \delta) x_{t-1} + \varepsilon_t - x_{t-1} \hat{a}_{t-1}) \\ R_t &= R_{t-1} + \frac{1}{t} (x_{t-1} x_{t-1} - R_{t-1})\end{aligned}$$

To get system with only lags on RHS, let $R_t = S_{t-1}$

$$\begin{aligned}\hat{a}_t &= \hat{a}_{t-1} + \frac{1}{t} S_{t-1}^{-1} x_{t-1} ((\rho \hat{a}_{t-1} + \delta) x_{t-1} + \varepsilon_t - x_{t-1} \hat{a}_{t-1}) \\ S_t &= S_{t-1} + \frac{1}{t} (x_t x_t - S_{t-1}) \frac{t}{t+1}\end{aligned}$$

Updating beliefs: LS learning

This can be written as

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \frac{1}{t} Q(\hat{\theta}_{t-1}, x_t, x_{t-1}, \varepsilon_t)$$

or

$$\Delta \hat{\theta}_t = T(\hat{\theta}_{t-1}, x_t, x_{t-1}, \varepsilon_t, t)$$

Key question

- If

$$\Delta \hat{\theta}_t = \frac{1}{t} Q(\hat{\theta}_{t-1}, x_t, x_{t-1}, \varepsilon_t, t)$$

then what can we "expect": about $\hat{\theta}_t$?

- In particular, can we "expect" that

$$\lim_{t \rightarrow \infty} \hat{a}_t = a_{re}$$

Corresponding differential equation

Much can be learned from following differential equation

$$\frac{d\theta}{d\tau} = h(\theta(\tau))$$

where

$$h(\theta) = \lim_{t \rightarrow \infty} E[Q(\theta, x_t, x_{t-1}, \varepsilon_t)]$$

Corresponding differential equation

In our example

$$\begin{aligned}h(\theta) &= \lim_{t \rightarrow \infty} \mathbb{E} [Q(\theta, x_t, x_{t-1}, \varepsilon_t)] \\&= \lim_{t \rightarrow \infty} \mathbb{E} \left[\begin{array}{c} S^{-1} x_{t-1} ((\rho a + \delta) x_{t-1} + \varepsilon_t - x_{t-1} a) \\ (x_t x_t - S) \frac{t}{t+1} \end{array} \right] \\&= \left[\begin{array}{c} MS^{-1} ((\rho - 1) a + \delta) \\ M - S \end{array} \right]\end{aligned}$$

where

$$M = \lim_{t \rightarrow \infty} \mathbb{E} [x_t^2]$$

Analyze the differential equation

$$\frac{d\theta}{d\tau} = h(\theta(\tau)) = \left[\begin{array}{c} MS^{-1}((\rho - 1)a + \delta) \\ M - S \end{array} \right]$$

$$\frac{d\theta}{d\tau} = 0 \text{ if } M = S \ \& \ a = \frac{\delta}{1 - \rho}$$

Thus, the (unique) rest point of $h(\theta)$ is the rational expectations solution

E-stability

$$\hat{\theta}_t - \hat{\theta}_{t-1} = \frac{1}{t} Q \left(\hat{\theta}_{t-1}, x_t, x_{t-1}, \varepsilon_t, t \right)$$

Limiting behavior can be analyzed using

$$\frac{d\theta}{d\tau} = h(\theta(\tau)) = \lim_{t \rightarrow \infty} E [Q(\theta, x_t, x_{t-1}, \varepsilon_t)]$$

A solution θ^* , e.g. $[a_{RE}, M]$, is "E-stable" if $h(\theta)$ is stable at θ^*

E-stability

- $h(\theta)$ is stable if real part of the eigenvalues is negative:
- Here:

$$h(\theta) = \frac{(\rho - 1)a + \delta}{M - S}$$

\implies convergence of differentiable system if $\rho - 1 < 0$

- \implies convergence even if $\rho < -1!$

Implications of E-stability?

- Recursive least-squares: stochastics in $T(\cdot)$ mapping
 - \implies what will happen is less certain, even with E-stability

General implications of E-stability?

