Learning and Monetary Policy

Lecture 4 – Recent Applications of Learning to Monetary Policy

George W. Evans (University of Oregon)

University of Paris X -Nanterre (September 2007)

Lecture 4 Outline

- This lecture describes a selection of recent applications:

- (i) Monetary policy under perpetual learning

- (ii) Learning and inflation persistence

- (iii) Explaining hyperinflations

- (iv) Liquidity Traps

- (v) Dynamic predictor selection and endogenous volatility

(i) Monetary policy under perpetual learning.

Orphanides and Williams (2005a)

– Lucas-type aggregate supply curve for inflation π_t:

$$\pi_{t+1} = \phi\pi^e_{t+1} + (1-\phi)\pi_t + \alpha y_{t+1} + e_{t+1},$$

– Output gap y_{t+1} is set by monetary policy up to white noise control error

$$y_{t+1} = x_t + u_{t+1}.$$

– Policy objective function $\mathcal{L} = (1-\omega)Var(y) + \omega Var(\pi - \pi^*)$ gives rule

$$x_t = -\theta(\pi_t - \pi^*).$$

where under RE $\theta = \theta^P(\omega, \phi, \alpha)$.

Learning: Under RE inflation satisfies

$$\pi_t = \bar{c}_0 + \bar{c}_1 \pi_{t-1} + v_t.$$

Under learning private agents estimate coefficients by **constant gain** (or **discounted**) least squares. Older data dated discounted at rate $(1 - \kappa)$. κ is called the "gain." In RLS just replace $1/t$ by κ.

- Discounting of data natural if agents are concerned to track structural shifts.

- There is some empirical support for constant gain learning.

- With constant gain, LS estimates fluctuate randomly around $(\bar{c}_0, \bar{c}_1)$: there is "**perpetual learning**" and

$$\pi^e_{t+1} = c_{0,t} + c_{1,t}\pi_t.$$

Results:

- Perpetual learning increases inflation persistence.

- Naive application of RE policy leads to inefficient policy. Incorporating learning into policy response can lead to major improvement.

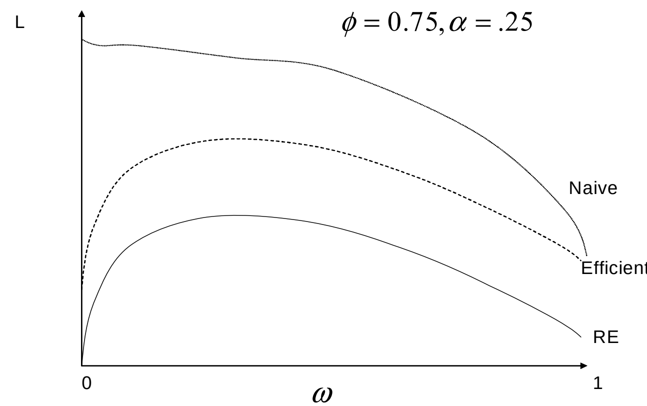

Policymaker's loss

- Efficient policy is more hawkish, i.e. under learning policy should increase θ to reduce persistence. This helps guide expectations.

- Following a sequence of unanticipated inflation shocks, inflation doves (i.e. policy-makers with low θ) can do very poorly, as expectations become detached from RE.

- If agents know π^* and only estimate the AR(1) parameter the policy trade-off is more favorable.

(ii) Learning and inflation persistence

Preceding analysis was an example of empirics based on calibration. Now we consider work that **estimates learning dynamics** (Milani 2005, 2007).

- The source of **inflation persistence** is subject to dispute. Empirically, a backward-looking component is needed in the NK Phillips curve.

- The bulk of the literature assumes RE, but learning could be a reason for inflation persistence.

- Incorporate indexation to Calvo price setting: non-optimized prices indexed to past inflation. This yields

$$\pi_t - \gamma\pi_{t-1} = \delta x_t + \beta E_t(\pi_{t+1} - \gamma\pi_t) + u_t,$$

  where x_t is the output gap. Some earlier work under RE finds γ near 1.

