Adaptive Learning in Macroeconomics

George W. Evans

University of Oregon and University of St. Andrews

June 2019

Second Behavioral Macroeconomics Workshop:

Heterogeneity and Expectations in Macroeconomics and Finance

Introduction

– Macroeconomic models are usually based on **optimizing** agents in dynamic, stochastic setting and can be summarized by a **dynamic system,** e.g., in the simplest case of point expectations and representative agents,

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

$y_t =$ 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, or $E_t^* f\left(y_{t+1}, w_{t+1}\right)$, 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 economists, e.g.

—— model agents like **economic theorists** – an **eductive** approach, or

—— model them as **econometricians** – an **adaptive** approach,

—— one can also blend these, combining strategic and adaptive elements.

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

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

– In **this talk** I follow the **adaptive** approach.

Adaptive Learning: Key Concepts

The benchmark RE (rational expectations) approach assumes agents (i) optimally forecast future variables, and (ii) can solve dynamic optimization problems. The AL (adaptive learning) approach relaxes these, but still aims for a relatively high degree of (bounded) rationality.

Key elements of the AL approach are:

- **Temporary equilibrium.** This idea goes back to Hicks (1946).
  – At t agents make (conditional) decisions based on the current state, realizations of exogenous shocks, and expectations of relevant variables.
  – Aggregation (of possibly heterogeneous agents) and market clearing $\rightarrow$ the TE outcome for the economy's endogenous variables.
  – At $t+1$ expectations are revised and the process repeats.
  – The path of the economy is generated recursively by the sequence of TE.

- **Stability.** The equilibium stochastic processes to which the TE paths converge over time should be stable against small perturbations in initial conditions and how expectations are formed and decisions are made.

- **Boundedly-rational expectations.** The (broad) AL approach emphasizes deviations of agents' forecasts from RE. This can include
  – LS learning, in which forecast rule parameters are updated over time.
  – Choosing between forecast rules, e.g. simple behavioral rules, based on past relative performance.
  – Social learning based on the genetic algorithm approach.

- **Bounded optimality.** For long-lived agents, we consider alternative approaches to how agents attempt to solve dynamic optimization problems.

- **Agent-level**. Our preferred approach is to model aggregate agent-level decision-making and explicitly aggregate these to get the TE. However, AL is often a fruitful approach even in simple, stylized models.

- The **cognitive consistency principle**. Agents should at a minimum be aiming to make good forecasts and attempting to optimize.

- **Restricted Perceptions Equilibria** (RPE). Under AL, REE (RE equilibrium) may be too strong, even as the equilibrium to which the TE path converges. An RPE is a generalization of REE in which agents make optimal forecasts within the class of forecast rules they consider.

We start with a benchmark example: LS learning in the "cobweb" model.

The Cobweb–type Model

Consider a simple univariate reduced form TE equation:

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

$E^*_{t-1} p_t$ 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's 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 simple 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 $0 < \alpha = \lambda/(1+\lambda) < 1$.

Equation (RF)

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

is the TE equation.

One bounded rationality approach is to specify expectations $E^*_{t-1} p_t$ as following *behavioral rules* or *rules of thumb*. In lab experiments Hommes and coauthors have found subjects often use one of several simple rules (naive, trend-chasing, adaptive expectations, etc.). The composition evolves over time.

AL usually focuses on a forecasting equation that nests RE but with *Least Squares learning* of the relevant coefficients.

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.

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:

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

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 multivariate linear models, nonlinear models, and if there are multiple equilibria.

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

It is **not** always the case that REE are stable under learning. When there are multiple REE, E-stability provides a **selection criterion**.

E-Stability in Multivariate Linear Models.

Often macro models can be set up in a standard form

$$y_t = ME_t^* y_{t+1} + Ny_{t-1} + Pv_t.$$

The usual RE solution takes the form $y_t = \bar{a} + \bar{b}y_{t-1} + \bar{c}v_t$, with here $\bar{a} = 0$.

Under LS learning agents use a PLM to make forecasts:

$$\begin{aligned} y_t &= a + by_{t-1} + cv_t \\ E_t^* y_{t+1} &= (I+b)a + b^2 y_{t-1} + (bc + cF)v_t, \end{aligned}$$

based on estimates (a_t, b_t, c_t) which they update using LS.

Inserting the forecasts into the model yields the ALM

$$y_t = M(I+b)a + (Mb^2 + N)y_{t-1} + (Mbc + NcF + P)v_t$$

This gives a **mapping from PLM to ALM:**

$$T(a,b,c) = (M(I+b)a, Mb^2 + N, Mbc + NcF + P).$$

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

$$d/d\tau(a,b,c) = T(a,b,c) - (a,b,c)$$

is locally asymptotically stable at the REE it is said to be **E-stable**. See EH, Chapter 10, for details. The **E-stability conditions** can be stated in terms of the derivative matrices

$$\begin{aligned} DT_a &= M(I+\bar{b}) \\ DT_b &= \bar{b}' \otimes M + I \otimes M\bar{b} \\ DT_c &= F' \otimes M + I \otimes M\bar{b}, \end{aligned}$$

where $\otimes$ denotes the Kronecker product and $\bar{b}$ denotes the REE value of b.

E-stability governs stability under LS learning. **This issue is distinct from the "determinacy" question.**

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 (small) 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 (Cho, Williams and Sargent (2002)).

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?

- Under LS learning convergence would now be to a **Restricted Perceptions Equilibrium** (RPE). There are many types of RPE:
  – Omitted variables, e.g. in the cobweb model if $w_t' = (w_{1t}, w_{2t})$, PLM_1 might include only w_{1t}.
  – Omitted lags, e.g. using a VAR(1) when the REE is VAR(p), $p \geq 2$
  – Functional form, e.g. if the TE map is nonlinear, but agents use a linear forecasting rule.

- In an RPE agents use the best (minimum MSE) econometric model within the class considered.

Variation 3: heterogeneous expectations

In practice, there is **heterogeneity** of expectations across agents.

– This arises if different agents have **different** initial expectations (**priors**), different (possibly random) **gains**, and/or **asynchronous updating**.

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

– Dynamic predictor selection **can be combined with LS learning** of parameters of alternative forecasting models. (Branch and Evans (2006, 2007, . . .), 'misspecification equilibria').

– **Social learning** (genetic algorithm learning as in Arifovic (1994, 1995, . . .) yields heterogeneous expectations.

General Implications of Adaptive Learning

- 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 models with an indeterminate steady state

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

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

Bounded Optimality: Short- vs. long-horizon decision-making

- Most macromodels (RBC, NK, DSGE, etc) assume **infinitely-lived** (or long-lived) **agents** who need to solve dynamic optimization problems.

- **Short-horizon decision-making.** Based on 1-step ahead forecasts agents make decisions that satisfy a necessary condition for optimal decisions.
  – **Shadow-Price Learning** is developed in Evans and McGough (2018). In LQ models SP-learning converges to fully optimal decisions.
  – **Euler-equation Learning** (e.g. Evans and Honkapohja (2006)) can be viewed as a special case of SP-learning.
  – **Value-function Learning** also in Evans and McGough (2018). See Evans D, Evans G and McGough (2019) for an application.
  – SP, EE and value-function learning are boundedly optimal as well as boundedly rational in forecasts. These approaches are tractable.

- **Infinite-horizon decision-making.** Agents solve their dynamic decision problems each period, given their forecasts over the infinite horizon of variables that are exogenous to their decisions.
  – Agents are fully optimizing given their forecasts but use adaptive learning to update their boundedly rational forecast rules.
  – See Preston (2005, 2006), Eusepi and Preston (2011, 2018).
  – **IH-learning** particularly useful if agents foresee a future change in policy.

- **Internal rationality.** Adam, Marcet & Beutel (2017). Agents solve their dynamic decision problem with Bayesian updating based on a prior that may not be externally valid.

- **Finite-horizon decision-making** also possible. See Branch, Evans and McGough (2013).

The New Keynesian Model and Monetary Policy

- Log-linearized New Keynesian model (CGG 1999, Woodford 2003 etc.) under EE learning

  1. "IS" equation (IS curve)

  $$x_t = -\varphi(i_t - E_t^* \pi_{t+1}) + E_t^* x_{t+1} + g_t$$

  2. the "New Phillips" equation (PC curve)

  $$\pi_t = \lambda x_t + \beta E_t^* \pi_{t+1} + u_t,$$

where x_t =output gap, π_t =inflation, i_t = nominal interest rate. $E_t^* x_{t+1}$, $E_t^* \pi_{t+1}$ are expectations. Parameters $\varphi, \lambda > 0$ and $0 < \beta < 1$.

- Observable shocks follow independent stationary AR(1) processes.

- Under Euler-equation learning these are behavioral (temporary equilibrium) equations. IH-learning can also be used.

- Interest rate setting by a standard **Taylor rule,** e.g.

$$
\begin{aligned}
i_t &= \chi_\pi \pi_t + \chi_x x_t \text{ where } \chi_\pi, \chi_x > 0 \text{ or} \\
i_t &= \chi_\pi \pi_{t-1} + \chi_x x_{t-1} \text{ or} \\
i_t &= \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}
\end{aligned}
$$

- Bullard and Mitra (JME, 2002) studied determinacy and E-stability for each rule.

Results for Taylor-rules in NK model (Bullard & Mitra, JME 2002)

- $i_t = \chi_\pi \pi_t + \chi_x x_t$ yields **determinacy and stability** under LS learning **if** $\lambda(\chi_\pi - 1) + (1-\beta)\chi_x > 0$. Note $\chi_\pi > 1$ is sufficient (Taylor principle).

- With $i_t = \chi_\pi \pi_{t-1} + \chi_x x_{t-1}$, **determinacy & E-stability** for $\chi_\pi > 1$ and $\chi_x > 0$ small. There are also **explosive** and **determinate E-unstable** regions.

- For $i_t = \chi_\pi E_t^* \pi_{t+1} + \chi_x E_t^* x_{t+1}$, **determinacy & E-stability** for $\chi_\pi > 1$ and $\chi_x > 0$ small. **Indeterminate & E-stable** for $\chi_\pi > 1$ and χ_x large. Honkapohja and Mitra (JME, 2004) and Evans and McGough (JEDC, 2005) find **stable sunspot solutions** in that region.

The zero lower bound (ZLB), stagnation and deflation

Evans, Guse, Honkapohja (EER, 2008), "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.

