TOPICS IN MACROECONOMICS: MODELLING INFORMATION, LEARNING AND EXPECTATIONS

KRISTOFFER P. NIMARK

Bounded rationality and adaptive learning

In these notes 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 sector is learning, and more recently, learning has also been proposed as a mechanism for

Date: November 24, 2010.

generating dynamics that match aggregate data (e.g. Milani 2006 and Eusepi and Preston 2008)

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

$$X_t = AX_{t-1} + C\mathbf{u}_t \tag{0.1}$$

$$Z_t = DX_t + \mathbf{v}_t$$

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.

0.2. **A Brief History of Expectations in Macroeconomics.** In the 1960's many macro models contained expectations terms and their importance for wages, inflation and output was already well understood.[1] Expectations were often modeled as adaptive, in the sense that expectations of future values of a variable, depended on lagged values of the same variable. This approach had some advantages: It seemed to fit the data rather well, and with appropriate restrictions on the lag coefficients, expectations converged to steady state levels in models where such a steady state existed. This was in spite of the fact that the seed to the rational expectations revolution had already been sowed by John Muth (1961). In the paper that coined the term "rational expectations", Muth asked the following question: for a given adaptive expectations scheme, what stochastic system would make these expectations optimal? Instead of taking an expectations scheme and finding a model that would make

[1] See for instance Samuelson and Solow (1960).

it rational, Sargent (1971) and Lucas (1972) reversed the process and let the optimal expectations be determined by the model. Expectations then disappear as free variables from macroeconomic models.

The rational expectations revolution has had a great impact on macroeconomics, and the reasons are many: It is a beautiful application of Occam's razor in that it reduces the number of free parameters (and variables) while increasing the internal consistency of models by extending rationality also to how expectations are formed and how available information is used. This economy of parameters comes from the fact that in rational expectations equilibria, the true model coincide with the model used by the agents inside the model to form expectations. This is an appealing feature of rational expectations if one believes that we as economists should not build models where we know more about the economy than the agents inside the model. Particularly, and as forcefully argued by Lucas (1976), we should not use such models for policy advice that presumes that policy makers will be able to systematically fool the public.

The "communism of models" also gave a lot of power to rational expectations econometrics, since it placed stringent cross equation restriction on the data. In addition, because rational expectations econometrics also economized on free parameters relative to previous approaches, it allowed for sharpened inference.[2]

A lot of the power of rational expectations thus comes from that it equates the expectations of agents inside the model to the mathematical expectations implied by the model (and perhaps subject to informational constraints). This seemed to attribute a great deal of knowledge to the agents inside the model since it is assumed that not only do agents know the structure of the economy, but they also know the parameters. A natural question to ask is "How do agents inside the model come about this knowledge?" This question was the initial motivation for the learning literature. Perhaps we should build models where agents

[2]This sharp inference often led to rejections of the rational expectations models of the first generation. In an interview Sargent is quoting Lucas as saying that "These likelihood ratios are rejecting too many good models!", see Evans and Honkapohja (2005).

know as much as we do as economists, but not *more* than we do as *econometricians*? If given a functional form, can agents inside the model discover the rational expectations equilibrium if given a sufficiently long history of data? That is, if we populate our models with econometricians, can they discover and converge to the rational expectations equilibrium? This is a more involved question than what may appear at first glance. As economists, or econometricians, we use asymptotic theory to motivate consistency of estimators etc. However, these statistical techniques are not directly transferable to self-referential learning models where expectations determine outcomes which in turn affect the learning process which in turn affect expectations and so on. In an early contribution, Bray (1983) asks exactly this question in an asset pricing model similar to that of the well-known model of Grossman and Stiglitz (1983). Bray's contribution will be discussed in more detail below. We first give an overview of some important concepts and how learning has most commonly implemented in the literature.[3] Readers who are already familiar with terms such as anticipated utility, recursive least squares learning, constant gain learning and self-confirming equilibria can skip this overview.

0.3. **Anticipated Utility.** Most economic models that deal with uncertainty assume that agents maximize expected utility (i.e. von Neumann-Morgenstern utility). In dynamic models, expected utility is often expressed as the discounted present value of future period-by-period utility, where the expected discounted value depend on (among other things) the transition laws of the economy and the stochastic processes that shocks the model. In the rational expectations literature, expectations of future utility are conditional on the true model and the information set of the agent. In that framework, an action is optimal if it maximizes the expected utility of the agent according to this measure.