Inflation under learning

- The preceding can be written as

$$\pi_t = \frac{\gamma}{1+\beta\gamma}\pi_{t-1} + \frac{\beta}{1+\beta\gamma}E_t^*\pi_{t+1} + \frac{\delta}{1+\beta\gamma}x_t + u_t.$$

For expectations assume a PLM

$$\pi_t = \phi_{0,t} + \phi_{1,t}\pi_{t-1} + \varepsilon_t$$

Agents use data $\{1, \pi_i\}_0^{t-1}$ to estimate ϕ_0, ϕ_1 using constant gain LS.

- The implied ALM is

$$\pi_t = \frac{\beta\phi_{0,t}(1+\phi_{1,t})}{1+\beta\gamma} + \frac{\gamma+\beta\phi_{1,t}^2}{1+\beta\gamma}\pi_{t-1} + \frac{\delta}{1+\beta\gamma}x_t + u_t.$$

- Alternatively, could use with real marginal cost as the driving variable.

Empirical results

- Data: GDP deflator, output gap is detrended GDP, real marginal cost is proxied by deviation of labor income share from 1960:01 to 2003:04.

- Initialization: agents' initial parameter estimates obtained by using pre-sample data 1951-1959.

- **Methodology:** PLM estimated from constant-gain learning using

$$\kappa = 0.015.$$

This provides estimates of $\phi_{0,t}$, $\phi_{1,t}$ and $E_t^* \pi_{t+1}$. Then estimate ALM using nonlinear LS, which separates learning effects from structural effects. **Note:** Simultaneous estimation of learning rule and the model would be much more ambitious.

- PLM parameters:

(i) $\phi_{1,t}$ initially low in 1950s and 60s, then higher (up to 0.958), then some decline to values above 0.8.
(ii) $\phi_{0,t}$ initially low, then became much higher and then gradual decline after 1980.

- ALM structural estimates: Degree of indexation $\gamma = 0.139$ (with output gap & $\kappa = 0.015$) and declining for higher κ.

- Model fit criterion (Schwartz' BIC) suggest values $\kappa \in (0.015, 0.03)$, with best fit at $\kappa = 0.02$.

- **Conclusion:** Estimates for γ not significantly different from zero. Results very different from those obtained under RE, which finds γ near 1.

- Milani has also estimated full NK models under learning. He finds that also the degree of habit persistence is low in IS curve.

(iii) Explaining Hyperinflations

The **seigniorage model of inflation extended to open economies** and occasional exchange rate stabilizations explain hyperinflation episodes during the 1980s (Marcet and Nicolini 2003).

Basic hyperinflation model (seigniorage model of inflation)

$$M_t^d/P_t = \phi - \phi\gamma(P_{t+1}^e/P_t) \text{ if positive and 0 otherwise,}$$

gives money demand. This is combined with exogenous government purchases $d_t = d > 0$ financed entirely by seigniorage:

$$M_t = M_{t-1} + d_t P_t$$

$$\frac{P_t}{P_{t-1}} = \frac{1 - \gamma(P_t^e/P_{t-1})}{1 - \gamma(P_{t+1}^e/P_t) - d/\phi}.$$

Under RE/perfect foresight, for $d > 0$ not too large, there are **two steady states** $\beta = \frac{P_t}{P_{t-1}}$, $\beta_L < \beta_H$, with a continuum of paths converging to β_H.

Under learning the PLM is

$$\frac{P_{t+1}}{P_t} = \beta,$$

and the implied ALM is

$$\frac{P_t}{P_{t-1}} = \frac{1 - \gamma\beta}{1 - \gamma\beta - d/\phi} \equiv T(\beta).$$