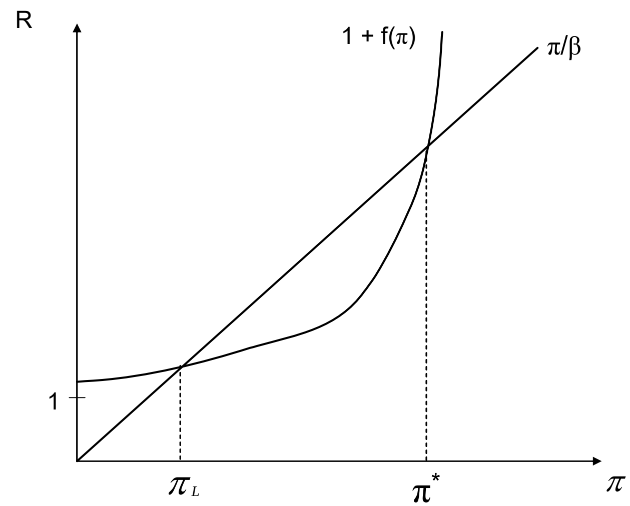

Multiple steady states with global Taylor rule.

– What happens under learning? EGH2008 consider a standard NK model. Monetary policy follows a global Taylor-rule, which implies two steady states.

– The key equations are the (nonlinear) PC and IS curves

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

– 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 deflation and falling output.

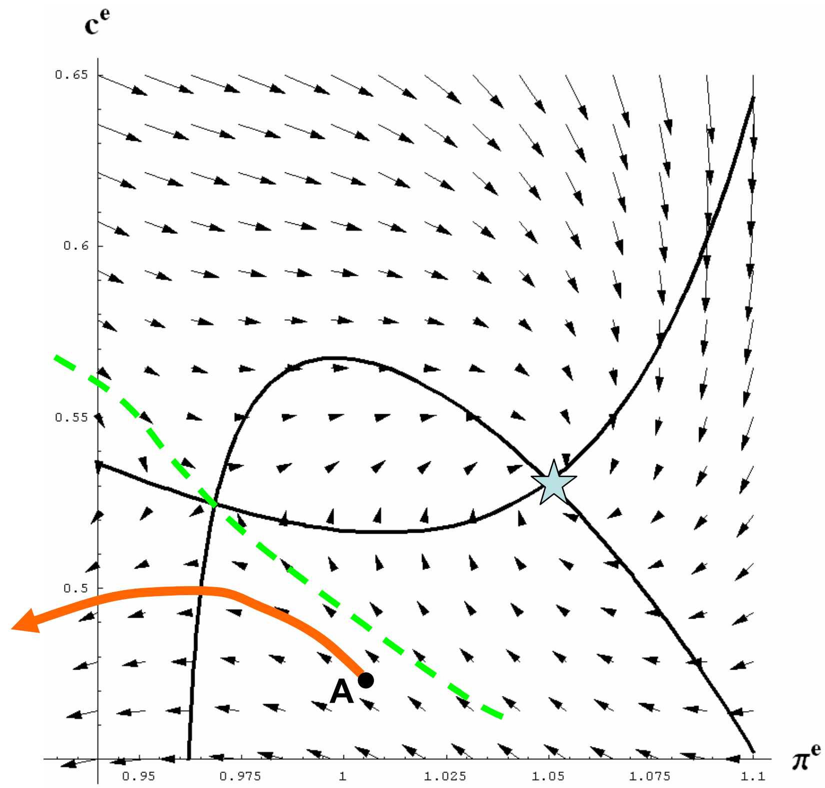

π^e and c^e dynamics under normal policy

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

- Benhabib, Evans and Honkapohja (JEDC, 2014) obtain similar qualitative results for an IH-learning specification and study impacts of alternative fiscal policies in greater detail.
  – Evans, Honkapohja and Mitra (2016) "Expectations, Stagnation and Fiscal Policy" further develop this model by adding **inflation and consumption lower bounds**. This generates an additional locally **stable stagnation steady state**.

Neo-Fisherian Monetary Policy
"Interest Rate Pegs in New Keynesian Models" (Evans and McGough, JMCB, 2018)

- Following the Financial Crisis of 2008-9, the US federal funds rate was essentially at the ZLB for the whole period 2009 – 2015.

- Beginning Dec. 2015 the Fed has started to normalize interest rates. This can be viewed as a return to Taylor rule.

- The **Neo-Fisherian view** (Cochrane, 2011, 2017/8 and Williamson, 2016) is that **normalization** should instead be **to a fixed interest rate peg** at the steady state level consistent with the **2%** inflation target.

- Evans and McGough (JMCB, 2018 and JME, 2019) argue using AL that the neo-Fisherian view is misguided.

- Neo-Fisherianism starts from the Fisher equation

$$R = r\pi$$

where R is the nominal interest rate factor, r is the real interest rate factor and π is the inflation factor. In steady state r is determined by β and the growth rate.

- The **neo-Fisherian argument** is: given r, if the inflation target is π^* then R should be set at $R^* \equiv r\pi^*$. In the basic NK model, and for simplicity ignoring exogenous shocks, the **steady state is an REE** and must satisfy $\pi^e = \pi = R^*/r = \pi^*$.

- The **neo-Fisherian policy conclusion**: if interest rates are low and if inflation and expected inflation are below target, then **announce a fixed interest rate peg** at the higher level $R^* = r\pi^*$. The Fisher equation ensures that π, π^e must increase in line R^*.

- This argument goes against conventional wisdom that low R increases π by increasing demand. **EMcG (2018) argue** the conventional view is right. **Neo-Fisherian policies** can **lead to instability and recession**.

- We show this for both short- and long-horizon AL. Here we use the **NK model with IH-learning** developed in Eusepi and Preston (AEJmacro, 2010) and extended in Evans, Honkapohja and Mitra (2016).

- Agents use linearized decision rules

$$\begin{aligned}
\tilde{c}_t^i &= (1-\beta)\hat{E}_t \sum_{s\geq 0} \beta^s \tilde{y}_{t+s}^i - \frac{\beta^2 \bar{y}}{\pi^*} \hat{E}_t \sum_{s\geq 0} \beta^s \tilde{R}_{t+s} + \frac{\bar{y}}{\pi^*} \hat{E}_t \sum_{s\geq 1} \beta^s \tilde{\pi}_{t+s}, \\
\tilde{\pi}_t^j &= (1-\gamma_1)\hat{E}_t \sum_{s\geq 0} (\beta\gamma_1)^s \tilde{\pi}_{t+s} + \frac{a_2 \pi^*}{\bar{y}} \hat{E}_t \sum_{s\geq 0} (\beta\gamma_1)^s \tilde{y}_{t+s}, \\
\tilde{y}_t &= \tilde{c}_t = \tilde{c}_t^i \text{ and } \tilde{\pi}_t = \tilde{\pi}_t^j.
\end{aligned}$$

- The interest rate follows a forward-looking Taylor rule subject to ZLB.

- We assume this is the full model: fiscal policy is "passive" (and here for simplicity $g = 0$). For cases of "active" fiscal policy, with possible regime switching, see McClung (2018).

Instability of fixed interest rate peg

Suppose **initially in steady state** with π^* target 1% per year (1.0025 quarterly). At $t = 10$ the CB increases the target to 3%. The steady state interest rate increases from 2% to 4% (i.e. from $R^* = 1.005$ to $R^* = 1.01$ quarterly).

Neo-Fisherian policy implements this by announcing new π^* of 3% and increasing R^* to a fixed 4%. Suppose agents immediately adjust π^e from 1% to 2.8%, i.e. almost all the way to 3%. Thereafter π^e is revised in response to observed inflation using AL. Figure 2 of EMcG (2018) gives the result.

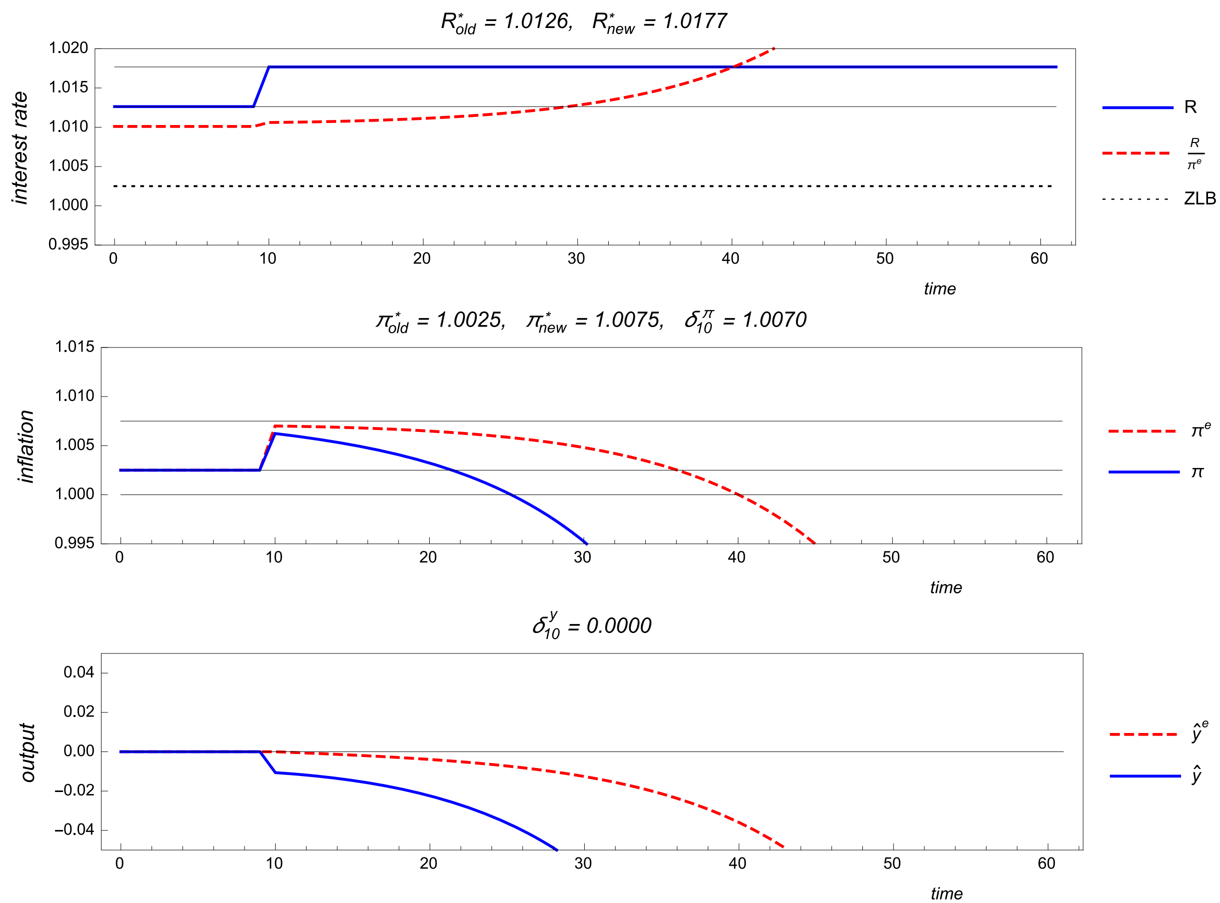

Increase in interest rate peg with almost full adjustment of inflation expectations.

