learning in macroeconomics

Expectations play a key role in macroeconomics. The assumption of rational expectations has been recently relaxed by explicit models of forecasting and model updating. Rational expectations can be assessed for stability under various types of learning, with least squares learning playing a prominent role. In addition to assessing the plausibility of an equilibrium, learning also provides a selection criterion when there are multiple equilibria. Monetary policy should be designed to avoid instability under learning and to facilitate coordination on desirable equilibria. Learning can also help to explain macroeconomic fluctuations as arising through either instabilities, stable indeterminacies or persistent learning dynamics.

Learning in macroeconomics refers to models of expectation formation in which agents revise their forecast rules over time, for example in response to new data. Expectations of future income, prices and sales play key roles in theories of saving and investment. Many other examples of the central role of expectations could be given.

1. Introduction

The current standard methodology for modelling expectations is to assume that the economy is in a rational expectations equilibrium (REE). REE is a model-consistent equilibrium in the two-way relationship between the influence of expectations on the economy and the dependence of expectations on the time path of the economy.

The standard formulation of REE makes strong assumptions on the information of economic agents. The true stochastic process of the economy is assumed known, with unforecastable random shocks constituting the remaining uncertainty. This assumption presupposes that the economic agents know much more than, say, the economists who in practice do not know the true stochastic structure and instead must estimate its parameters.

Recently, macroeconomic theory has been moving beyond the strict rational expectations (RE) hypothesis. Explicit models of imperfect knowledge and associated learning processes have been developed. In models of learning economic agents try to improve their knowledge of the stochastic process of the economy over time as new information becomes available.

Different approaches to modelling learning behaviour have been employed. Perhaps the most common has been ‘adaptive learning’, which views economic agents as econometricians who estimate the parameters of their model and make forecasts using their estimates. In adaptive learning economic agents have limited common knowledge since they estimate their own perceived laws of motion.

A second approach, called ‘eductive learning’, assumes common knowledge of rationality: economic agents engage in a process of reasoning about the possible outcomes knowing that other agents engage in the same process. Eductive learning takes place in logical time. A third approach has been ‘rational learning’, which employs a Bayesian viewpoint. Full knowledge of economic parameters is then replaced by priors and Bayesian updating under a correctly specified model, including common knowledge that all agents share this knowledge. Rational learning thus retains a form of REE at each point of time.

Basic theories of learning were developed largely in the 1980s and 1990s. See Sargent (1993, 1999), Evans and Honkapohja (2001), Guesnerie (2005) and Beck and Wieland (2002) for references. Recently, models of learning have been applied to issues of macroeconomic, and especially monetary, policy. In this overview, we focus on adaptive learning as it has been the most widely used approach. (For references to the pre-2001 literature, see Evans and Honkapohja, 2001.)

2. Least squares learning

In adaptive learning it is commonly assumed that agents estimate their model of the dynamics of economic variables, called the *perceived law of motion* (PLM), by *recursive least squares* (RLS), arguably the most common estimation method in econometrics.

2. 1 Overview

We illustrate the key concepts using the Cagan model of the price level $\hat{m} - p_t = -\psi(p^e_{t+1} - p_t) + \varphi' w_t + \varepsilon_t$, where p_t and $\hat{m}$ are logarithms of the price level and (constant) nominal money supply. Here $\psi > 0$ and p^e_{t+1} denotes the expectations of p_{t+1} formed at time *t*. w_t is a vector of observable exogenous variables, assumed to follow a stationary vector autoregression (VAR) process $w_t = Fw_{t-1} + e_t$, in which *F* is taken as known for simplicity. ε_t is an unobservable i.i.d. shock.

The reduced form of the Cagan model is

$$p_t = \alpha_0 + \alpha_1 p^e_{t+1} + \beta' w_t + v_t, \tag{1}$$

where $v_t = -(1 + \psi)^{-1}\varepsilon_t$ and α_0, α_1 and β depend on $\hat{m}$, ψ and φ. The model has a unique REE of the form $p_t = \bar{a} + \bar{b}' w_t + v_t$, where $\bar{a} = (1-\alpha_1)^{-1}\alpha_0$, $\bar{b} = (I - \alpha_1 F')^{-1}\beta$.

