TOPICS IN MACROECONOMICS: MODELLING INFORMATION, LEARNING AND EXPECTATIONS

LECTURE NOTES 8

KRISTOFFER P. NIMARK

1. Bounded rationality and adaptive learning

In these note we will discuss an approach to bounded rationality known as Recursive Least Squares (RLS) learning where we relax the assumption that agents know the structure of the economy and *"put the agents and the econometrician on the same footing"* (Sargent 1993). Evans and Honkapohja (2001) is a good (and exhaustive) textbook treatment of the topic. The basic idea is well capture by the previous quote: We build models populated with econometricians who run regressions to form expectations. It's original motivation was to find out whether rational expectations equilibria (REE) are learnable in the sense that agents equipped with a suitable functional form for their regression model will discover the true parameters of the model with access to a long enough history of data. This is more involved question than running recursive OLS in other settings, since the system is self-referential. That is, agents' actions depend on their regression estimates, which in turn depend on their actions since these influence observables through their effect on expectations. It turns out that in many setting, yes, agents will discover the REE by running recursive least squares learning. Roughly speaking, what is required for beliefs to converge to the REE is that agents use a model with a functional form that nest the data generating process in REE and that their initial beliefs are not too far away from those of the REE.

Examples of other questions that has been tackled within this framework are for instance how optimal monetary policy should be conducted in an environment where the private

Date: October 23, 2008.

sector is learning, and more recently, learning has also been proposed as a mechanism for generating dynamics that match aggregate data (e.g. Milani 2006 and Eusepi and Preston 2008)

1.1. **How is learning different from imperfect information.** In this class, we have extensively studied models that can be put in the form

$$X_t = AX_{t-1} + C\mathbf{u}_t \tag{1.1}$$
$$Z_t = DX_t + \mathbf{v}_t$$

and to solve these models we assumed that agents knew the structure of the economy, i.e. agents knew A, C, D and the Σ_{vv} but could not observe the state X_t directly. In the learning literature, these assumptions are flipped: Agents can observe the state X_t (or what they perceive to be the state) but do not completely know the structure of the economy, i.e. they do not know the matrices A, C, D and the Σ_{vv}. Instead, they will form expectations based on running recursive regressions and take actions as if these expectations were optimal.

1.2. **Two important concepts: PLM and ALM.** In the noisy RBC model we discussed in the imperfect information part of the course, we had to introduce a new state vector to the model in order to describe the complete dynamics of the system. This extra state vector was the agents' estimate of the true state. We need to do a similar thing here, in that we need to define both an Actual Law of Motion (ALM) and a Perceived Law of Motion (PLM). The PLM is a parameterized equation describing how agents form expectations and the ALM is an equation describing how the endogenous variables actually evolve. The ALM is a function of the PLM. In addition, in a rational expectations equilibria, the PLM and the ALM coincide.

2. Recursive Least Squares Learning: A simple example

As the an illustration, consider the price setting "model" (apparently, this can be thought of as a reduced form of the price process for a partial equilibrium model with production delays, see Evans and Honkpohja 2001)

$$p_t = \mu + \alpha E_{t-1} p_t + \eta_t \tag{2.1}$$

$$\eta_t \sim N\left(0, \sigma_\eta^2\right) \tag{2.2}$$

In the rational expectations equilibria, the price is a constant plus white nois error

$$p_t = \frac{\mu}{1-\alpha} + \eta_t \quad (REE) \tag{2.3}$$

To solve the model conjecture that the agents' perceived law of motion is

$$p_t = m_{t-1} + e_t \qquad (PLM) \tag{2.4}$$

that is, the agents has the correct functional form but need to estimate the constant m_t. The actual law of motion is then given by replacing the expectations in (2.1) with the price expectation implied by the PLM. This yields