The economy moves into recession, which becomes increasingly severe. π initially falls somewhat short of its target and this feeds back into π^e. This increases the real interest rate, which leads to contracting output. The result is a **cumulative self-fulfilling recession** with falling π, y:

$$\pi < \pi^e \longrightarrow\downarrow \pi^e \longrightarrow\uparrow R/\pi^e \longrightarrow\downarrow y \longrightarrow\downarrow \pi.$$

Because R is held at a **fixed peg** nothing impedes the recession.

Suppose, even more favorably to the neo-Fisherian hypothesis, π^e at $t = 10$ increases the *full way* to the target. Suppose at $t = 11$ there is small one-time negative shock to aggregate demand. This again sets off a cumulative process that leads to falling inflation and recession (Figure 3).

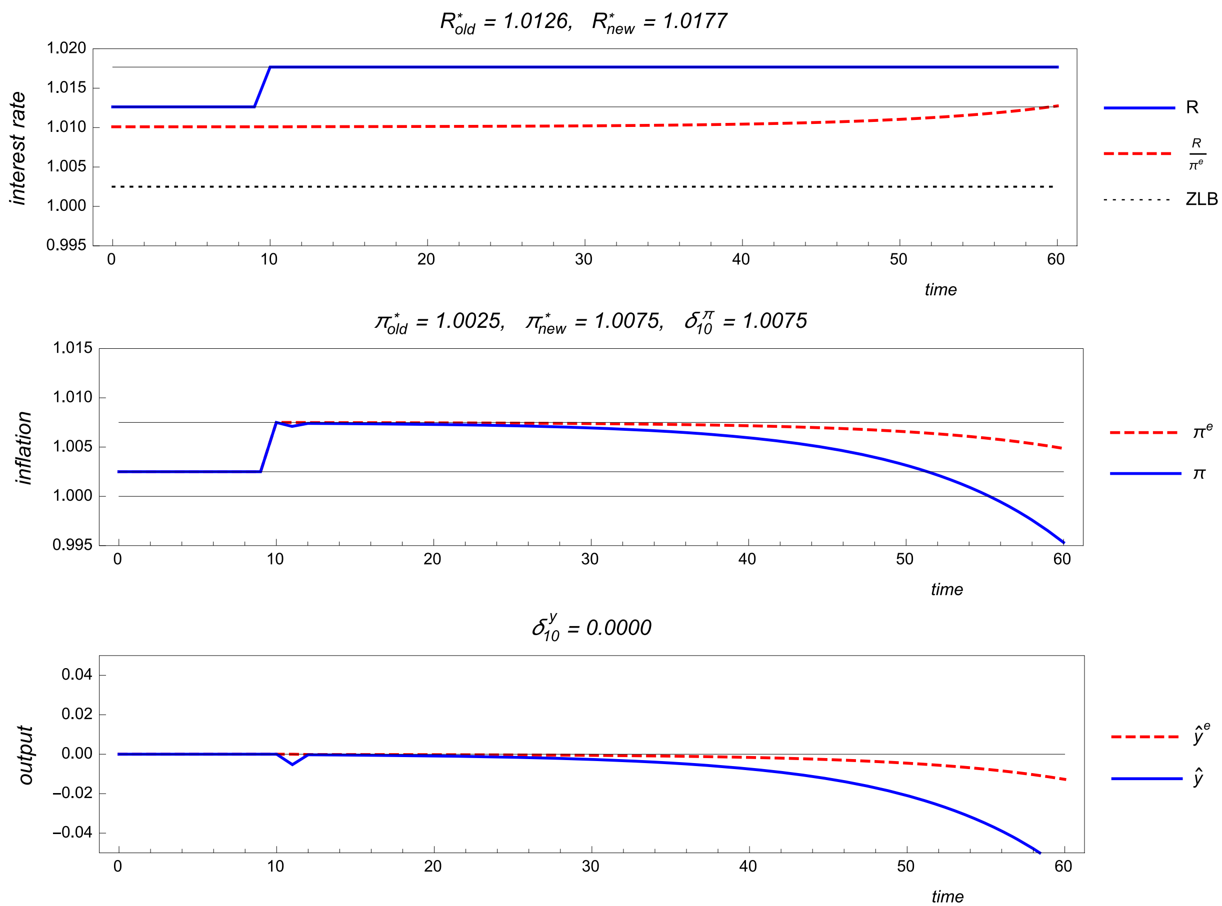

Increase in interest rate peg with full adjustment of inflation expectations.

- The neo-Fisherian policy necessarily generates instability if there is *any sensitivity whatsoever* of expectations to actual data.

- The central mechanism given here is essentially the same as in Howitt (1992), Bullard and Mitra (2002), Evans and Honkapohja (2003), Eusepi and Preston (2010), Benhabib, Evans and Honkapohja (2014) and Evans, Honkapohja and Mitra (2016).

- As in Bullard and Mitra (2002), the Taylor principle is key: to stabilize the economy the interest rate must be adjusted more than one-for-one in response to deviations of inflation or inflation expectations from target.

- One is tempted to say that under AL agents in the economy learn gradually, but there are some economists who never learn.

Near-rational sunspot equilibria in NK model

- Before leaving the NK model, recall that for a forward-looking Taylor rule, if monetary policy responds too aggressively to output there is indeterminacy with stable sunspot equilibria. These results were established for the linearized short-horizon NK model.

- Blanchard (IMF Blog, Dec. 11, 2011): "*. . . the world economy is pregnant with multiple equilibria – self-fulfilling outcomes of pessimism or optimism, with major macroeconomic implications.*" This view makes imperative understanding when and how sunspot equilibria, which represent and characterize the class of stationary multiple equilibria, are consistent with the modern DSGE paradigm.

- Evans and McGough (2019) have shown how to examine this issue for agents using **linear forecasting rules, but nonlinear short-horizon decision rules**. We think this a natural way to model agent's decision-making. We call these equilibria NRSE: **near-rational sunspot equilibria**.

- The nonlinear NK model is:

$$
\begin{aligned}
y_t^{-\sigma} &= \beta \cdot R_t E_t^* \pi_{t+1}^{-1} y_{t+1}^{-\sigma} \\
\gamma \cdot \pi_t(\pi_t - \pi^*) &= \beta \cdot \gamma \cdot E_t^* \pi_{t+1}(\pi_{t+1} - \pi^*) + \left(\frac{\nu}{\alpha}\right) y_t^{\frac{1+\chi}{\alpha}} + (1-\nu) y_t^{1-\sigma} \\
R_t &= R^* \left(E_t^* \left(\frac{\pi_{t+1}}{\pi^*} \right) \right)^{\alpha_\pi \cdot \pi^*} \left(E_t^* \left(\frac{y_{t+1}}{y^*} \right) \right)^{\alpha_y \cdot y^*} e^{v_t}, \\
v_t &= \rho v_{t-1} + u_t,
\end{aligned}
$$

If policymakers use relatively large values of α_π, α_y then the steady-state is indeterminate.

Agents forecast using

$$x_t = a + b\eta_t + cv_t, \text{ where } \eta_t = (\lambda + \xi)\eta_{t-1} + \varepsilon_t,$$

for $x' = (y, \pi)$ and sunspot η_t. We show convergence under AL to an NRSE $\longrightarrow$ extra y, π volatility. This can be avoided by reducing α_π, α_y.

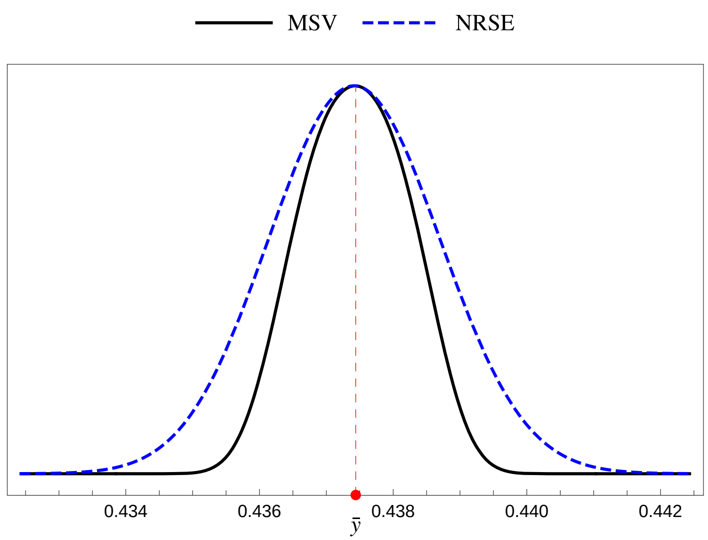

Output densities of near-rational MSV and NRSE

Learning about Risk and Return: a simple model of bubbles and crashes

(Branch and Evans, AEJ macro 2011)

Thus, this vast increase in the market value of asset claims is in part the indirect result of investors accepting lower compensation for risk. Such an increase in market value is too often viewed by market participants as structural and permanent . . . Any onset of increased investor caution elevates risk premiums and, as a consequence, lowers asset values and promotes the liquidation of the debt that supported higher asset prices. This is the reason that history has not dealt kindly with the aftermath of protracted periods of low risk premiums.

Alan Greenspan (2005).

We use a simple mean-variance linear asset pricing model, similar to DeLong, Shleifer, Summers and Waldman (1990) and add boundedly rational AL. Related model of asset-pricing, using the internal rationality approach, are Adam, Marcet, Nicolini (2016) and Adam, Marcet and Beutel (2017).