Steady state learning: agents estimate β based on past inflation:

$$\left(\frac{P_{t+1}}{P_t}\right)^e = \beta_t$$

$$\beta_t = \beta_{t-1} + t^{-1}(P_{t-1}/P_{t-2} - \beta_{t-1}).$$

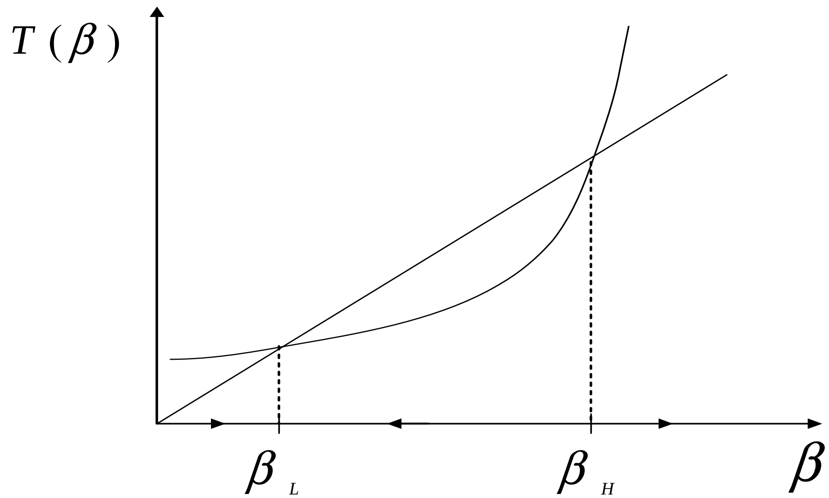

Steady state learning in hyperinflation model

Since $0 < T'(\beta_L) < 1$ and $T'(\beta_H) > 1$, β_L is E-stable, and therefore locally stable under learning, while β_H is not.

Hyperinflation stylized facts

Facts:

- Recurrence of hyperinflation episodes.
- ERR (exchange rate rules) stop hyperinflations, though new hyperinflations eventually occur.
- During a hyperinflation, seigniorage and inflation are not highly correlated.
- Hyperinflations only occur in countries where seigniorage is on average high.

These facts are difficult to reconcile with RE.

Marcet-Nicolini's Approach:

- The low inflation steady state is locally learnable.

- The gain sequence is decreasing for recent MSE large and constant κ for MSE large.

- A sequence of adverse shocks can create explosive inflation. When inflation rises above β^U inflation is stabilized by moving to an ERR.

- The learning dynamics lead to periods of stability alternating with occasional eruptions into hyperinflation.

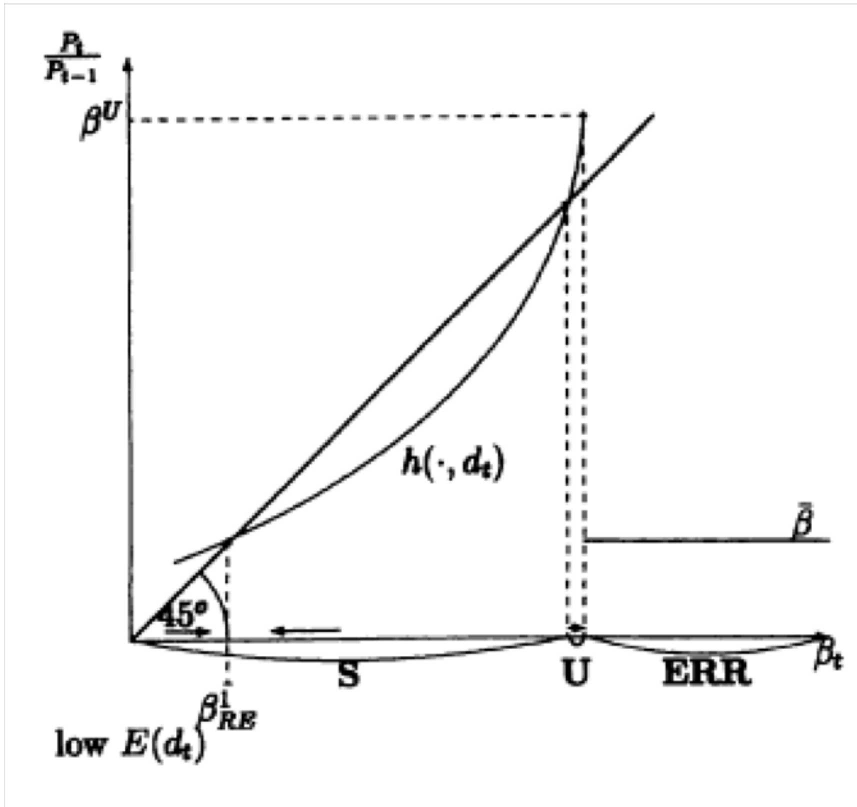
$\frac{P_t}{P_{t-1}}$
β^U
$h(\cdot, d_t)$
$\bar{\beta}$
45^o
S
U
ERR
β_t
β^1_{RE}
low $E(d_t)$