In the macro learning literature it is common to instead of expected utility use what has become known as anticipated utility (see Kreps 1998). It is in many ways similar to

[3]A lot of the material of the next section can be found in the text book by Evans and Honkapohja (2001).

expected utility, except for two properties: (i) Agents do not know the true model and, (ii) even though agents know that they are learning about the parameters, they do not take into account that they will continue to learn in the future when the choose actions today. An optimal action according to an anticipated utility maximizer is the action that would maximize expected utility if the current beliefs of the agent were the true model (and time invariant). Cogley and Sargent (2008) demonstrate in a life cycle model of consumption and saving that anticipated utility can approximate expected utility very closely. Kreps (1998) argues that it is a reasonable behavior in settings where it is difficult to figure out the true structure of the world and thus would constitute a rational action.

0.4. **Actual and Perceived Laws of Motion.** In a rational expectations model, there is only one model, that is, the "true" model. In a model of learning, the true model is referred to as the Actual Law of Motion (ALM). This name distinguishes it from the Perceived Law of Motion (PLM), which is the model used by the agents inside the model to form expectations. A simple example can make the distinction clear. Consider the price setting model

$$p_t = \mu + \alpha \widehat{E}_{t-1} p_t + \varepsilon_t \tag{0.2}$$

$$\varepsilon_t \sim N\left(0, \sigma^2\right) \tag{0.3}$$

where p_t is the price level, μ and α are parameters and $\widehat{E}_{t-1}$ is the operator denoting agents' (not necessarily rational) expectations at time $t-1$. In the rational expectations equilibria (REE), the price is a constant plus white noise error

$$p_t = \frac{\mu}{1-\alpha} + \varepsilon_t \quad (REE) \tag{0.4}$$

In a model of learning, instead of solving for the fixed point (1.4) to find the expectation of (1.2), it is common to conjecture that agents' perceived law of motion is of a specific functional form. Often, but not always, the functional form is assumed to be the same as the functional form of the rational expectations equilibrium. In our case, that means conjecturing a perceived law of motion of the form

$$p_t = m_{t-1} + e_t \qquad (PLM) \tag{0.5}$$

Even though agents here are endowed with a perceived law of motion of the same functional form as the rational expectations equilibrium, they still need to estimate the constant m_t. For a given m_t, the actual law of motion can be found by replacing the expectations in (1.2) with the price expectation implied by the perceived law of motion. This yields

$$p_t = \mu + \alpha m_{t-1} + \eta_t \qquad (ALM) \tag{0.6}$$

The role of economic theory is thus to determine the mapping between the perceived law of motion and the actual law of motion. In a rational expectations equilibria, the ALM equals the PLM so that

$$m_{t-1} = \frac{\mu}{1-\alpha} \tag{0.7}$$

for all t. The literature on learning deviates from the rational expectations literature in that at least temporarily, or for some t, the equality does not hold. Instead, agents inside the model form expectations m_t of the mean of the process using some type of learning scheme. Depending on how agents construct estimates of m_t, a learning schemes can be classified as either decreasing gain learning or constant gain learning.

0.5. **Recursive Least Squares/Decreasing gain learning.** A common type of learning is to assume that agents use simple OLS to form beliefs about the parameters in their perceived law of motion. In this simple example, the OLS estimate of the mean m is