There is a risky asset with dividend y_t and price p_t and a risk-free asset that pays the rate of return $R = \beta^{-1}$, where $0 < \beta < 1$. Demand for the risky asset is

$$z_{dt} = \frac{E_t^* \left(p_{t+1} + y_{t+1}\right) - \beta^{-1} p_t}{a\sigma_t^2},$$

where E_t^* are (possibly) non-rational expectations and

$$\sigma_t^2 = Var_t^*(p_{t+1} + y_{t+1} - Rp_t).$$

Writing z_{st} for risky asset supply and setting $z_{dt} = z_{st}$ we have

$$p_t = \beta E_t^* \left(p_{t+1} + y_{t+1}\right) - \beta a \sigma_t^2 z_{st}.$$

$a > 0$ measures risk-aversion.

This is a very simple model that incorporates risk. We keep it simple because we are going to add learning.

We also assume: (i) Dividends y_t are a constant plus white noise, and (ii) asset supply $z_{st} = z_0 + v_t$, white noise, unless price falls below a small proportion of its fundamental value. This implies that the price dynamics are entirely driven by learning.

Rational Expectations Equilibria

Under RE, with exogenous supply, there are two solution classes.

– Fundamentals solution:

$$p_t = \frac{\beta(y_0 - a\sigma^2 s_0)}{1-\beta} - \beta a\sigma^2 v_t$$

Here σ^2 is an equilibrium object.

– Rational bubbles solutions

$$p_t = a\sigma^2 s_0 - y_0 + \beta^{-1} p_{t-1} + a\sigma^2 v_{t-1} + \xi_t,$$

where ξ_t is an arbitrary MDS, i.e. $E_t \xi_{t+1} = 0$.

Since $0 < \beta < 1$ the bubbles solutions are explosive in conditional mean.

Stability under Learning

We give agents a PLM (perceived law of motion) that nests the fundamentals solution and also allows for the bubble term in p_{t-1},

$$\begin{aligned} p_t &= k + cp_{t-1} + \varepsilon_t, \\ \sigma^2 &= Var_t(p_{t+1} + y_{t+1}) \end{aligned}$$

where ε_t is perceived white noise with constant variance.

Under learning agents estimate k, c and σ^2 using an adaptive learning algorithm: (recursive) LS learning for k, c and a recursive estimate of σ^2.

Proposition: (1) The fundamentals REE is locally stable under learning. (2) The bubbles REE are unstable under learning.

However, the **transitional learning dynamics** exhibits paths in which the agents' PLM **escapes to a random walk**, $k = 0, c = 1$, with asset prices sensitive to **changed estimates of risk**, leading to bubbles and crashes.

The random walk PLM behaves like a near-rational bubble.

Discounted LS. Furthermore, under **discounted (or "constant gain") learning** (in which agents discount past data at a geometric rate) there can be **recurring bubbles and crashes**.

If the gains (discounting) are small, the dynamics stay near the fundamentals RE. **For larger gains (discounting) there are more frequent escapes**.

Stochastic simulations.

Frequent bubbles and crashes arise when the gain on the estimate of risk (γ_2) is relatively large. We vary $\gamma_2 = 0.001$ to 0.04.

Starting from the fundamentals RE, **crashes and bubbles can arise from various sequences of random shocks, e.g.**

$u_t \approx 0$,$v_t \approx 0 \rightarrow\downarrow \sigma_t^2 \rightarrow \uparrow p_t \rightarrow$ random-walk beliefs.

Random-walk beliefs are almost self-fulfilling and have price high volatility.

Explosive price bubbles $\rightarrow\uparrow \sigma_t^2 \rightarrow$ crashes.

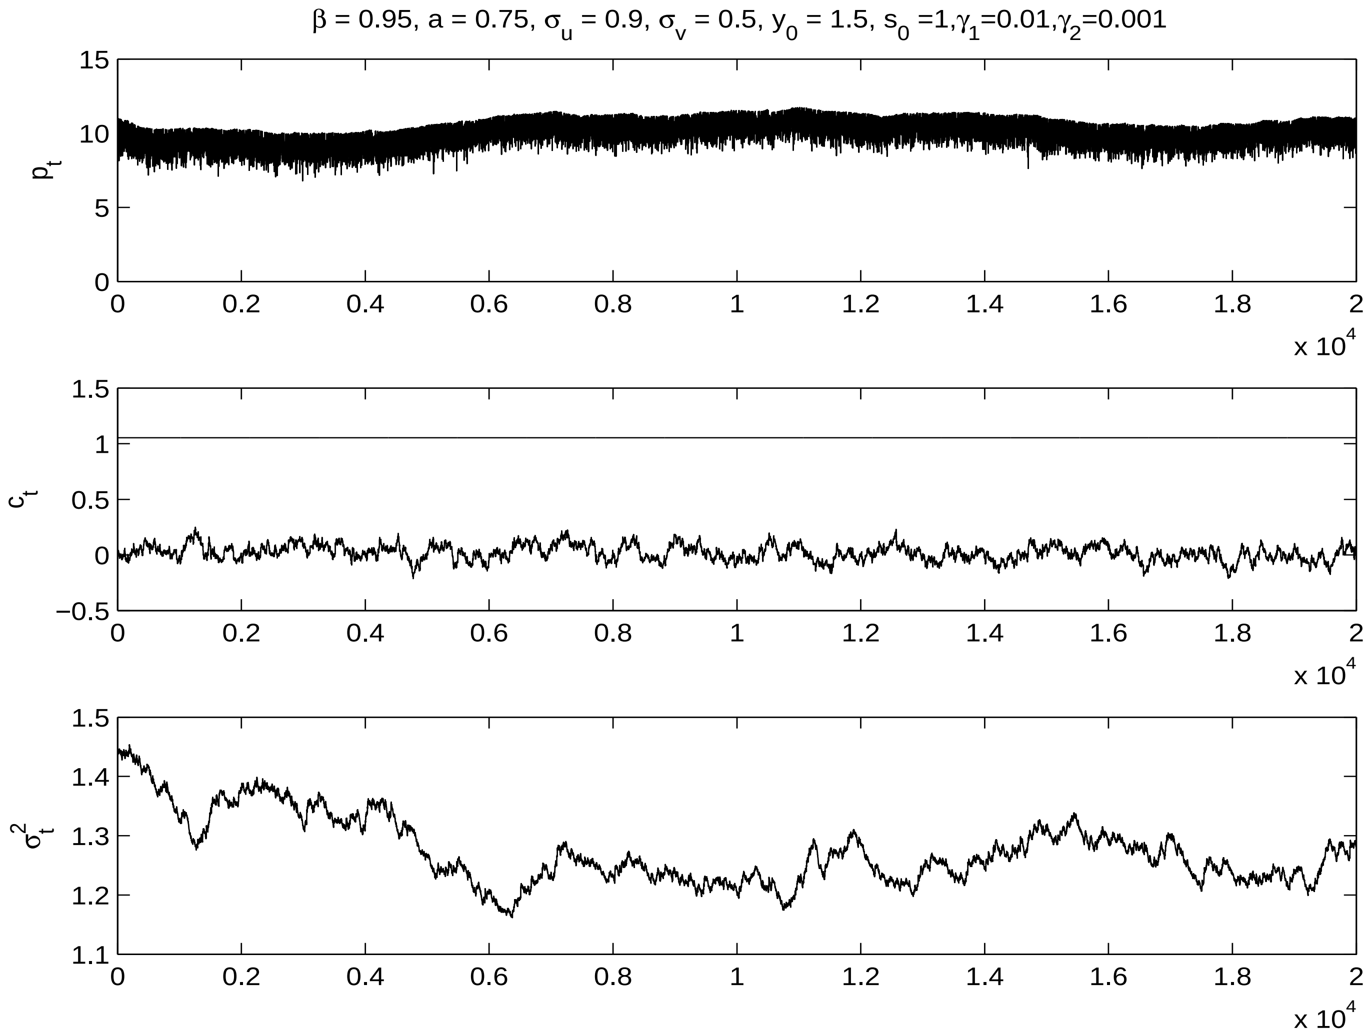

β = 0.95, a = 0.75, σu = 0.9, σv = 0.5, y0 = 1.5, s0 =1,γ1=0.01,γ2=0.001
pt
ct
σt2
x 10^4

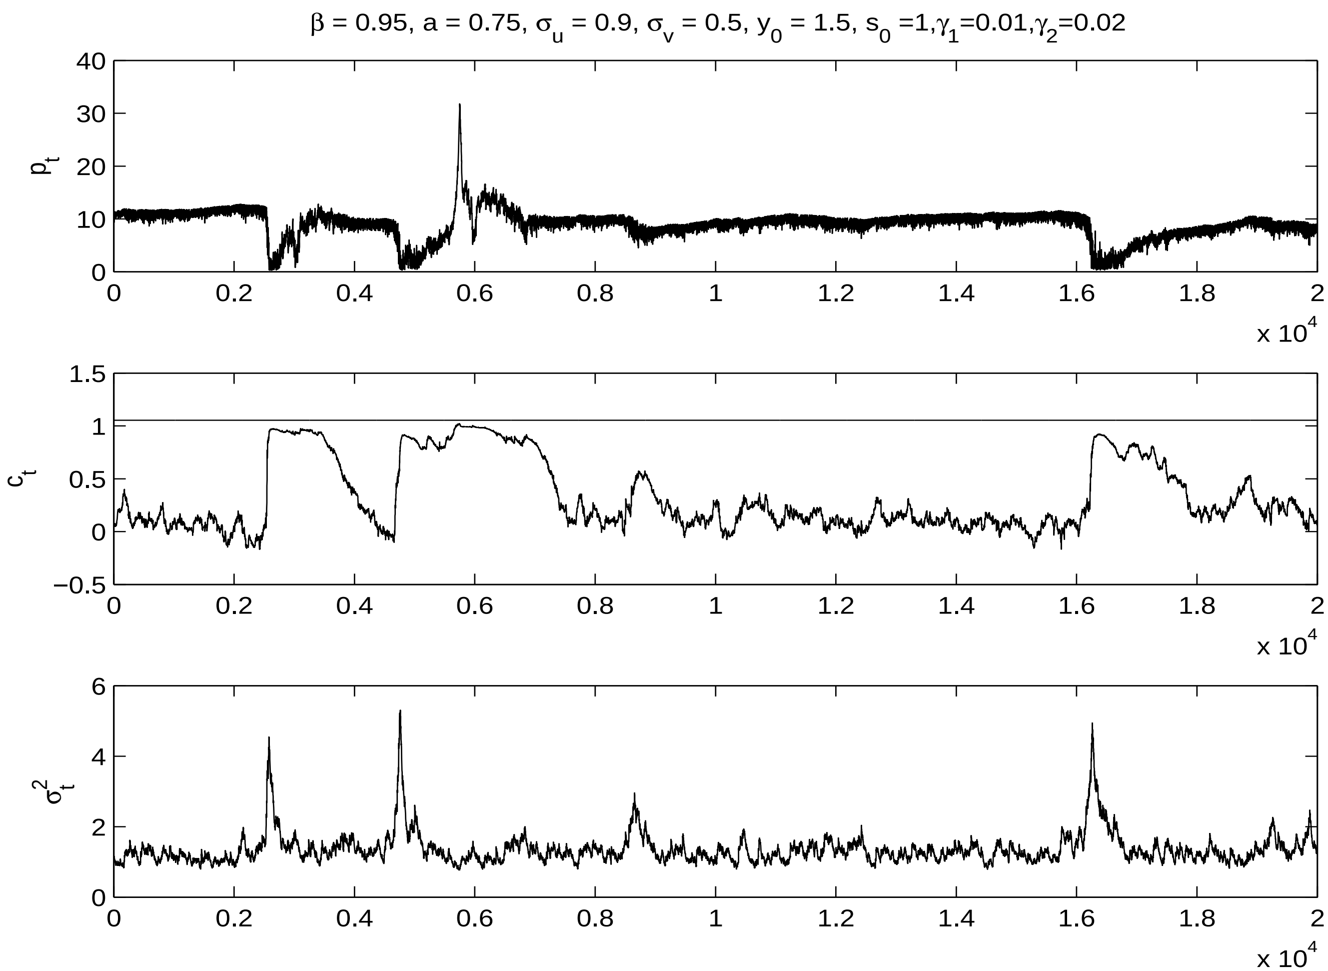
β = 0.95, a = 0.75, σ_u = 0.9, σ_v = 0.5, y_0 = 1.5, s_0 =1,γ_1=0.01,γ_2=0.02
p_t
c_t
σ^2_t
x 10^4