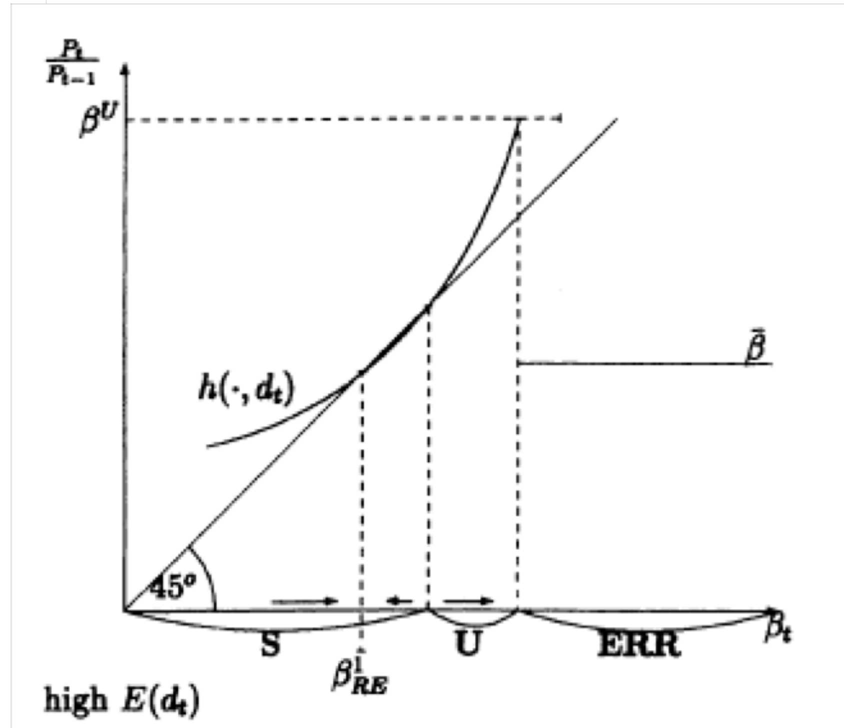
$\frac{P_t}{P_{t-1}}$
β^U
$\bar{\beta}$
$h(\cdot, d_t)$
45^o
S
U
ERR
β_t
β^1_{RE}
high $E(d_t)$

- The learning approach can explain all the stylized facts.