$$m_t = t^{-1} \sum_{s=0}^{t-1} p_{t-s} \tag{0.8}$$

It is common to use the equivalent recursive formulation

$$m_t = m_{t-1} + t^{-1} (p_t - m_{t-1}) \tag{0.9}$$

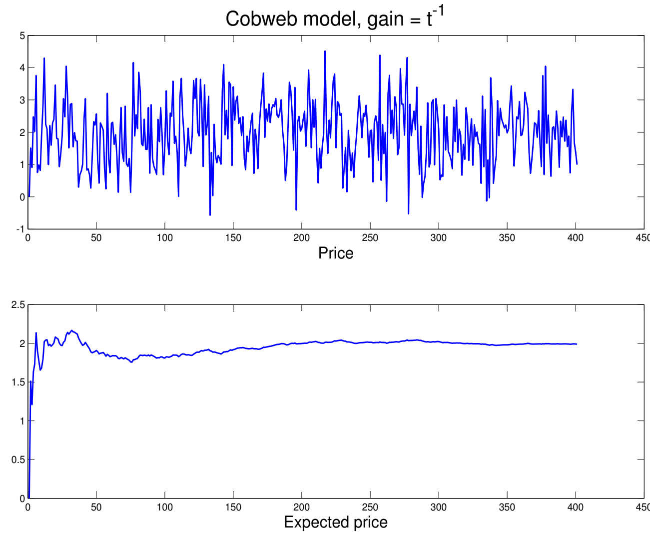

FIGURE 1. Convergence of beliefs with decreasing gain learning

(and hence the name recursive least squares). The term in brackets is the innovation, or the surprise component, of the price observation in period t. If agents by chance happen to observe exactly what they expected, i.e. if $p_t = m_{t-1}$, then the estimate does not change so that $m_t = m_{t-1}$, i.e. beliefs about the mean m only change when agents are surprised by what they observe. Since the weight t^{-1} put on each new observation is decreasing as t grows larger (as the history of observations grows longer), agents stop updating their belief about the mean asymptotically. This is why this type of learning is also known as *decreasing gain learning.*

To describe the complete model we put the ALM and the updating equation (1.9) of agents PLM into the single system

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu t^{-1} \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1 + \alpha t^{-1} - t^{-1} \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ t^{-1} \end{bmatrix} \eta_t \qquad (0.10)$$

A simulation of 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 top panel of 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 can be seen in the bottom panel of Figure 1. $\widehat{E}_{t-1}p_t = 2$ is also the rational expectation of the price.

Of course, recursive least squares learning can be implemented in more complex models than in this example. The important aspect of recursive least squares is that the weight on observations is decreasing as time goes on. This is what allows the coefficients in agents perceived laws of motion to converge.

0.6. **Constant gain learning.** 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 and can be motivated as an approximation to fully rational updating in a model with parameter drift (see Evans, Honkapohja and Williams, forthcoming). The recursive updating equation under constant gain learning for the simple example above is given by

$$m_t = m_{t-1} + \gamma\left(p_t - m_{t-1}\right) \tag{0.11}$$

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

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu\gamma \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha\gamma-\gamma \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ \gamma \end{bmatrix} \eta_t \tag{0.12}$$

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 positive

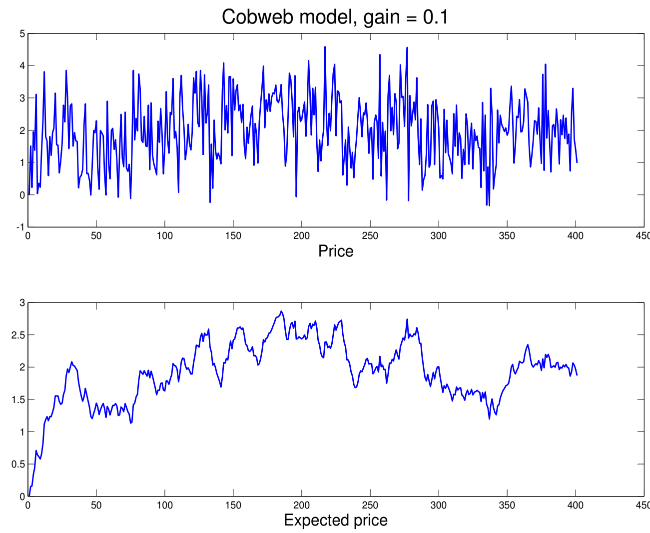

FIGURE 2

weight in the constant gain algorithm, while it has a weight that tends to zero as time passes in the decreasing gain learning algorithm.

Constant gain is common in empirical (or quasi-empirical) work since it solves two problems. The econometrician do not have to take a stand on when period zero is[4] and under appropriate parameter restrictions, constant gain learning also results in a system that converges to a stationary distribution of p_t and m_t, which may facilitate estimation.