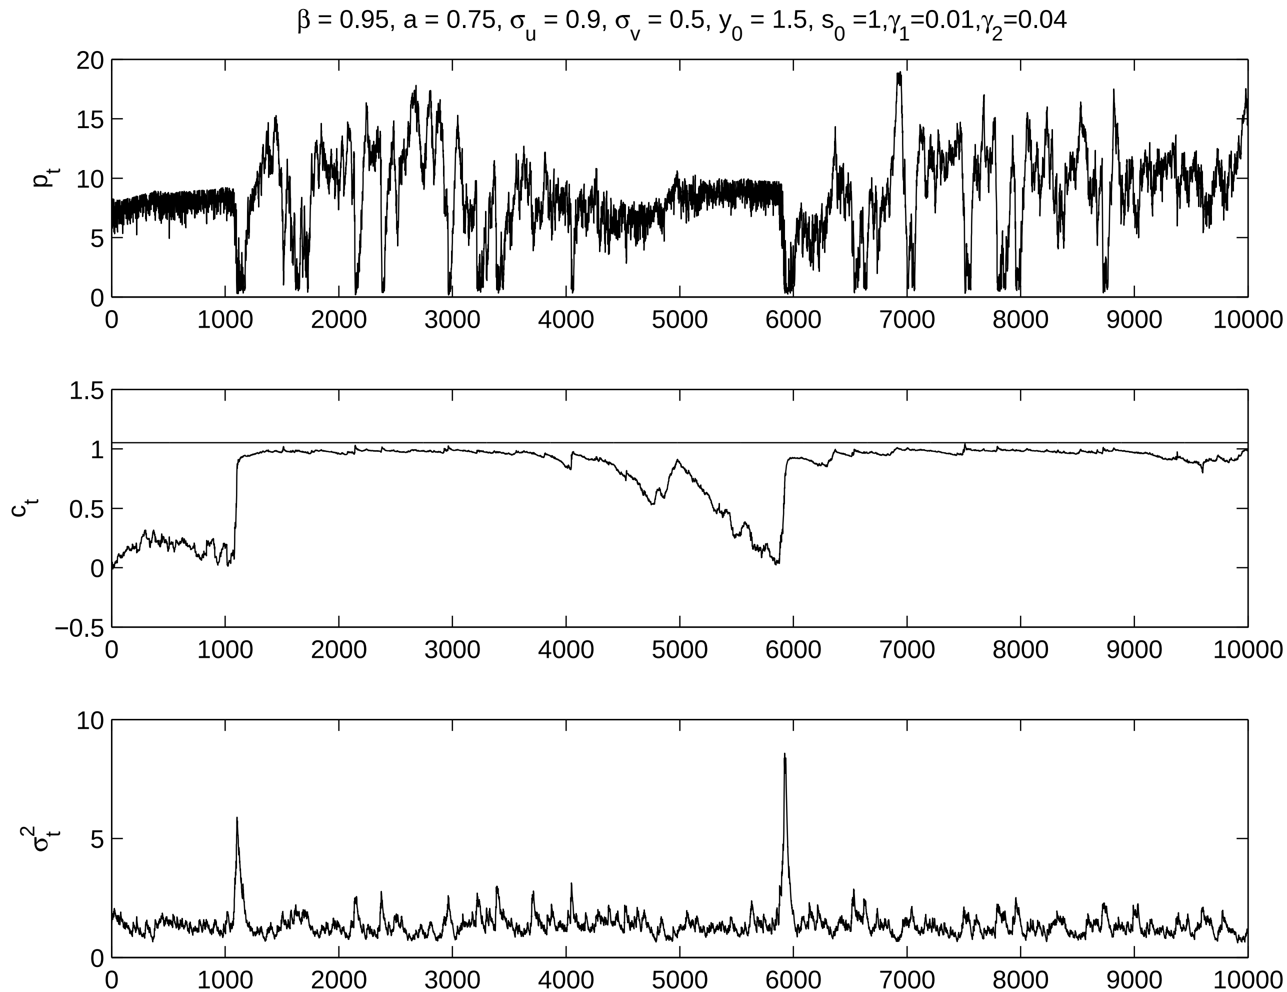
β = 0.95, a = 0.75, σu = 0.9, σv = 0.5, y0 = 1.5, s0 =1,γ1=0.01,γ2=0.04
pt
ct
σ2t
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

Macro Experiments 1: sunspots in the lab

Are Sunspots Learnable? An Experimental Investigation in a Simple Macroeconomic Model

Jasmina Arifovic, George Evans and Olena Kostyshyna

Can agents coordinate on an SSE (stationary sunspot equilibrium)? We investigate this in a simple, stylized macro model. Agents/subjects are firms with production function

$$y_t = \psi_t \sqrt{n_t},$$

where ψ_t indexes productivity. Profit is output minus labor costs and w is fixed

$$\Pi_t = \psi_t \sqrt{n_t} - w n_t.$$

Productivity ψ_t depends on average n across all other firms $\overline{N_t}$. The economy is a sequence of static market equilibria.

Each subject/firm decides on n_t, before knowing productivity, ψ_t because it does not know $\overline{N_t}$, when its decision is made. There is a positive production externality:

$$\begin{aligned} \psi_t &= 2.5 && \text{when} \quad \overline{N_t} \leq 11.5 \\ \psi_t &= 2.5 + (\overline{N_t} - 11.5) && \text{when} \quad 11.5 < \overline{N_t} < 13 \\ \psi_t &= 4 && \text{when} \quad 13 \leq \overline{N_t} \end{aligned}$$

There are three perfect foresight steady states: $n_L = 6.25$, $n_M = 12.54$ and $n_H = 16$.

n_L and n_H are stable under learning, while n_M is not.

Each period there is a public announcement, either “High employment is forecasted this period” or “Low employment is forecasted this period”. The announcement is known to be exogenous and random with $p_{HH} = 0.8$ and $p_{LL} = 0.7$.

In addition to the steady states there are SSEs in which agents choice depends on the announcement. SSEs switching between n_L and n_H are stable under learning.

Each period subjects are asked to forecast $\overline{N_t}$, and based on this their choice of optimal n_t is made, which determines actual aggregate N_t. Payoffs depend on either their profits or on MSE..

Will agents coordinate on SSEs or only on steady states?

See Figures. The experiments find that often they appear to converge on the stable SSE, but sometimes they converge on n_L or n_H.

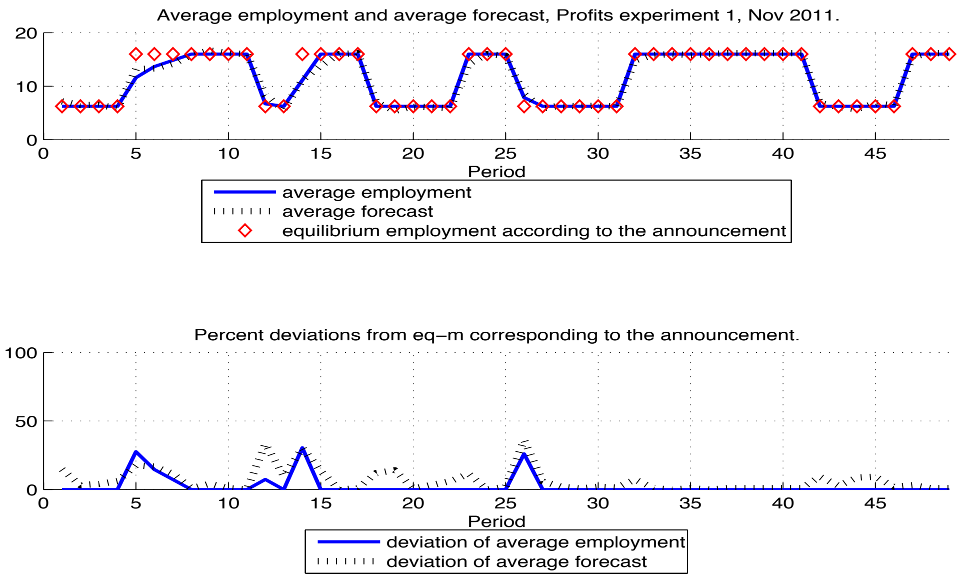

Figure 4: Session 1 of profits treatment.