- If a solution **is not** E-stable:
 - \implies non-convergence is a probability 1 event
- If a solution **is** E-stable:
 - the presence of stochastics make the theorems non-trivial
 - in general only info about *mean dynamics*

Mean dynamics

See Evans and Honkapohja textbook for formal results.

- Theorems are a bit tricky, but are of the following kind:
If a solution f^ is E -stable, then the time path under learning will either leave the neighborhood in finite time or will converge towards f^* . Moreover, the longer it does not leave this neighborhood, the smaller the probability that it will*
- So there are two parts
 - mean dynamics: convergence towards fixed point
 - escape dynamics: (large) shocks may push you away from fixed point

Importance of Gain

$$\hat{\gamma}_T = \hat{\gamma}_{T-1} + \omega(T)R_T^{-1}x_T (y_T - x_T'\hat{\gamma}_{T-1})$$

- Gain in least squares updating formula, $\omega(T)$, plays a key role in theorems
- $\omega(T) \rightarrow 0$ too fast: you may end up in something that is not an equilibrium
- $\omega(T) \rightarrow 0$ too slowly: you may not converge towards it
- So depending on the application, you may need conditions like

$$\sum_{t=1}^{\infty} \omega(t)^2 < \infty \text{ and } \sum_{t=1}^{\infty} \omega(t) = \infty$$

Special cases

- In simple cases, stronger results can be obtained
- Evans (1989) shows that the system of equations (2) and (3) with standard recursive least squares (gain of $1/t$) converges to rational expectations solution if $\rho < 1$ (so also if $\rho < -1$).

Bayesian learning

- LS learning has some disadvantages:
 - why "least-squares" and not something else?
 - how to choose gain?
 - why don't agents incorporate that beliefs change?
- Beliefs are updated each period
 - ⇒ Bayesian learning is an obvious thing to consider

Bayesian versus LS learning

- LS learning \neq Bayesian learning with uninformed prior at least not always
- Bullard and Suda (2009) provide following nice example

Bayesian versus LS learning

Model:

$$p_t = \rho_L p_{t-1} + \rho_0 \hat{E}_{t-1} [p_t] + \rho_1 \hat{E}_{t-1} [p_{t+1}] + \varepsilon_t \quad (4)$$

- Key difference with earlier model:
 - two extra terms

Bayesian versus LS learning

Solution:

-

$$p_t = bp_{t-1} + \varepsilon_t$$

where b is a solution to

$$b = \rho_L + \rho_0 b + \rho_1 b^2$$

Bayesian learning - setup

- PLM:

$$p_t = \hat{b}_{t-1} p_{t-1} + \varepsilon_t$$

and ε_t has a known distribution

- plug PLM into (4) \implies ALM
 - but a Bayesian learner is a bit more careful

Bayesian learner understands he is learning

$$\begin{aligned}\hat{\mathbf{E}}_{t-1} [p_{t+1}] &= \hat{\mathbf{E}}_{t-1} [\rho_L p_t + \rho_0 \hat{\mathbf{E}}_t [p_{t+1}] + \rho_1 \hat{\mathbf{E}}_t [p_{t+2}]] \\ &= \rho_L p_t + \hat{\mathbf{E}}_{t-1} [\rho_0 \hat{\mathbf{E}}_t [p_{t+1}] + \rho_1 \hat{\mathbf{E}}_t [p_{t+2}]] \\ &= \rho_L p_t + \hat{\mathbf{E}}_{t-1} [\rho_0 \hat{b}_t p_t + \rho_1 \hat{b}_t p_{t+1}]\end{aligned}$$

- and he realizes, for example, that \hat{b}_t and p_t are both affected by ε_t !