0.7. **Self-Confirming Equilibrium.** A rational expectations equilibrium is a fixed point of a mapping from a perceived law of motion to an actual law of motion. A *self confirming equilibrium* is a somewhat weaker concept that nests rational expectations as a special case. In a self-confirming equilibrium, agents' perceived law of motion coincide with the actual law of motion along the equilibrium path. Off-equilibrium behavior of the actual and perceived law of motion may differ though, but since off-equilibrium behavior is never observed, there is nothing to indicate to agents that their model is misspecified. The distinction between

[4]Having to choose when period zero is not a problem that bothers everyone. Ed Prescott has confidently stated in another context that 1947 was year zero. (Sorry, this is conversationally transmitted knowledge, so no reference.)

rational expectations equilibrium and a self-confirming equilibrium is particularly important if a large agent, say a central bank, has a misspecified model in mind and knowledge of the true model would imply different behavior.

Lucas (1973) can be interpreted as a story about a self confirming equilibrium where policy makers estimate a Phillips curve type of relationship and mistakenly believe that there is a trade off between systematic inflation and unemployment when in fact there according to Lucas is no such trade off. Lucas argued that this was concealed to policymakers by the fact that there is a relationship between *surprise* inflation and unemployment. However, if policy makers never choose a policy that would reveal this fact to them, this cannot be discovered through sample correlations, no matter how large the available sample is.

Self-confirming equilibria does not have to be this subtle though. It has also been used to model learning with misspecified models where enough data would indeed reveal the misspecification, for instance if agents fitted a different functional form to the data.[5]

0.8. **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

$$\begin{bmatrix} p_t \\ m_t \end{bmatrix} = \begin{bmatrix} \mu \\ \mu\gamma \end{bmatrix} + \begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha\gamma-\gamma \end{bmatrix} \begin{bmatrix} p_{t-1} \\ m_{t-1} \end{bmatrix} \tag{0.13}$$

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

$$\begin{bmatrix} 0 & \alpha \\ 0 & 1+\alpha\gamma-\gamma \end{bmatrix} \tag{0.14}$$

[5]See Chapter 13 in Evans and Honkapohja (2001) for some examples.

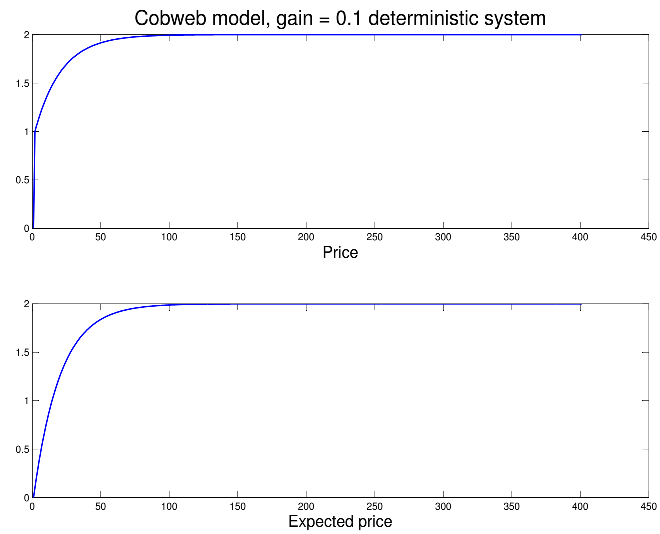

FIGURE 3

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

0.9. **The Cob-Web Model.** We can use the somewhat more complicated, but still rather old-skool, Cob-Web model (1.15)

$$p_t = \mu + \alpha E_{t-1} p_t + \delta w_{t-1} + \eta_t \tag{0.15}$$

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

$$p_t = \frac{\mu}{1-\alpha} + \frac{\delta}{1-\alpha} w_{t-1} + \eta_{t-1} \quad (REE) \tag{0.16}$$

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

$$p_t = \mu + \alpha \left(a_{t-1} + b_{t-1} w_{t-1}\right) + \delta w_{t-1} + \eta_t \qquad (ALM) \tag{0.17}$$

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

$$\phi_t = \begin{pmatrix} a_t & b_t \end{pmatrix}' \tag{0.18}$$

and vector of observables as

$$z_t = \begin{pmatrix} 1 & w_t \end{pmatrix}' \tag{0.19}$$

Normally we would estimate ϕ_t by OLS

$$\phi_t = \left(\sum_{s=1}^{t} z_s z_s' \right)^{-1} \left(\sum_{s=1}^{t} z_s p_s \right) \tag{0.20}$$

but we can equivalently estimate ϕ_t recursively using

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(p_t - \phi_{t-1}' z_{t-1} \right) \tag{0.21}$$

$$R_t = R_{t-1} + t^{-1} \left(z_{t-1} z_{t-1}' - R_{t-1} \right) \tag{0.22}$$

where

$$R_t^{-1} = \left(\sum_{s=1}^{t} z_s z_s' \right)^{-1} \tag{0.23}$$

1. Convergence to Rational Expectations Equilibria

In the late 1970s and early 1980s there was still a fair amount of scepticism about rational expectations, as it seemed to impute a lot knowledge to agents. Informal justifications of rational expectations based on the argument that agents will discover if they make systematic errors can be exemplified by a quote from Grossman and Stiglitz (1976)