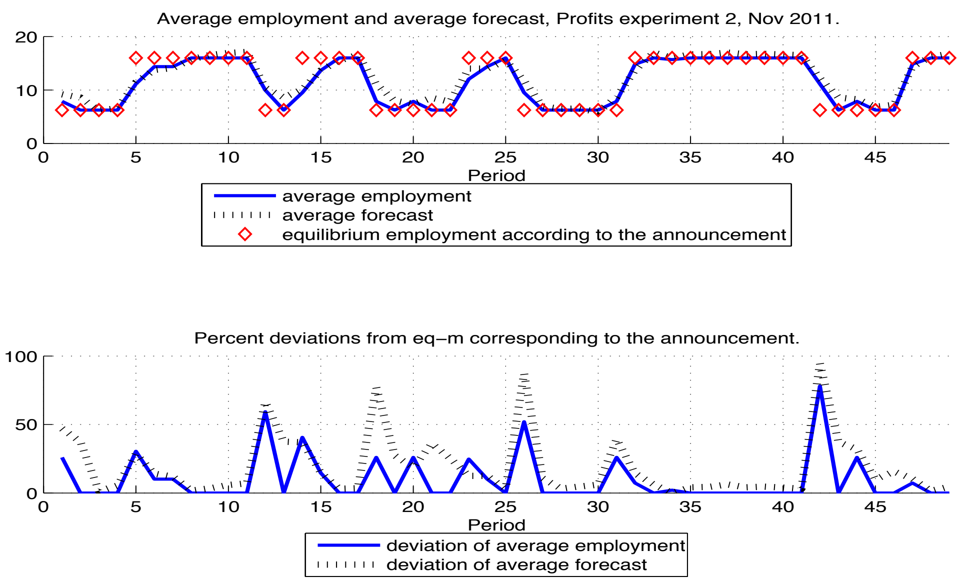

Figure 5: Session 2 of Profits treatment.

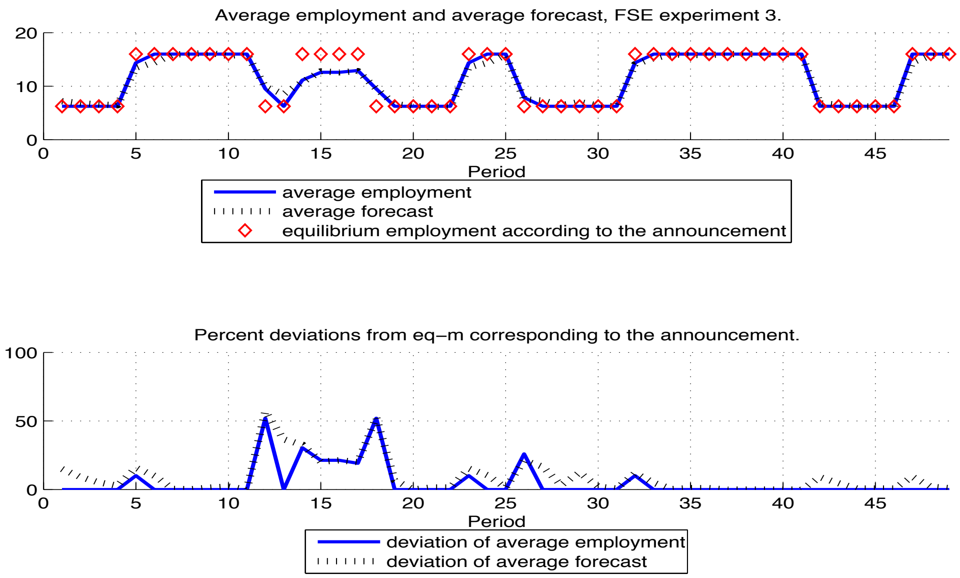

Figure 12: Session 3 of FSE treatment.

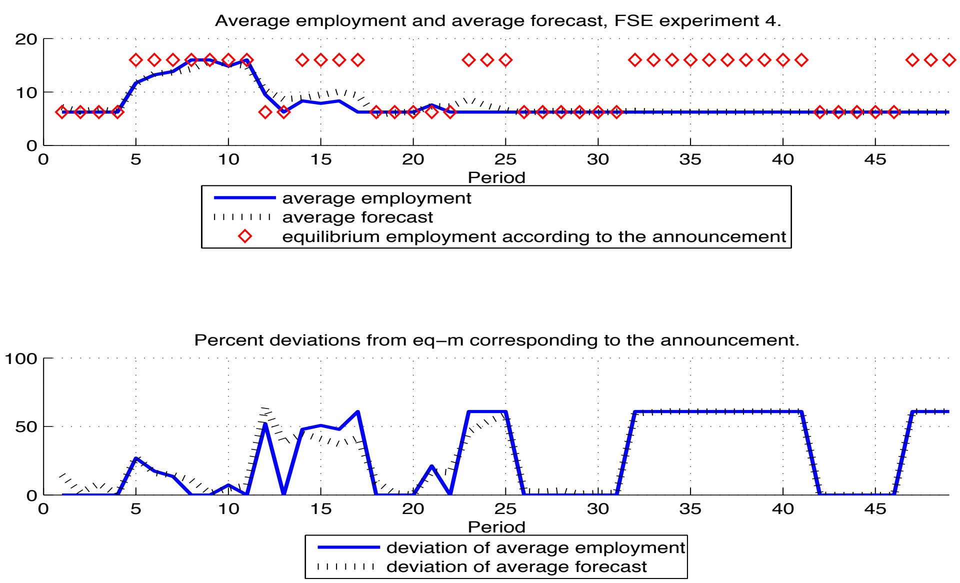

Figure 13: Session 4 of FSE treatment.

Macro Experiments 2: asset-price dynamics in the lab

Are long-horizon expectations (de-)stabilizing? A lab experiment
George Evans, Cars Hommes, Bruce McGough, Isabelle Salle

This experiment uses a Lucas "tree" type set-up. There is a single asset: chickens. The dividend of the asset is the single consumption good: eggs. Each period chickens can be traded for eggs at a market-clearing price. The economy has a constant small probability of ending. Payoffs are based on either utility of consumption or on MSFE (picked at random).

There are two types of subjects: short-horizon ($T = 1$) and long-horizon ($T = 10$). We give each type a boundedly rational trading decision, which depends both on the market price and the expected average price over their horizon.

There are 4 treatments: L, S, M70 (70% S) and M30 (30% S).

The expectational feedback parameter (from p^e to p) is positive and less than one, but higher for S than for L. AL theory predicts markets will be more likely to converge to fundamentals and be less volatile when it is populated with long-horizon agents.

See Figures. Asset prices are less volatile and tend to converge to the fundamental price for treatments with L, while with S only there are frequent divergences from fundamentals and there is more price volatility..

B Individual Plots

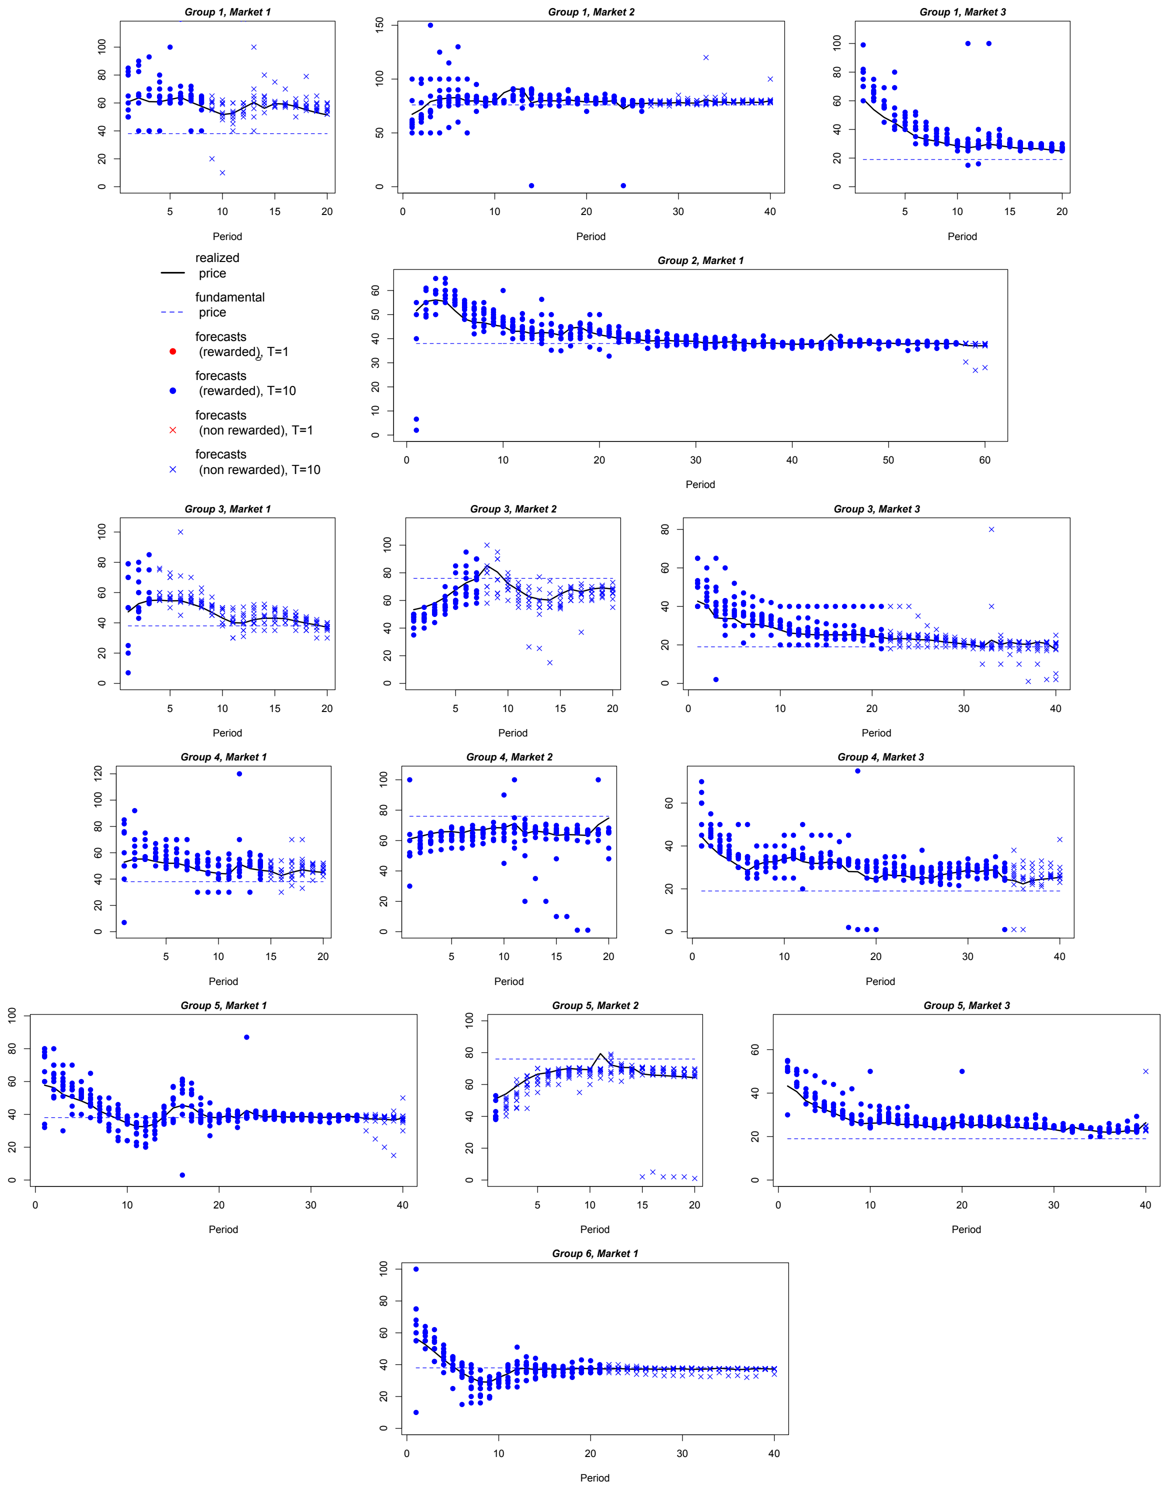

Figure 10: Treatment L: 100% $T = 10$ (6 groups)