Bayesian learner understands he is learning

- Bayesian learner realizes that

$$\hat{\mathbf{E}}_{t-1} [\hat{\mathbf{b}}_t p_{t+1}] \neq \hat{\mathbf{E}}_{t-1} [\hat{\mathbf{b}}_t] \hat{\mathbf{E}}_{t-1} [p_{t+1}]$$

and calculates $\hat{\mathbf{E}}_{t-1} [\hat{\mathbf{b}}_t p_{t+1}]$ explicitly

- LS learner forms expectations thinking that

$$\begin{aligned} \hat{\mathbf{E}}_{t-1} [\hat{\mathbf{b}}_t p_{t+1}] &= \hat{\mathbf{E}}_{t-1} [\hat{\mathbf{b}}_{t-1} p_{t+1}] \\ &= \hat{\mathbf{b}}_{t-1} \hat{\mathbf{E}}_{t-1} \left[\left(\rho_L p + \rho_0 \hat{\mathbf{b}}_{t-1} + \rho_1 \hat{\mathbf{b}}_{t-1} \right) p_t \right] \end{aligned}$$

Bayesian versus LS learning

- Bayesian learner cares about a covariance term
- Bullard and Suda (2009) show that Bayesian is similar to LS learning in terms of E-stability
- Such covariance terms more important in nonlinear frameworks
- Unfortunately not much done with nonlinear models

Learning what?

Model:

$$P_t = \beta E_t [P_{t+1} + D_{t+1}]$$

- Learning can be incorporated in many ways.
- Obvious choices here:
 - ❶ learn about *dgp* D_t and use true mapping for $P_t = P(D_t)$
 - ❷ know *dgp* D_t and learn about $P_t = P(D_t)$
 - ❸ learn about both

Learning what?

- ① Adam, Marcet, Nicolini (2009): one can solve several asset pricing puzzles using a simple model if learning is learning about $E_t [P_{t+1}]$ (instead of learning about *dgp* D_t)
- ② Adam and Marcet (2011): provide micro foundations that this is a sensible choice

Simple model

Model:

$$P_t = \beta \mathbf{E}_t [P_{t+1} + D_{t+1}]$$

$$\frac{D_{t+1}}{D_t} = a \varepsilon_{t+1}$$

with

$$\mathbf{E}_t [\varepsilon_{t+1}] = 0$$

ε_t i.i.d.

Model properties REE

- Solution:

$$P_t = \frac{\beta a}{1 - \beta a} D_t$$

- P_t/D_t is constant
- P_t/P_{t-1} is i.i.d.

Adam, Marcet, & Nicolini 2009

PLM:

$$\hat{E}_t \left[\frac{P_{t+1}}{P_t} \right] = \gamma_t$$

ALM:

$$\frac{P_t}{P_{t-1}} = \frac{1 - \beta\gamma_{t-1}}{1 - \beta\gamma_t} a \varepsilon_t = \left(a + \frac{a\beta\Delta\gamma_t}{1 - \beta\gamma_t} \right) \varepsilon_t$$
$$\gamma_{t+1} = \left(a + \frac{a\beta\Delta\gamma_t}{1 - \beta\gamma_t} \right)$$

Model properties with learning

- Solution is quite nonlinear
 - especially if γ_t is close to β^{-1}
- Serial correlation.
 - in fact there is momentum. For example:

$$\gamma_t = a \ \& \ \Delta\gamma_t > 0 \implies \Delta\gamma_{t+1} > 0$$

$$\gamma_t = a \ \& \ \Delta\gamma_t < 0 \implies \Delta\gamma_{t+1} < 0$$

- P_t/D_t is time varying

Adam, Marcet, & Nicolini 2011

Agent i does following optimization problem

$$\max \widehat{E}_{i,t} [\cdot]$$

- $\widehat{E}_{i,t}$ is based on a sensible probability measure
- $\widehat{E}_{i,t}$ is not necessarily the true conditional expectation

Adam, Marcet, & Nicolini 2011

- Setup leads to standard first-order conditions but with $\widehat{E}_{i,t}$ instead of E_t
- For example

$$P_t = \beta \widehat{E}_{i,t} [P_{t+1} + D_{t+1}]$$

if agent i is not constrained

- Key idea:
 - price determination is difficult
 - agents do not know this mapping
 - \implies they forecast $\widehat{E}_{i,t} [P_{t+1}]$ directly
 - \implies law of iterated expectations cannot be used because next period agent i may be constrained in which case the equality does not hold

Topics - Overview

- ① E-stability and sunspots
- ② Learning and nonlinearities
Parameterized expectations
- ③ Two representations of sunspots