"an individual will eventually observe that the frequency distribution of returns, conditional on the observable variables, is different from the subjective distribution, and accordingly, ought to revise his expectation"

This quote was taken as the starting point of a seminal paper by Margaret Bray (1982). In that paper, Bray demonstrates that under some conditions, and using a particular model in which (only a fraction of) agents are boundedly rational and behaves like econometricians, the system can converge to the rational expectations equilibrium. Bray proved her results for two different learning technologies. In the first, agents use a fixed coefficient rule to forecast while using the resulting outcomes to estimate a ”new” rule. Once the estimates of the parameters of the new rule have converged, the new rule is used as the new forecasting rule that decisions from then on are based on. These decisions influence outcomes and these outcomes are used to estimate a new rule, and so on. This procedure is then repeated until convergence. In the second set-up, agents revise their forecasting rule every time a new data point becomes available. (This second set up thus resembles recursive least squares learning as described above.) This was an important result and showed formally that the intuition conveyed in the quote above could hold up to a more rigorous analysis. Bray took this result to give some support for the rational expectations hypothesis, but cautioned that her results hold only asymptotically and that “rational expectations are, if anything, a long run rather than a short run phenomenon”.

It is also a more involved analysis than suggested by the quote above. The fact that learning affects decisions, which in turn affect observables that are used in estimation makes the system self-referential and invalidates a lot of statistical asymptotic theory. This fact also distinguishes boundedly rational learning from Bayesian (or rational) learning which by repeated application of Bayes' rule can be shown to asymptotically converge to the truth. However, the knowledge required for fully Bayesian learning are almost as large as the requirement for full information rational expectations, which is one reason why the result of Bray was considered important as a justification for rational expectations.

While the results of Bray (1982) were important, they were also derived in a rather “simple and special” model, to use Bray's own description. Only a fraction of agents are boundedly rational, the other fraction are fully informed rational agents. The dynamic structure of the system is such that all exogenous processes are white noise, which simplifies the analysis considerably. Bray and Savin (1986) and Fourgeaud et al (1986) analyzes the convergence properties of the somewhat more complex cobweb model and derive results that give parameter restrictions for the cobweb model that makes the model converge to the REE with probability 1.

1.1. **A General Method to Analyze Convergence of Learning Models.**

In a series of paper written by Marcet and Sargent (1988, 1989a, 1989b), results from stochastic approximation theory due to Ljung (1977) are used to derive a general procedure for analyzing convergence of learning models. The method is naturally easiest to explain using a simple model, though it is also applicable to more complex systems.

A simple example

Consider the Cob-web model

$$p_t = \mu + \alpha \widehat{E}_{t-1} p_t + \delta w_{t-1} + \eta_t \tag{1.1}$$

where w_t is an observable exogenous stochastic process and η_t is an unobservable shock. Conjecture the perceived law of motion

$$p_t = a_{t-1} + b_{t-1} w_{t-1} + e_t \quad (PLM) \tag{1.2}$$

with the resulting actual law of motion

$$p_t = \mu + \alpha \left(a_{t-1} + b_{t-1} w_{t-1} \right) + \delta w_{t-1} + \eta_t \qquad (ALM) \tag{1.3}$$

Agents are assumed to estimate the parameters in the vector $\phi = \begin{bmatrix} a & b \end{bmatrix}'$ of their perceived law of motion using least squares. This could either be done by OLS

$$\phi_t = \left(\sum_{s=1}^{t} z_s z_s' \right)^{-1} \left(\sum_{s=1}^{t} z_s p_s \right) \tag{1.4}$$

but for our purposes it usually more convenient to use the recursive but equivalent updating scheme