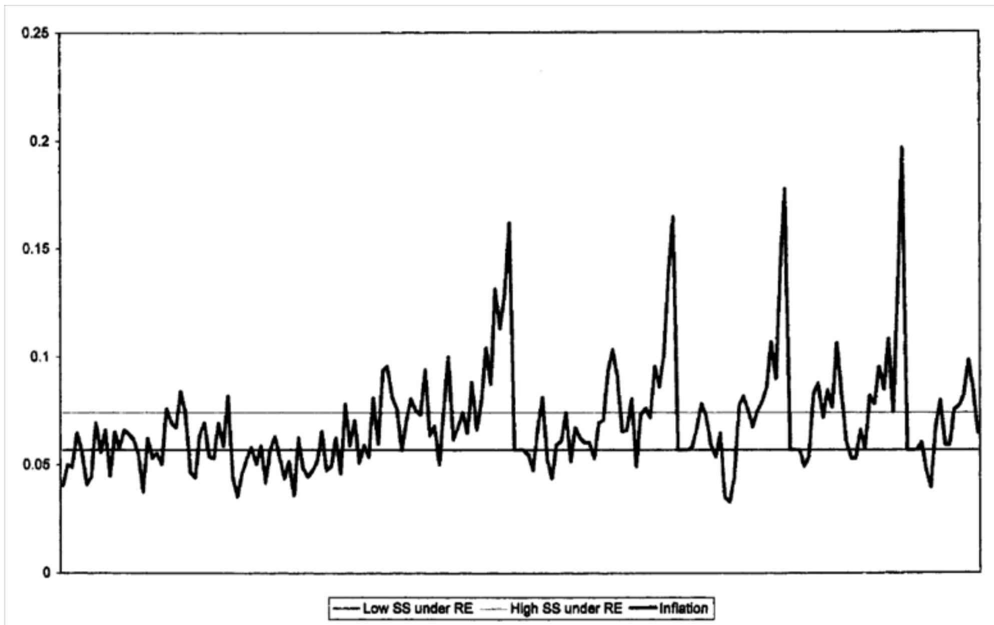

Hyperinflations under learning

(iv) Liquidity Traps

Evans, Guse, Honkapohja (2007), "Liquidity Traps, Learning and Stagnation" consider issues of **liquidity traps and deflationary spirals** under learning.

Possibility of a "liquidity trap" under a global Taylor rule subject to zero lower bound. Benhabib, Schmitt-Grohe and Uribe (2001, 2002) analyze this for RE.

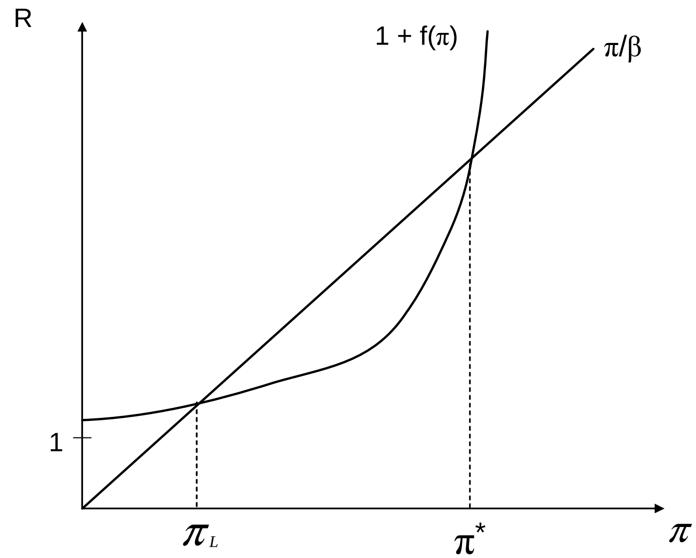

Multiple steady states with global Taylor rule.

What happens under learning?

- Evans and Honkapohja (2005b) analyze a flexible-price perfect competition model:
  - deflationary paths possible
  - switch to aggressive money supply rule at low π avoids liquidity traps.

- Evans, Guse and Honkapohja (2007) consider a model with (i) monopolistic competition (ii) price-adjustment costs. Monetary policy follows a global Taylor-rule. Fiscal policy is standard: exogenous government purchases g_t and Ricardian tax policy that depends on real debt level.

- The key equations are the PC and IS curves

$$
\begin{aligned}
\frac{\alpha\gamma}{\nu}\left(\pi_t - 1\right)\pi_t &= \beta\frac{\alpha\gamma}{\nu}\left(\pi^e_{t+1} - 1\right)\pi^e_{t+1} \\
&\quad + (c_t + g_t)^{(1+\varepsilon)/\alpha} - \alpha\left(1 - \frac{1}{\nu}\right)(c_t + g_t)c_t^{-\sigma_1} \\
c_t &= c^e_{t+1}(\pi^e_{t+1}/\beta R_t)^{\sigma_1},
\end{aligned}
$$

  There are also money and debt equations.

- Two stochastic steady states at π_L and π^*. Under "steady-state" learning, π^* is locally stable but π_L is not.

- Pessimistic expectations c^e, π^e can lead to a deflationary spiral and stagnation.