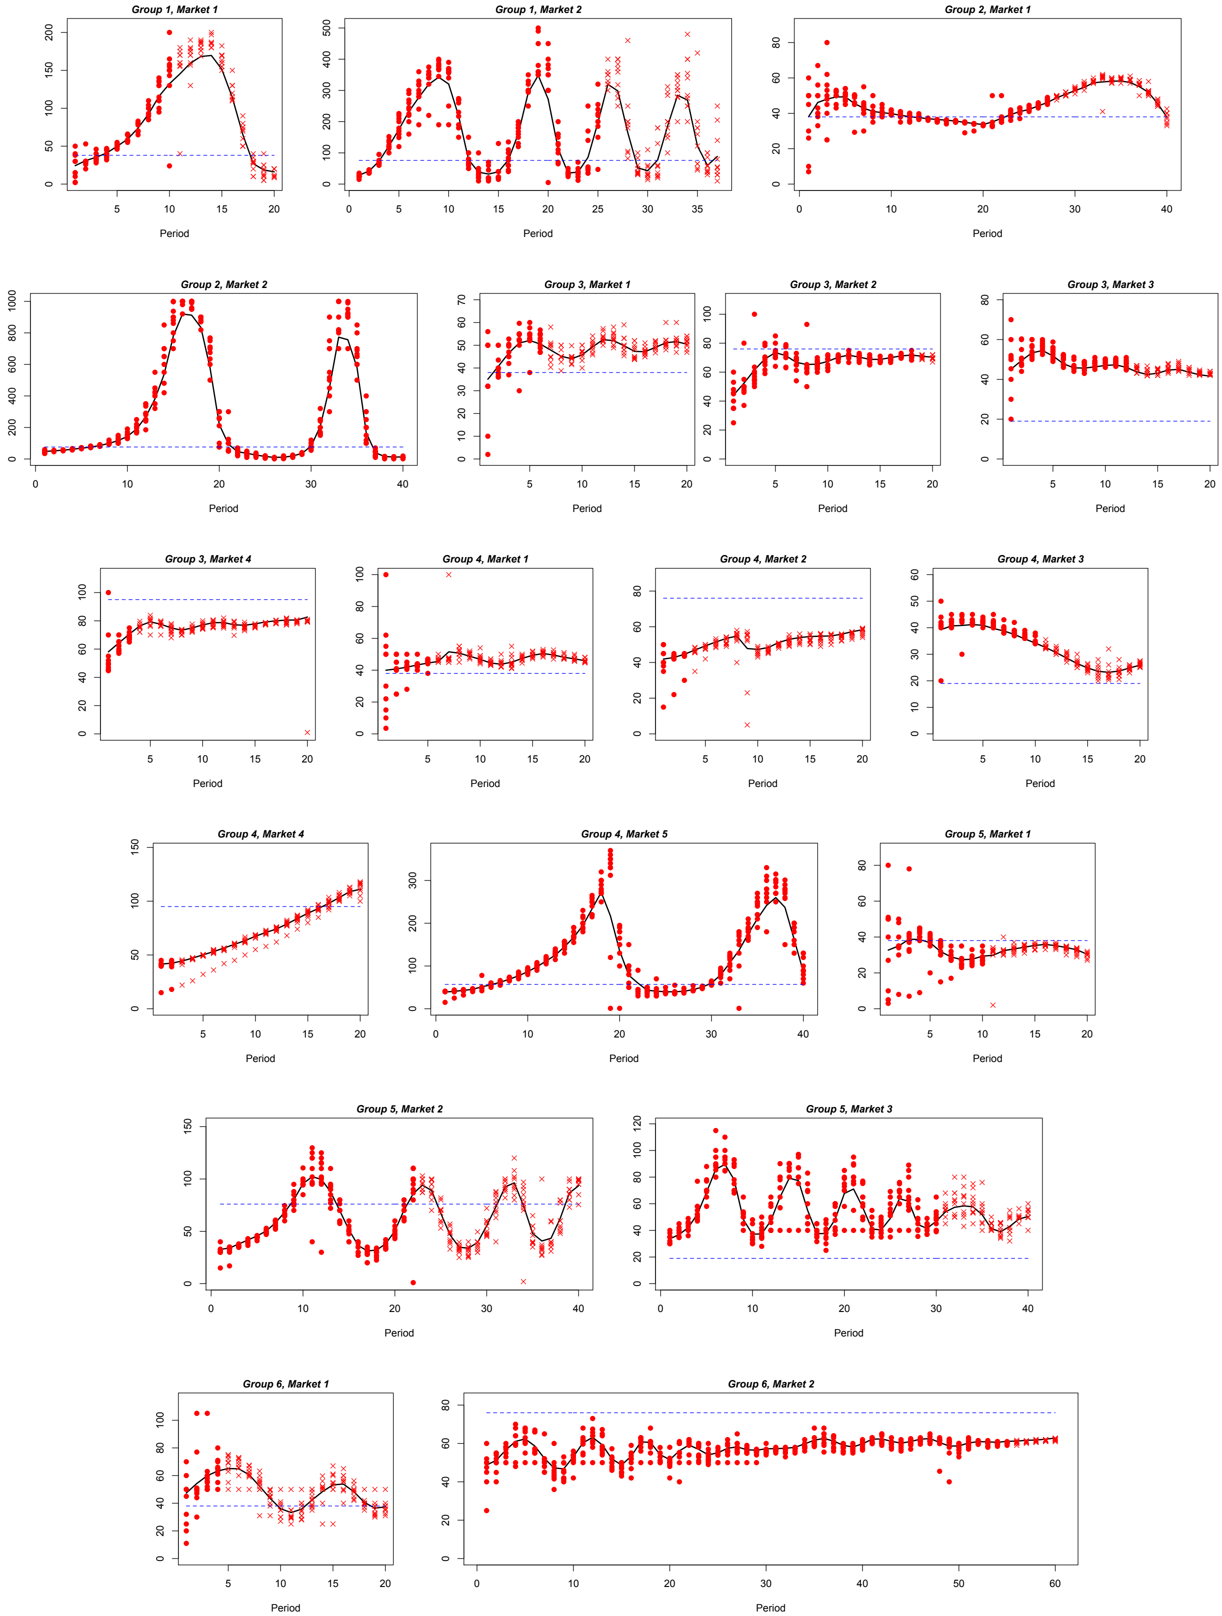

Figure 11: Treatment S: 100% $T = 1$ (6 groups)