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(p_t - \phi_{t-1}' z_{t-1} \right) \tag{1.5}$$

$$R_t = R_{t-1} + t^{-1} \left(z_{t-1} z_{t-1}' - R_{t-1} \right) \tag{1.6}$$

where

$$R_t^{-1} = \left(\sum_{s=1}^{t} z_s z_s' \right)^{-1} \tag{1.7}$$

The actual law of motion (2.3) and the recursive updating equations (2.5) - (2.6) completely describe the dynamics of the system. However, the properties of the system in this form are hard to analyze, partly because of the non-linear mapping between observations and the actual law of motion, and partly because the dependence of the law of motion on a complete history of endogenous innovations. The contribution of Marcet and Sargent (1989a) was to demonstrate that some properties, e.g. the asymptotic behavior, of the system (2.3), (2.5) and (2.6) is shared by the differential equation

$$\frac{d\phi}{d\tau} = T(\phi) - \phi \tag{1.8}$$

where $T(\phi)$ is defined as the operator that maps the parameters of the perceived law of motion into a vector of the corresponding of parameters of the actual law of motion. The difference equation (2.8) turns out to be much more amenable to analysis. For instance, to study the convergence properties of the system we can substitute in the vector of parameters from the perceived law of motion (2.2)

$$\phi = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{1.9}$$

and the vector of parameters from the actual law of motion (2.3)

$$T(\phi) \equiv \begin{pmatrix} \mu \\ \delta \end{pmatrix} + \begin{pmatrix} \alpha & 0 \\ 0 & \alpha \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{1.10}$$

into the differential equation (2.8) to get

$$\frac{d\phi}{d\tau} = T(\phi) - \phi \tag{1.11}$$

$$= \begin{pmatrix} \mu \\ \delta \end{pmatrix} + \begin{pmatrix} \alpha - 1 & 0 \\ 0 & \alpha - 1 \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{1.12}$$

The main result of Marcet and Sargent (1989a) was to show that the system converges to a resting point with probability 1 if

$$eig \begin{pmatrix} \alpha - 1 & 0 \\ 0 & \alpha - 1 \end{pmatrix} < 0 \tag{1.13}$$

that is if the differential equation (2.8) is stable. Clearly, a resting point $\frac{d\phi}{d\tau} = 0$ implies that $T(\phi) = \phi$ or that the perceived law of motion coincides with the actual law of motion. Marcet and Sargent (1989a) also proved that the only resting point of the system is a self-confirming equilibria (though they did not call it that at the time).

Checking the stability of the matrix in (2.12) turns out to be significantly and generally more tractable than analyzing the system (2.3), (2.5) and (2.6) directly. Marcet and Sargent demonstrated through a series of examples of increasing complexity that the method is generally applicable and also works for larger systems. First, a simple non-self-referential system is analyzed followed by the models of Bray (1982), Bray and Savin (1986), a model of hyperinflation in stock prices due to Fourgeaud et al (1986), and finally, a version of Lucas and Prescott's (1971) model of investment under uncertainty. The last example is significant since it shows that the method is also applicable to analyze convergence of models where agents learn from endogenous variables, something that had not been demonstrated before.

In a closely related paper, Marcet and Sargent (1989b) demonstrate that the method can also be used to find an approximate solution to a hard-to-solve model due to Townsend (1983). The 1988 AER P&P paper (Marcet and Sargent 1988) summarizes these two papers using less technical language.

2. 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.

$$X_{t|t} = AX_{t-1|t-1} + K_t \left(Z_t - AX_{t-1|t-1} \right) \tag{2.1}$$

The Kalman updating equation (3.1) looks a bit like the RLS updating equation

$$\phi_t = \phi_{t-1} + t^{-1} R_t^{-1} z_{t-1} \left(p_t - \phi_{t-1}' z_{t-1} \right) \tag{2.2}$$

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

$$\phi_t = \phi_{t-1} + \varepsilon_t \tag{2.3}$$

and that as in the cobweb model agents observe prices

$$p_t = z_t' \phi_t + e_t \tag{2.4}$$

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

$$K_t = P_{t|t-1}z_t\left(z_t'P_{t|t-1}z_t + t^{-1}\Sigma\left[e_te_t\right]\right)^{-1} \quad (2.5)$$

$$P_{t|t-1} = P_{t-1|t-2} - P_{t-1|t-2}z_t\left(z_t'P_{t-1|t-2}z_t + t^{-1}\Sigma\left[e_te_t\right]\right)^{-1}z_t'P_{t-1|t-2} \quad (2.6)$$