Agents are assumed to use the PLM $p_t = a + b'w_t + \eta_t$, where η_t is a disturbance term. The PLM has the same functional form as the REE but possibly different coefficients since agents do not know the REE. To estimate the PLM, agents use data $\{p_i, w_i\}_{i=0}^{t-1}$ and forecast using the estimated model $E^*_t p_{t+1} = a_{t-1} + b'_{t-1} F w_t$.

These forecasts lead to a temporary equilibrium or *actual law of motion* (ALM) $p_t = T(\phi_{t-1})' z_t + v_t$, where $T(\phi)' = (\alpha_0 + \alpha_1 a, \alpha_1 b'F + \beta')$. The REE $(\bar{a}, \bar{b}')$ is a fixed point of the mapping $T(\phi)$ from the PLM to the ALM. If we let $\phi'_t = (a_t, b'_t)$ and $z'_t = (1, w'_t)$, RLS estimation is given by

$$\begin{aligned} \phi_t &= \phi_{t-1} + t^{-1} R_t^{-1} z_t (p_t - \phi'_{t-1} z_t) \\ R_t &= R_{t-1} + t^{-1}(z_t z'_t - R_{t-1}). \end{aligned} \tag{2}$$

where p_t is given by the ALM. We say that the REE is *stable under RLS learning* if $(a_{t-1}, b'_{t-1}) \to (\bar{a}, \bar{b}')$ over time.

This model of learning involves bounded rationality. Each period agents maximize their objective, given their forecasts. However, agents treat the economy as having constant parameters, which is true only in the REE. Outside the REE the PLMs are misspecified, but misspecification vanishes as learning converges to the REE.

A key result, which holds in numerous models, is that RLS learning converges to RE under certain conditions on model parameters. Thus, the REE can be learned even though economic agents initially have limited knowledge and are boundedly rational.

Expectational stability (*E-stability*) is a convenient way for establishing the convergence conditions for RLS learning. Define the differential equation $d\phi/d\tau = T(\phi) - \phi$, which describes partial adjustment in virtual time τ. The REE is *E-stable* if it is locally stable under the differential equation. For models of the form (1), convergence is guaranteed if $0 < \alpha_1 < 1$, which is satisfied in the Cagan model since $\alpha_1 = \psi(1 + \psi)^{-1}$. Evans and Honkapohja (2001) contains a detailed discussion of convergence of RLS learning.

2.2 The roles of learning

Adaptive learning has several other important roles besides being a stability theory for REE. RE models can have multiple stationary equilibria, that is, *indeterminacy of equilibrium*. In such situations learning stability acts as a *selection criterion* to determine the plausibility of a particular REE.

As an example consider the non-stochastic Cagan model with government spending financed by seigniorage, with nonlinear reduced form $x_t = G(x^e_{t+1}),$ where x_t denotes inflation (see Evans and Honkapohja, 2001, chs. 11 and 12, for details). This model has two (interior) steady state solutions $\hat{x} = G(\hat{x})$. The low-inflation steady state x_L is stable under learning and the high-inflation steady state x_H is not. Learning selects a unique REE x_L in this model. In more general models, learning stability does not necessarily select a unique REE, but the set of 'plausible' REE is usually significantly smaller than the set of all REE.

The roles of RLS learning are not restricted to stability of REE and equilibrium selection. Learning can also provide new forms of dynamics as discussed below.

3. Monetary policy design

Indeterminacy of equilibria and instability of REE under RLS learning mean that the economy can be subject to persistent fluctuations. These instabilities can arise in the New Keynesian (NK) model (Woodford, 2003), which is widely used for studying monetary policy. Policy design has an important role in eliminating these instabilities and facilitating convergence to 'desirable' equilibria.

Consider the linearized NK model. The IS and PC curves $x_t = -\varphi(i_t - E^*_t \pi_{t+1}) + E^*_t x_{t+1} + g_t$ and $\pi_t = \lambda x_t + \beta E^*_t \pi_{t+1} + u_t$ summarize private sector behaviour. Here x_t, π_t and i_t denote the output gap, inflation and the nominal interest rate. φ and λ are positive parameters while $0 < \beta < 1$ is the discount factor. The shocks g_t and u_t are assumed to be observable and follow a known *VAR*(1) process.