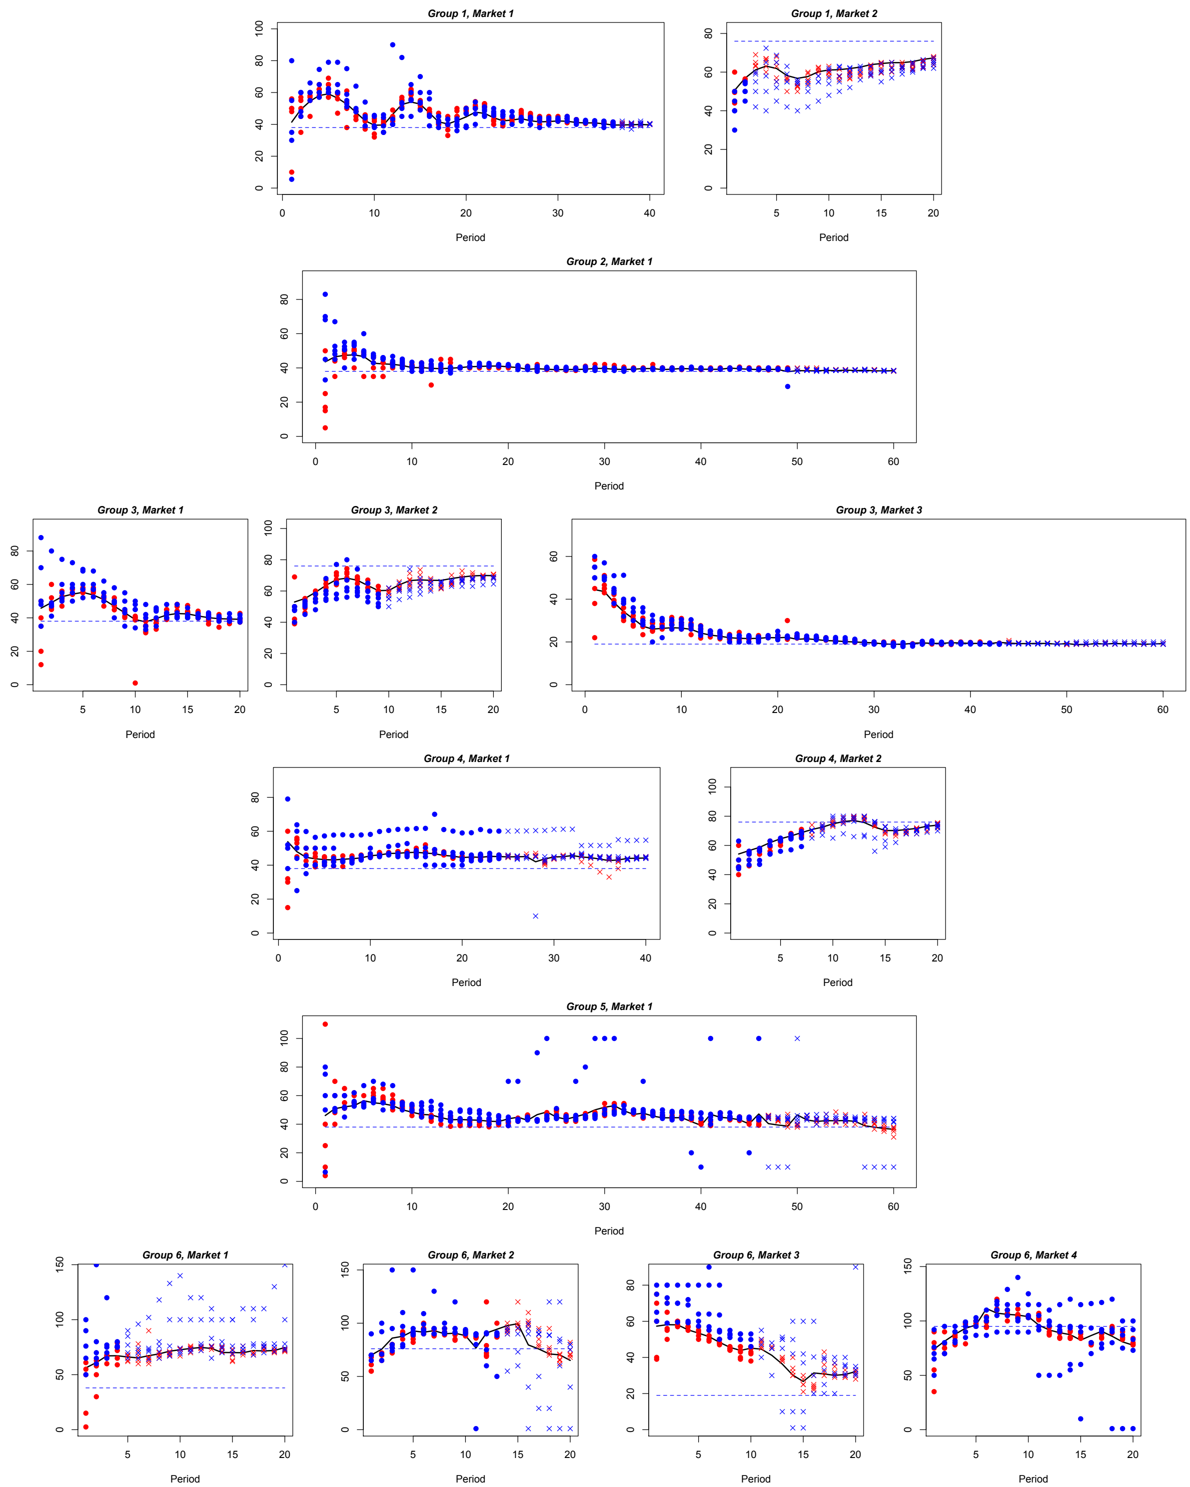

Figure 12: Treatment M50: 50% $T = 1$/50% $T = 10$ (6 groups)

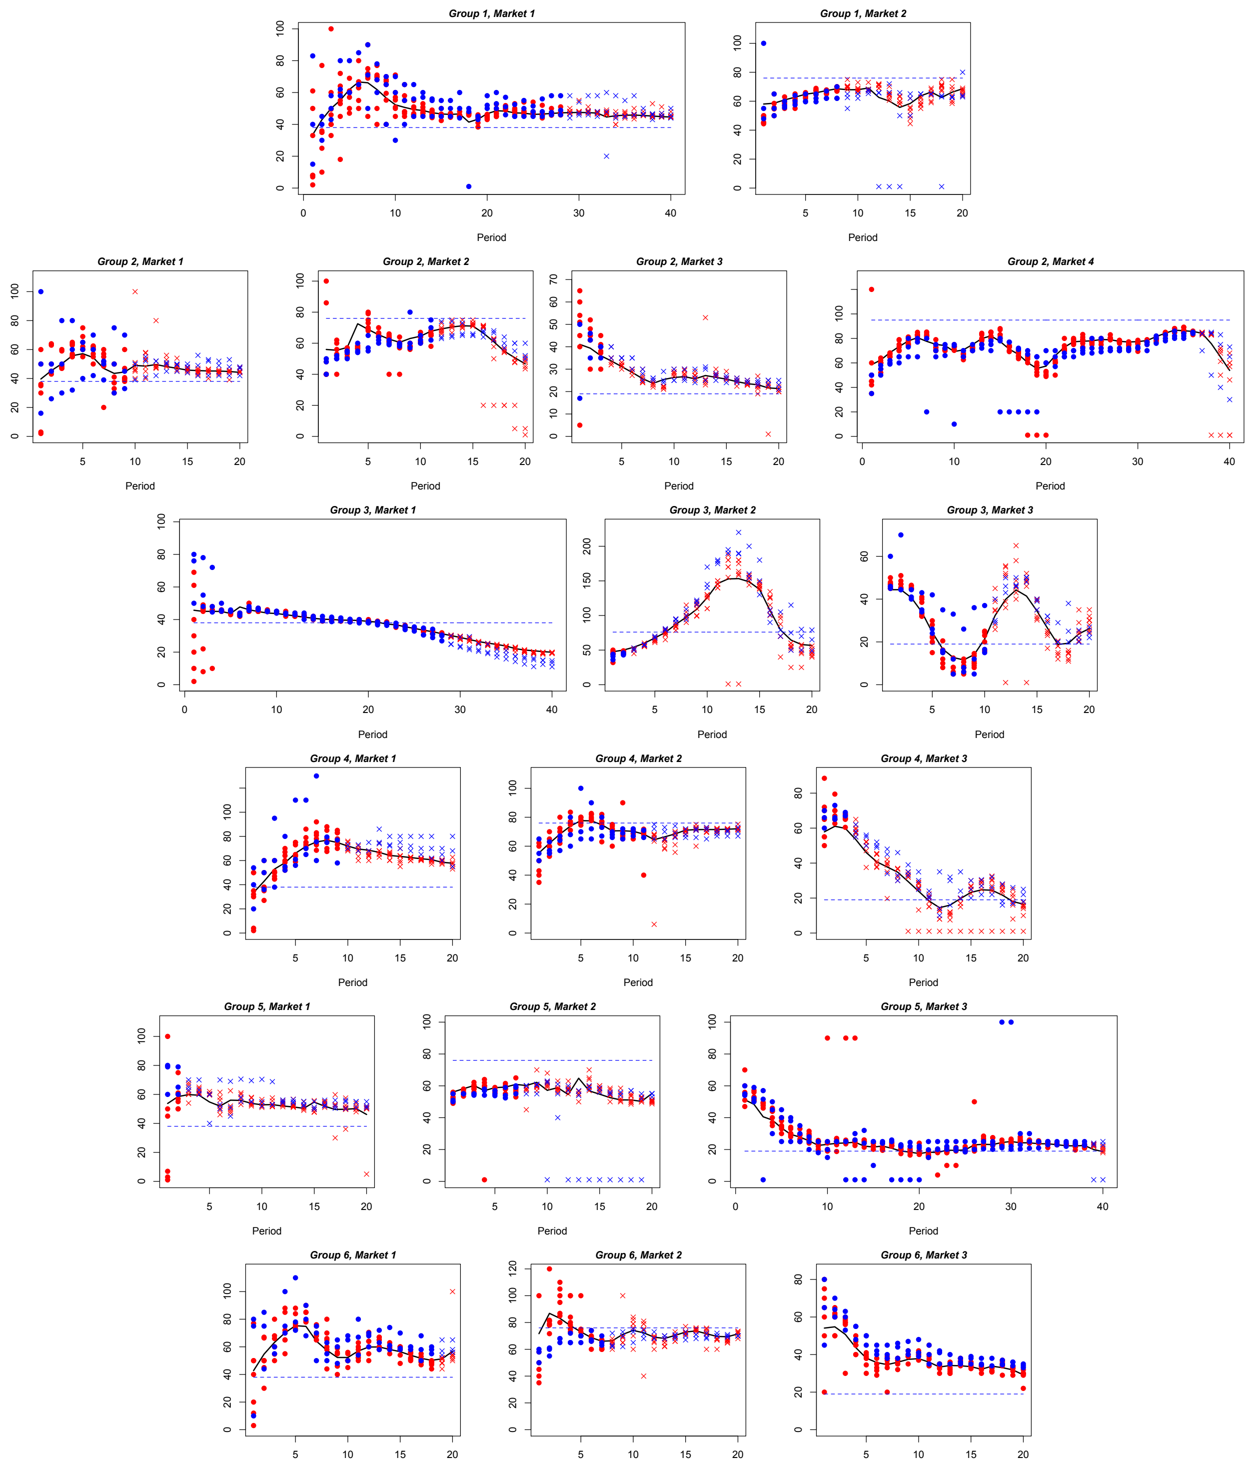

Figure 13: Treatment M70: 70% $T = 1$/30% $T = 10$.

Conclusions

- REE requires a story for how expectations are coordinated.
- Adaptive least-squares learning (AL) by agents is one natural way to implement CCP.
- The E-stability tools make assessment of local stability of an REE under adaptive learning straightforward.
- Additional dynamics arising from the learning transition, constant gain, misspecification and model selection can give interesting and plausible learning dynamics.

Implications of AL for policy:

– For interest-rate rules the Taylor principle $\chi_\pi > 1$ is important for enhancing stability under AL.

– The unintended low-inflation steady state created by the ZLB is not locally stable under AL. Large pessimistic expectations shocks can lead to large recessions, deflation and stagnation.

– In severe recessions monetary policy may needed to be supplemented by fiscal stimulus.

– The Neo-Fisherian policy of pegging the interest rate at a level consistent with the desired π^* leads to instability under AL.

Implications of AL for endogenous fluctuations

– There are cases in the NK model in which stationary sunspot equilibria are stable under learning (NRSE).

– Forward-looking monetary policy should obey the Taylor principle but must avoid over-reaction to prevent NRSE.

– Experiments indicate that sunspot equilibria can arise in the lab.

AL and asset pricing

– The fundamental price is stable under AL, but if agents put significant weight on recent data (constant-gain learning) then "escapes" can occur in which asset price bubbles periodically arise and crash.

– Experiments indicate that deviations of asset prices from the fundamental price is less likely if enough agents have long as opposed to short-horizon decision rules.