$$p_t = \mu + \alpha m_{t-1} + \eta_t \qquad (ALM) \tag{2.5}$$

Agents use least squares to form a belief about m_t

$$m_t = t^{-1} \sum_{s=0}^{t-1} p_{t-s} \tag{2.6}$$

which we can formulate recursively as

$$m_t = m_{t-1} + t^{-1}\left(p_t - m_{t-1}\right) \tag{2.7}$$

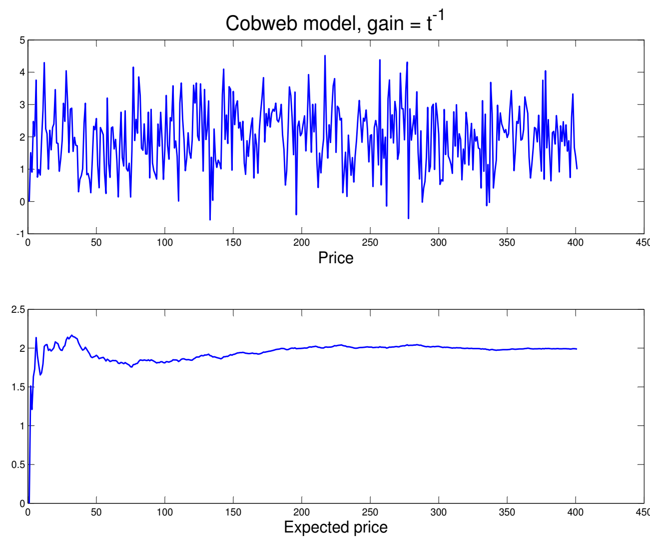

FIGURE 1

We can the nput the ALM and the updating equation of agents PLM into a single system

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu t^{-1} \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha t^{-1} - t^{-1} \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ t^{-1} \end{bmatrix} \eta_t \tag{2.8}$$

and simulate it. The actual price and the evolution of m_t is illustrated in Figure 1 (with $\mu = 1$ and $\alpha = .5$). As we can see from the figure, the price moves around quite a bit, but this is mostly due to the innovation η_t. Agents estimate of the mean (which is also the expected price in this simple model) converges quite rapidly towards 2, which is also the rational expectation of the price.

2.1. **Constant gain.** Another learning mechanism, that is similar to recursive least squares is "constant gain learning". Instead of putting equal weight on all observations, constant gain learning discounts old observations. This makes sense if agents suspect that they live in an unstable environment with drifting parameters. The recursive updating equation under

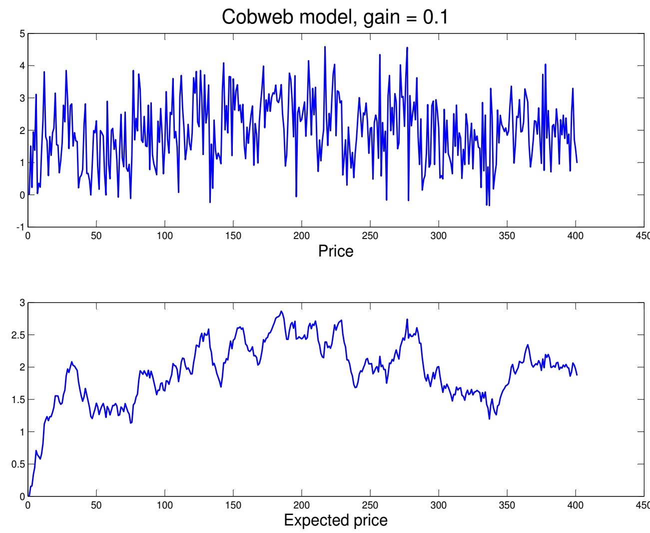

FIGURE 2

constant gain learning for the simple example above is given by

$$m_t = m_{t-1} + \gamma \left(p_t - m_{t-1}\right) \tag{2.9}$$

where $\gamma : 0 < \gamma < 1$. We have thus replaced the weight t^{-1} on the innovation with a constant (which explains the name). We can simulate the model with constant gain learning in the same way as before

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu\gamma \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1 + \alpha\gamma - \gamma \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ \gamma \end{bmatrix} \eta_t \tag{2.10}$$

As we can see from Figure 2, the expected price does not settle down, but instead moves around also in the later periods. That is because innovations always have some postive weight in the constant gain algorithm, while it has a weight that tends to 0 as time passes in the RLS learning algortihm.

Constant gain is common in empirical (or quasi-empirical) work since it solves two problems:

- With constant gain we do not need to take a stand on when $t = 0$ is (though Ed Prescott has confidently stated in another context that 1947 was year zero).
- Constant gain results in a stationary distribution of p_t and m_t, which sometimes makes estimation easier.

2.2. **Constant gain and deterministic systems.** We saw above that the agents' estimate m_t did not converge under constant gain learning. There is one exception to this that might be useful to know about: When there are no "true" innovations in the model, that is, if $\sigma_\eta^2 = 0$, the system converges also with constant gain. This can be demonstrated by simulating the system

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu\gamma \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha\gamma-\gamma \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} \tag{2.11}$$

and as we can see in Figure 3 the sequence $\{m_t\}_{t=1}^T$ tends to the REE solution $\frac{\mu}{1-\alpha}$ as $T \to \infty$. It can also be seen by the fact that the eigenvalues of the matrix

$$\begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha\gamma-\gamma \end{bmatrix} \tag{2.12}$$

are zero and $1+\alpha\gamma-\gamma$ which is smaller than unity in absolute value as long as $0 \leq \gamma, \alpha < 1$