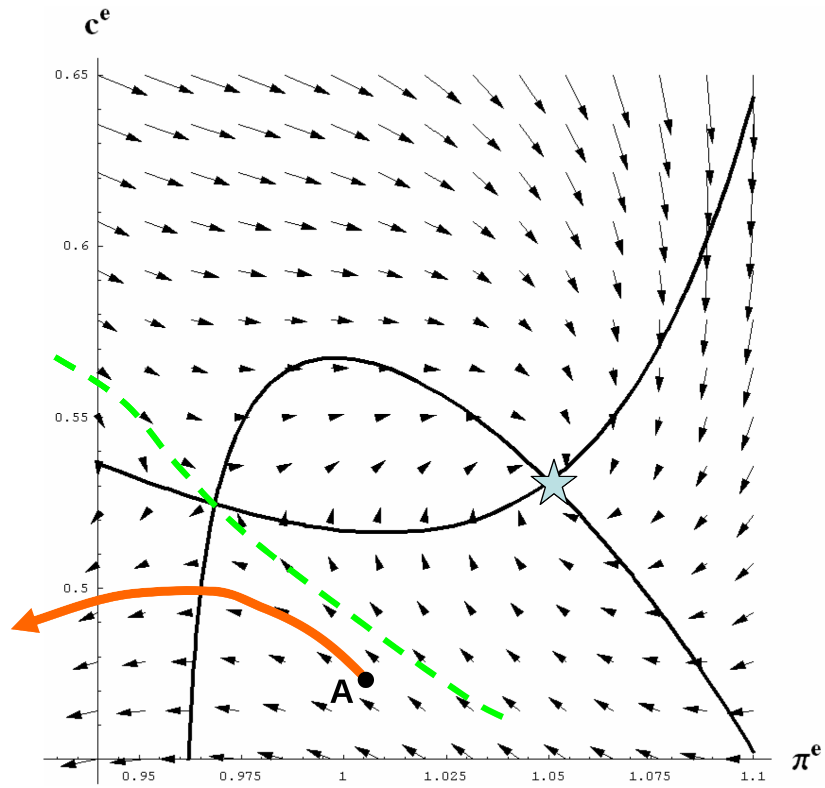

π^e and c^e dynamics under normal policy

- To avoid this we recommend adding aggressive policies at an **inflation threshold** $\tilde{\pi}$, where $\pi_L < \tilde{\pi} < \pi^*$.

- Setting the R-rule so that π_L is a deflationary rate, a natural choice is $\tilde{\pi} = 1$, i.e. zero net inflation.

- If π_t falls to $\tilde{\pi}$ then

  - R_t should be reduced as needed to near the zero lower bound $R = 1$.

  - If necessary, then g_t should also be increased.

Thus *both* aggressive fiscal *and* monetary policy may be needed. Policy needs to focus on inflation, not expansionary spending *per se*.