E-stability and sunspots

Model:

$$x_t = \rho E_t [x_{t+1}]$$

x_t cannot explode

no initial condition

Solution:

$$|\rho| < 1 : x_t = 0 \quad \forall t$$

$$|\rho| \geq 1 : x_t = \rho^{-1} x_{t-1} + e_t \quad \forall t$$

where e_t is the sunspot (which has $E_t [e_t] = 0$)

Adaptive learning

PLM:

$$x_t = \hat{a}_t x_{t-1} + e_t$$

ALM:

$$\begin{aligned} x_t &= \hat{a}_t \rho x_{t-1} \\ \implies \hat{a}_{t+1} &= \hat{a}_t \rho \end{aligned}$$

- thus divergence when $|\rho| > 1$ (sunspot solutions)

Adaptive learning

PLM:

$$x_t = \hat{a}_t x_{t-1} + e_t$$

ALM:

$$\begin{aligned} x_t &= \hat{a}_t \rho x_{t-1} \\ \implies \hat{a}_{t+1} &= \hat{a}_t \rho \end{aligned}$$

- thus divergence when $|\rho| > 1$ (sunspot solutions)

Stability puzzle

- There are few counter examples and not too clear why sunspots are not learnable in RBC-type models
 - sunspot solutions are learnable in some New Keynesian models (Evans and McGough 2005)
 - McGough, Meng, and Xue 2011 provide a counterexample and show that an RBC model with *negative* externalities has learnable sunspot solutions

PEA and learning

- Learning is usually done in linear frameworks
- PEA parameterized the conditional expectations in nonlinear frameworks
- \implies PEA is a natural setting to do
 - adaptive learning as well as
 - recursive learning

Model

$$P_t = \mathbb{E} \left[\beta \left(\frac{D_{t+1}}{D_t} \right)^{-\nu} (P_{t+1} + D_{t+1}) \right] = G(X_t)$$

X_t : state variables

Conventional PEA in a nutshell

- Start with a guess for $G(X_t)$, say $g(x_t; \eta_0)$
 - $g(\cdot)$ may have wrong functional form
 - x_t may only be a subset of X_t
 - η_0 are the coefficients of $g(\cdot)$

Conventional PEA in a nutshell

- Iterate to find fixed point for η_i
 - 1 use η_i to generate time path $\{P_t\}_{t=1}^T$

2 let

$$\hat{\eta}_i = \arg \min_{\eta} \sum_t (y_{t+1} - g(x_t; \eta))^2$$

where

$$y_{t+1} = \beta \left(\frac{D_{t+1}}{D_t} \right)^{-v} (P_{t+1} + D_{t+1})$$

3 Dampen if necessary

$$\eta_{i+1} = \omega \hat{\eta}_i + (1 - \omega) \eta_i$$

Interpretation of conventional PEA

- Agents have beliefs
- Agents get to observe long sample generated with these beliefs
- Agents update beliefs
- Corresponds to adaptive expectations
 - no stochastics if T is large enough

Recursive PEA

- Agents form expectations using $g(x_t; \eta_t)$
- Solve for P_t using

$$P_t = g(x_t; \eta_t)$$

- Update beliefs using this one additional observation
- Go to the next period using η_{t+1}

Recursive methods and convergence

Look at recursive formulation of LS:

$$\hat{\gamma}_t = \hat{\gamma}_{t-1} + \frac{1}{t} R_t^{-1} x_t (y_t - x_t' \hat{\gamma}_{t-1})$$

- !!! $\Delta \hat{\gamma}_t$ gets smaller as t gets bigger

General form versus common factor representation

sunspot literature distinguishes between:

- 1 General form representation of a sunspot
- 2 Common factor representation of a sunspot

First consider non-sun-spot indeterminacy

Model:

$$k_{t+1} + a_1 k_t + a_2 k_{t-1} = 0 \text{ or}$$

$$(1 - \lambda_1 L)(1 - \lambda_2 L) k_{t+1} = 0$$

Also:

- k_0 given
- k_t has to remain finite

Multiplicity

Solution:

$$\begin{aligned}k_t &= b_1 \lambda_1^t + b_2 \lambda_2^t \\ k_0 &= b_1 + b_2\end{aligned}$$