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

$$\phi_{t|t} = \phi_{t-1|t-1} + K_t\left(p_t - \phi_{t-1|t-1}\right) \quad (2.7)$$

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 (3.1) can be approximated by the constant gain updating equation

$$\phi_t = \phi_{t-1} + \gamma R_t^{-1}z_{t-1}\left(p_t - \phi_{t-1}'z_{t-1}\right) \quad (2.8)$$

with γ constant. That is if

$$E\left[\varepsilon_t\varepsilon_t'\right] << E\left[e_te_t'\right] \quad (2.9)$$

we have that

$$\phi_{t-1|t-1} + K_t\left(p_t - \phi_{t-1|t-1}\right) \approx \phi_{t-1} + \gamma R_t^{-1}z_{t-1}\left(p_t - \phi_{t-1}'z_{t-1}\right) \quad (2.10)$$

for an appropriately chosen γ.

3. 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] Adam, K., A.Marcet, and J. P. Nicolini, 2008, Stock Market Volatility and Learning, *mimeo.*

[2] Barillas, F., L.P. Hansen and Thomas J. Sargent, *forthcoming*, Doubts or variability?, *Journal of Economic Theory.*

[3] Bernanke, B., Mihov, I., 1998b. Measuring monetary policy. *Quarterly Journal of Economics* 113, 869–902.

[4] Blume, L. M. Bray and D. Easley, 1986, Introduction to Stability of rational Expectations Equilibrium, *Journal of Economic Theory* 26, p313-317.

[5] Bray, Margaret, 1982, Learning, Estimation, and the Stability of Rational Expectations, *Journal of Economic Theory* 26, p318-339.

[6] Bray, M. and N.E.. Savin, 1986, Rational Expectations Equilibria, Learning and model Specification, *Econometrica* 54, p1129-1160.

[7] Bullard, J., 2006, "The Learnability Criterion and Monetary Policy, *Federal Reserve Bank of St. Louis Review*, 88, pp203—217.

[8] Cho, I-K, N. Williams, and T.J. Sargent, 2002, Escaping Nash Inflation, *Review of Economic Studies* 69, p1-40.

[9] Cogley, T. and T.J. Sargent, 2005a, Drifts and volatilities: monetary policies and outcomes in th epost WWII US, *Review of Economic Dynamics* 8, p262-302.

[10] Cogley, T. and T.J. Sargent, 2005b, The Conquest of US Inflation: Learning and Robustness to Model Uncertainty, *Review of Economic Dynamics* 8, p528-563.

[11] Cogley, T., Colacito, R. and T.J. Sargent, 2007, Benefits from U.S. Monetary Experimentation in the Days of Samuelson and Solw and Lucas, *Journal of Money Credit and Banking* 8, p67-100.

[12] Doan T, RB Litterman and CA Sims, 1984, 'Forecasting and Conditional Projection Using Realistic Prior Distributions', *Econometric Reviews*, 3(1), pp 1-100.

[13] Evans, G. and s. Honkpohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton University Press.

[14] Evans, George W. & Honkapohja, Seppo, 2005, "An Interview With Thomas J. Sargent," *Macroeconomic Dynamics*, Cambridge University Press, vol. 9(04), p561-583.

[15] Evans, G., S. Honkapohja and T.J. Sargent, 1993, On the Preservation of Deterministic Cycles when some Agents Perceive them to be Random Fluctuations, *Journal of Economic Dynamics and Control* 17, p705-721.

[16] Evans, G., S. Honkapohja and N. Williams, forthcoming, Generalized Stochastic Gradient Learning, *International Economic Review.*

[17] Eusepi, S., and B. Preston (2008): "Expectations, Learning and Business Cycle Fluctuations, *NBER working paper Nr* 14181.

[18] Fourgeaud, C., C. Gourieroux, and J. Pradel, 1986, Learning Procedures and Convergence to Rationality, *Econometrica* 54, p846–868.

[19] Fudenberg, D. and D.K. Levine, 1993, Steady State Learning and Nash Equilibrium, *Econometrica*, 61 (1993), 547-574.

[20] Gaspar, V., O. Issing, O. Tristani and D.vestin, 2006, *Imperfect knowledge and monetary policy*, Cambridge University Press.