3. The Cob-Web Model

We can use the somewhat more complicated, but still rather old-skool, Cob-Web model (3.1)

$$p_t = \mu + \alpha E_{t-1} p_t + \delta w_{t-1} + \eta_t \tag{3.1}$$

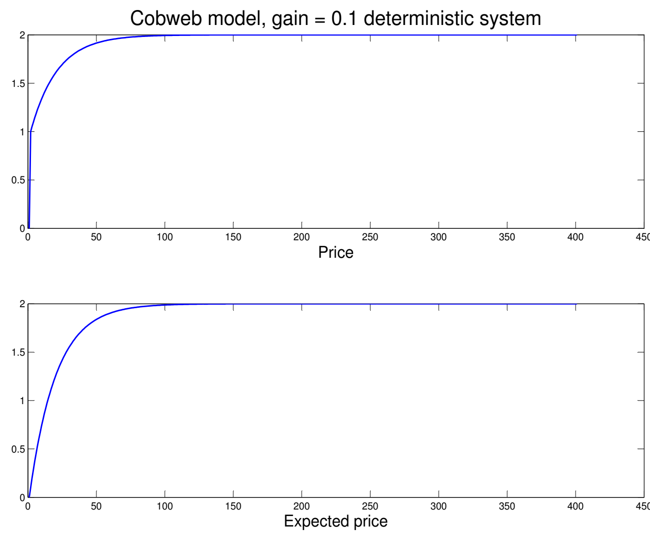

FIGURE 3

to illustrate how the RLS algorithm works when agents' PLM has more than one parameter. The REE of the Cob-Web model is given by

$$p_t = \frac{\mu}{1-\alpha} + \frac{\delta}{1-\alpha} w_{t-1} + \eta_{t-1} \quad (REE) \tag{3.2}$$

Agents fit a PLM that again nests the REE

$$p_t = a_{t-1} + b_{t-1} w_{t-1} + e_t \quad (PLM)$$

Plugging the PLM into the structural model gives the ALM

$$p_t = \mu + \alpha \left(a_{t-1} + b_{t-1} w_{t-1} \right) + \delta w_{t-1} + \eta_t \qquad (ALM) \tag{3.3}$$

To set up the RLS algorithm, define a vector ϕ_t containing the parameters of the PLM

$$\phi_t = \begin{pmatrix} a_t & b_t \end{pmatrix}' \tag{3.4}$$

and vector of observables as

$$z_t = \begin{pmatrix} 1 & w_t \end{pmatrix}' \tag{3.5}$$

Normally we would estimate ϕ_t by OLS

$$\phi_t = \left(\sum_{s=1}^{t} z_s z_s'\right)^{-1} \left(\sum_{s=1}^{t} z_s p_s\right) \tag{3.6}$$

but we can equivalently estimate ϕ_t recursively using

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(p_t - \phi_{t-1}' z_{t-1}\right) \tag{3.7}$$

$$R_t = R_{t-1} + t^{-1} \left(z_{t-1} z_{t-1}' - R_{t-1}\right) \tag{3.8}$$

where

$$R_t^{-1} = \left(\sum_{s=1}^{t} z_s z_s'\right)^{-1} \tag{3.9}$$

3.1. **E-Stability.** E-stability is a property learning models may or may not have that is closely related to whether agents can discover a REE through least squares learning. Checking for E-stability is a way to check if the parameters in the PLM converge or not and involves stochastic approximation methods, and it works in the following way. Take our PLM and ALM from the Cob-Web model above

$$p_t = a_{t-1} + b_{t-1} w_{t-1} + e_t \quad (PLM) \tag{3.10}$$

$$p_t = \mu + \alpha \left(a_{t-1} + b_{t-1} w_{t-1}\right) + \delta w_{t-1} + \eta_t \qquad (ALM) \tag{3.11}$$

Stochastic approximation techniques can then be used to show that the asymptotic behavior of ϕ in our discrete time setting is shared by a differential equation given by

$$\frac{d\phi}{d\tau} = T(\phi) - \phi \tag{3.12}$$

where $T(\phi)$ is defined as the vector of parameters in the ALM and ϕ is defined as the vector of parameters of the PLM so that

$$T(\phi) \equiv \begin{pmatrix} \mu \\ \delta \end{pmatrix} + \begin{pmatrix} \alpha & 0 \\ 0 & \alpha \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{3.13}$$

and

$$\phi = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{3.14}$$

Plugging into the differential equation (3.12)

$$\begin{aligned} \frac{d\phi}{d\tau} &= T(\phi) - \phi & (3.15) \\ &= \begin{pmatrix} \mu \\ \delta \end{pmatrix} + \begin{pmatrix} \alpha - 1 & 0 \\ 0 & \alpha - 1 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} & (3.16) \end{aligned}$$