Thus many possible choices for b_1 and b_2 if $|\lambda_1| < 1$ and $|\lambda_2| < 2$

Multiplicity

- What if we impose recursivity?

$$k_t = \bar{d}k_{t-1}$$

- Does that get rid of multiplicity? No, but it does reduce the number of solutions from ∞ to 2

$$\begin{aligned} (\bar{d}^2 + a_1\bar{d} + a_2) k_{t-1} &= 0 \quad \forall t \\ &\implies \\ (\bar{d}^2 + a_1\bar{d} + a_2) &= 0 \end{aligned}$$

the 2 solutions correspond to setting either λ_1 or λ_2 equal to 0

Back to sunspots

Doing the same trick with sunspots gives a solution with following two properties:

- ① it has a *serially correlated* sunspot component with the same factor as the endogenous variable (i.e. the common factor)
- ② there are two of these

General form representation

Model:

$$\begin{aligned} E_t [k_{t+1} + a_1 k_t + a_2 k_{t-1}] &= 0 \text{ or} \\ E_t [(1 - \lambda_1 L) (1 - \lambda_2 L) k_{t+1}] &= 0 \end{aligned}$$

General form representation:

$$\begin{aligned} k_t &= b_1 \lambda_1^t + b_2 \lambda_2^t + e_t \\ k_0 &= b_1 + b_2 + e_0 \end{aligned}$$

where e_t is serially uncorrelated

Common factor representation

Model:

$$\begin{aligned} E_t [k_{t+1} + a_1 k_t + a_2 k_{t-1}] &= 0 \text{ or} \\ E_t [(1 - \lambda_1 L) (1 - \lambda_2 L) k_{t+1}] &= 0 \end{aligned}$$

Common factor representation:

$$\begin{aligned} k_t &= b_1 \lambda_i^t + \zeta_t \\ \zeta_t &= \lambda_i \zeta_{t-1} + e_t \\ k_0 &= b_i + \zeta_0 \\ \lambda_i &\in \{\lambda_1, \lambda_2\} \end{aligned}$$

where e_t is serially uncorrelated

References

- Adam, K., G. Evans, and S. Honkapohja, 2006, Are hyperinflation paths learnable, *Journal of Economic Dynamics and Control*.
 - Paper shows that the high-inflation steady state is stable under learning in the seignorage model as discussed in slides.
- Adam, K., A. Marcet, and J.P. Nicolini, 2009, Stock market volatility and learning, manuscript.
 - Paper shows that learning about endogenous variables like prices gives you much more "action" than learning about exogenous processes (i.e. they show that learning with feedback is more interesting than learning without feedback).

References

- Adam, K., and A. Marcet, 2011, Internal rationality, imperfect market knowledge and asset prices, Journal of Economic Theory.
 - Paper motivates that the thing to be learned is the conditional expectation in the Euler equation.
- Bullard, J. and J. Suda, 2009, The stability of macroeconomics systems with Bayesian learners, manuscript.
 - Paper gives a nice example on the difference between Bayesian and Least-Squares learning.

References

- Evans, G.W., and S. Honkapohja, 2001, Learning and expectations in macroeconomics.
 - Textbook on learning dealing with all the technical stuff very carefully.
- Evans, G.W., and S. Honkaphja, 2009, Learning and macroeconomics, Annual Review of Economics.
 - Survey paper.
- Evans, G.W., and B. McGough, 2005, Indeterminacy and the stability puzzle in non-convex economies, Contributions to macroeconomics.
 - Paper argues that sunspot solutions (or indeterminate solutions more generally) are not stable (not learnable) in RBC-type models.
- Evans, G.W., and B. McGough, 2005, Monetary policy, indeterminacy, and learning, Journal of Economic Dynamics and Control.
 - Paper shows that the NK model with some Taylor rules has learnable sunspot solutions
- McGough, B., Q. Meng, and J. Xue, 2011, Indeterminacy and E-stability in real business cycle models with factor-generated externalities, manuscript.
 - Paper provides a nice example of a sunspot in a linearized RBC-type model that