[21] Grossman, S.I. and J.E. Stiglitz, 1983, On the impossibility of informationally efficient markets, *American Economic Review* 70, p393-408.

[22] Hansen, L.P. and T.J. Sargent, 2007, *Robustness*, Princeton University Press.

[23] Hansen, Lars Peter and Sargent, Thomas J., 1980. "Formulating and estimating dynamic linear rational expectations models," *Journal of Economic Dynamics and Control*, pp7-46.

[24] Kreps, David, 1998, Anticipated Utility and Dynamic Choice, in D.P. Jacobs, E. Kalai, and M. Kamien, eds., Frontiers of Research in Economic Theory, Cambridge University Press, p242–274.

[25] Kydland, Finn E & Prescott, Edward C, 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, vol. 85(3), pp 473-91.

[26] Ljungqvist, L. and T.J. Sargent, 2004, Recursive Macroeconomic Theory, second edition, MIT press.

[27] Lucas, Robert Jr., 1972. "Expectations and the neutrality of money," *Journal of Economic Theory*, vol. 4(2), p103-124.

[28] Lucas, Robert E, Jr, 1973. "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review*, pp326-34

[29] Lucas, Robert Jr, 1976,Econometric policy evaluation: A critique, *Carnegie-Rochester Conference Series on Public Policy* vol. 1(1), p19-46.

[30] Marcet, A., and J. P. Nicolini (2003): "Recurrent Hyperinflations and Learning," *American Economic Review*, 93, pp1476-1498.

[31] Marcet, Albert and Sargent, Thomas J, 1988. "The Fate of Systems with "Adaptive" Expectations," *American Economic Review*, vol. 78(2), p168-72.

[32] Marcet, A. and T.J. Sargent, 1989a, Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models,' *Journal of Economic Theory*, vol. 48, no. 2.

[33] Marcet, Albert and Sargent, Thomas J, 1989b, "Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information," *Journal of Political Economy*, vol. 97(6), pp1306-22.

[34] Milani, Fabio, 2007, Expectations, Learning and Macroeconomic Persistence, *Journal of Monetary Economics*, 54, pp2065-2082.

[35] Muth, John A., 1961, Rational Expectations and the Theory of Price Movements", *Econometrica* 29, no. 6 : 315-35.

[36] Samuelson, P.A., Solow, R.M., 1960, Analytical aspects of anti-inflation policy. *American Economic Review* 50, p177–184.

[37] Sargent, Thomas J, 1971, A Note on the 'Accelerationist' Controversy, *Journal of Money, Credit and Banking* vol. 3(3), p721-25.

[38] Sargent, Thomas J., 1979, Macroeconomic Theory, Academic Press, New York.

[39] Sargent, Thomas J., 1979. "A note on maximum likelihood estimation of the rational expectations model of the term structure," *Journal of Monetary Economics*, pp133-143.

[40] Sargent, Thomas J., 1987, Dynamic Macroeconomic Theory, Harvard University Press.

[41] Sargent, Thomas J., 1991, "Equilibrium with Signal Extraction from Endogenous Variables", *Journal of Economic Dynamics and Control* 15, pp245-273.

[42] Sargent, T. J. (1999), The Conquest of American Inflation. Princeton, NJ: Princeton University Press.

[43] Sargent, Thomas J., 1991, Equilibrium with Signal Extraction from Endogenous Variables, *Journal of Economic Dynamics and Control,* 15, p245-273.

[44] Sargent, T.J. and N. Williams, 2005, Impacts of Priors on Convergence and Escapes from Nash Inflation, *Review of Economic Dynamics* 8, p360-391.

[45] Sargent, T.J., N. Williams and Tao Zha, 2006, Shocks and Government Beliefs: The Rise and Fall of American Inflation, *American Economic Review* 96, p1193-1223.

[46] Sims, C.A., 1999. Drifts and breaks in monetary policy. Unpublished manuscript. Department of Economics, Princeton University.

[47] Timmermann, A. G. (1993): "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," Quarterly Journal of Economics, 108, pp1135-1145.

[48] Timmermann, A. G. (1996): "Excessive Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning," Review of Economic Studies, pp63, 523—557.

[49] Townsend, Robert M., 1983, Forecasting the Forecasts of Others, *Journal of Political Economy*, vol 91, pp546-588.