The system will be stable if

$$eig \begin{pmatrix} \alpha - 1 & 0 \\ 0 & \alpha - 1 \end{pmatrix} < 0 \tag{3.17}$$

so that $PLM \rightarrow ALM$ with probability 1 if $T(\phi) - \phi$ is a stable system and with probability 0 if it is unstable. (For more details see Evans and Honkapohja (2001)). It is worth noting that even though REE $\implies ALM = PLM$ the converse is not true. In some models, there can exist self-confirming equilibria that are different from the REE. That is, agents do not observe anything that suggest that their model is misspecified, but if they knew the true model they would choose a different action. Apart from checking convergence properties of the PLM, E-stability has also been proposed as a equilibrium selection device in models with multiple equilibria. If there are many equilibria, of which only one is E-stable, the E-stable equilibria is argued to be the more likely equilibrium to be observed in reality.

4. In what sense (if any) is learning optimal?

When agents use RLS, even though they do not know the true model they use the information in the history of observations optimally. This is not necessarily the case with constant gain learning. However, it can be shown that constant gain learning is a close approximation to optimal learning when parameters are truly time varying. Perhaps these two points can be understood better by reformulating the RLS algorithm as a Kalman filter problem.

$$X_{t|t} = AX_{t-1|t-1} + K_t \left(Z_t - AX_{t-1|t-1} \right) \tag{4.1}$$

The Kalman updating equation (4.1) looks a bit like the RLS updating equation

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(p_t - \phi_{t-1}' z_{t-1} \right) \tag{4.2}$$

Can we redefine state space to conform to estimating ϕ_t instead of X_t?

Assume that the true parameters ϕ_t follow a random walk

$$\phi_t = \phi_{t-1} + \varepsilon_t \tag{4.3}$$

and that as in the cobweb model agents observe prices

$$p_t = z_t' \phi_t + e_t \tag{4.4}$$

The Kalman gain for this system is just a special case of what we have done before

$$\begin{aligned} K_t &= P_{t|t-1} z_t \left(z_t' P_{t|t-1} z_t + t^{-1} \Sigma \left[e_t e_t \right] \right)^{-1} && (4.5) \\ P_{t|t-1} &= P_{t-1|t-2} - && (4.6) \\ & \quad P_{t-1|t-2} z_t \left(z_t' P_{t-1|t-2} z_t + t^{-1} \Sigma \left[e_t e_t \right] \right)^{-1} z_t' P_{t-1|t-2} \end{aligned}$$

and the updating equation for the parameterrs in the agents PLM $\phi_{t|t}$ is given by

$$\phi_{t|t} = \phi_{t-1|t-1} + K_t \left(p_t - \phi_{t-1|t-1} \right) \tag{4.7}$$

Evans, Honkapohja and Williams (2008) show that if the innovations to the parameters are much smaller than the innovations to the the price, the optimal Kalman gain updating equation (4.1) can be approximated by the constant gain updating equation

$$\phi_t = \phi_{t-1} + \gamma R_t^{-1} z_{t-1} \left(p_t - \phi'_{t-1} z_{t-1} \right) \tag{4.8}$$

with γ constant. That is if

$$E\left[\varepsilon_t \varepsilon'_t\right] << E\left[e_t e'_t\right] \tag{4.9}$$

we have that

$$\phi_{t-1|t-1} + K_t \left(p_t - \phi_{t-1|t-1} \right) \approx \phi_{t-1} + \gamma R_t^{-1} z_{t-1} \left(p_t - \phi'_{t-1} z_{t-1} \right) \tag{4.10}$$

for an appropriately chosen γ.

5. Summing up:

- Agents behave as econometricians.
- Agents can discover REE if fitting the correct functional form and model is E-stable.
- Learning is optimal in the sense of no information wasted if $\gamma_t = t^{-1}$ and actual parameters fixed.
- Constant gain not optimal, but makes more sense if there are structural breaks or parameter drift.

References

[1] Evans and Honkpohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton University Press

[2] Evans, Honkpohja and Williams, forthcoming, Generalized Stochastic Gradient Learning, *International Economic Review.*

[3] Eusepi and Preston, 2008, Expectations, Learning and Business Cycle Fluctuations ,*NBER Working Paper* no. 14181.

[4] Milani, 2006, Expectations, Learning and Macroeconomic Persistence, *Journal of Monetary Economics*.

[5] Sargent, Thomas J., 1993 *Bounded Rationality in Macroeconomics*, Arne Ryde lecture series.