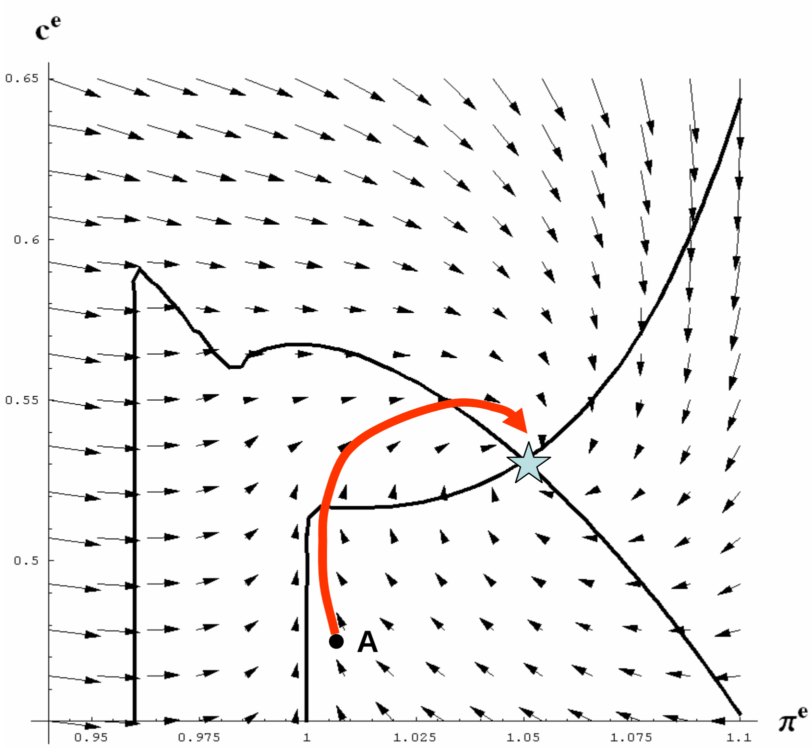

Inflation threshold $\tilde{\pi}$, $\pi_L < \tilde{\pi} < \pi^*$, for aggressive monetary policy and, if needed, aggressive fiscal policy.

(v) Dynamic predictor selection & endogenous volatility

Throughout the lectures we have assumes all agents are using the same econometric model: any **heterogeneity in expectations** has been "mild."

There are several papers that consider heterogeneity in the sense that different groups of **agents use different forecasting models**.

In this topic we start from the approach introduced by Brock and Hommes (1997) in which agents entertain competing forecasting models – naive cheap models and more costly sophisticated models.

The proportions of agents using the different models at t depends on recent forecasting performance. These **proportions evolve over time**.

Branch and Evans (2007) look at agents choosing between **alternative misppecified models** that are each updated using LS learning, and develop an application to macroeconomics that is able to generate **endogenous volatility**.

EMPIRICAL OVERVIEW

In many countries there is substantial evidence of **stochastic volatility** in output and inflation.
– Cogley and Sargent emphasize **parameter drift**, while
– Sims and Zha emphasize **regime switching**.

Our paper provides a theoretical explanation based on learning and dynamic predictor selection.

THE MODEL

We use a simple Lucas-style AS curve with a "quantity theory" AD curve:

$$
\begin{aligned}
AS : q_t &= \phi \left(p_t - p_t^e\right) + \beta_1' z_t \\
AD : q_t &= m_t - p_t + \beta_2' z_t + w_t, \\
z_t &= A z_{t-1} + \varepsilon_t.
\end{aligned}
$$

where w_t, z_t are exogenous and w_t, ε_t are iid. This model can be microfounded along the lines of Woodford (2003). The components of z_t depend on preference, cost and productivity shocks. We assume money supply m_t follows

$$
m_t = p_{t-1} + \delta' z_t + u_t,
$$

where u_t is iid.

Combining equations leads to the reduced form

$$\pi_t = \theta \pi_t^e + \gamma' z_t + \nu_t,$$

where $0 < \theta = (1+\phi)^{-1}\phi < 1$ and v_t depends on w_t, u_t.

The unique REE is

$$\pi_t = (1-\theta)^{-1}\gamma' A z_{t-1} + \gamma' \varepsilon_t + \nu_t.$$

MODEL MISSECIFICATION

– The world is complex. We think econometricians typically misspecify models.

– By the cognitive consistency principle we therefore believe economic agents misspecify their models.

– To model this simply we assume that z_t is 2×1 and agents choose between two models

$$\pi_t^e = b^1 z_{1,t-1} \text{ and } \pi_t^e = b^2 z_{2,t-1}.$$

If the proportion n_1 uses model 1 then

$$\pi_t^e = n_1 b^1 z_{1,t-1} + (1 - n_1) b^2 z_{2,t-1}.$$

– We impose the RPE (restricted perceptions equilibrium) requirement that, given n, each forecast model satisfies

$$E z_{i,t-1}(\pi_t - b^i z_{i,t-1}) = 0, \text{ for } i = 1, 2.$$

– To close the model we follow Brock-Hommes & assume that n depends on the relative MSE of the two models:

$$n_i = \frac{\exp\{\alpha E u_i\}}{\sum_{j=1}^{2} \exp\left\{\alpha E u_j\right\}} \text{ where } Eu = -E\left(\pi_t - \pi_t^e\right)^2 .$$

Here $\alpha > 0$ is the BH "intensity of choice" parameter. We pick α large.

– We show that for α large **there can be two ME (Misspecification Equilibria)** for appropriate z_t processes and other parameters. This can happen even though there is a unique RE.

– In one ME n_1 is near 1 and in the other n_1 is near zero.

REAL-TIME LEARNING WITH CONSTANT GAIN

– Now assume agents **update** their **forecasting using constant gain** learning:

(i) constant gain learning of parameter values b^1 and b^2, and

(ii) constant gain estimates of $Eu_1 - Eu_2$.

– Simulations exhibit both "**regime-switching**" as n_1 moves quickly between values near 1 and 0 and then stay at these values for an extended period, and **parameter drift** as the estimated coefficients b_t^1 and b_t^2 move around.

– **Simulations** strongly exhibit **endogenous volatility** that is absent under RE.

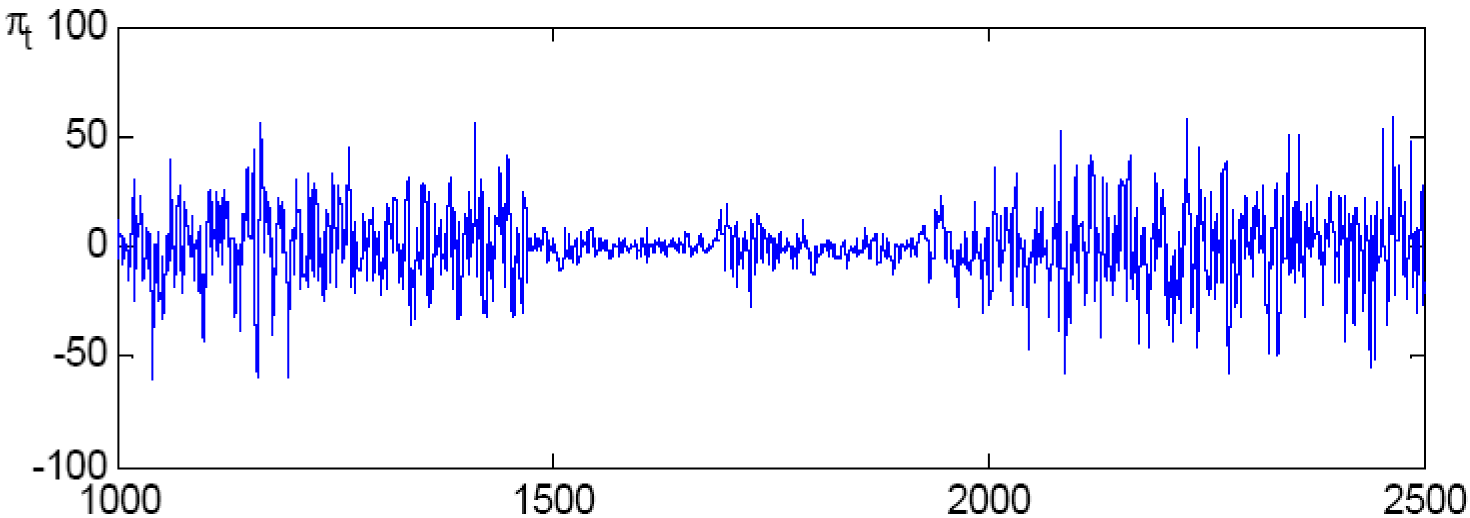

Simulation under constant gain learning and dynamic predictor selection.

Conclusions to Lectures

- Expectations play a large role in modern macroeconomics. People are smart, but boundedly rational. Cognitive consistency principle: economic agents should be about as smart as (good) economists, e.g. model agents as econometricians.

- Stability of RE under private agent learning is not automatic. Monetary policy must be designed to ensure both determinacy and stability under learning.

- Policymakers may need to use policy to guide expectations. Under learning there is the possibility of persistent deviations from RE, hyperinflation, and deflationary spirals with stagnation. Appropriate monetary and fiscal policy design can minimize these risks.

- Learning has the potential to explain various empirical phenomena difficult to explain under RE, e.g. persistence and stochastic volatility.