Central bank (CB) behaviour is described by an interest-rate rule. CB may use an instrument rule that is not based on explicit optimization. Examples are Taylor rules that depend on current data or forecasts, $i_t = \chi_\pi \pi_t + \chi_x x_t$ or $i_t = \chi_\pi E^*_t \pi_{t+1} + \chi_x E^*_t x_{t+1}$, where χ_π, χ_x >0.

The IS and PC equations, together with either Taylor rule, lead to a bivariate reduced form in (x_t,π_t), which can be examined for determinacy (uniqueness of equilibrium) and E-stability. Bullard and Mitra (2002) show that current-data Taylor rules yield both E-stability and determinacy iff $\lambda(\chi_\pi - 1) + (1 - \beta)\chi_x > 0$. Under forward-looking rules $\chi_\pi > 1$ and small χ_x yield E-stability and determinacy.

Optimal monetary policy under discretion and commitment has been examined by Evans and Honkapohja (2003a; 2003b; 2006). Various ways to implement optimal policy have been suggested. Some commonly suggested interest-rate rules, based on fundamental shocks and variables, can lead to E-instability and/or indeterminacy. Evans and Honkapohja advocate appropriate expectations-based rules that deliver both E-stability and determinacy.

Other aspects of learning are also important for monetary policy. One practical concern is the observability of private forecasts needed for forecast-based rules. Results by Honkapohja and Mitra (2005) show that using internal CB forecasts in place of private sector expectations normally delivers E-stability.

Another difficulty for optimal monetary policy is that it requires knowledge of structural parameters, which are in practice unknown. CB can learn the values of φ and λ by estimating IS and PC equations. Expectations-based optimal rules continue to deliver stability under simultaneous learning by private agents and the CB (see Evans and Honkapohja, 2003a; 2003b).

4. Fluctuations

A major issue in macroeconomics is economic fluctuations, for example, business cycles and asset price movements. Can learning help to explain these phenomena?

4.1 Stable sunspot fluctuations

One theory of macroeconomic fluctuations interprets them as rational ‘sunspot’ equilibria. Although many macroeconomic models – for example, the real business cycle (RBC) model or Taylor’s overlapping contracts model – have a unique stationary solution under RE, other models can have indeterminacy. Examples include the overlapping generations (OLG) model and RBC models with increasing returns and monopolistic competition or tax distortions.

When multiple equilibria are present, some solutions may depend on variables, ‘sunspots’, that are completely extraneous to the economy. Such stationary sunspot equilibria (SSEs) exhibit self-fulfilling prophecies with the sunspot acting as a coordinating device: if expectations depend on a sunspot variable, then the actual economy, since it depends on expectations, can also depend rationally on the sunspot.

As already noted, learning stability is a selection device. Suppose agents’ forecasts are a linear function of both the macroeconomic state and a sunspot variable. If the forecast functions have coefficients close to but not equal to SSE values, and if agents update the estimated coefficients using RLS, can the coefficients converge to SSE values? If not, this casts doubt on the plausibility of SSEs.

SSEs appear not to be stable under learning in indeterminate RBC models but are learnable in some other models. We first describe results for the NK model and then discuss the possibility of stable SSE in other models.

4.1.1 SSEs in the NK model

Consider again the linearized NK model augmented by either the current-data or forward-looking Taylor rule. As noted above, indeterminacy is likely when the ‘Taylor principle’ $\chi_\pi > 1$ is violated.

In practice CBs are said to use forward-looking rules, and Clarida, Gali and Gertler (2000) argue that empirical estimates of χ_π are less than 1 in the period before 1984, while they are greater than 1 for the subsequent period. Could SSEs explain the higher economic volatility in the earlier period?

Honkapohja and Mitra (2004) and Evans and McGough (2005) approach this question by asking when SSEs are stable under learning in the NK model. Surprisingly, SSEs appear never to be stable under learning for current-data Taylor rules. When the forward-looking Taylor rule is employed, stable SSEs occur not when $\chi_\pi < 1$, but rather when $\chi_\pi > 1$ and χ_π and χ_x are sufficiently large, that is, *overly* aggressive rules lead to learnable SSEs. However, this does not rule out the Clarida, Gali, Gertler explanation for pre-1984 instability because, if $\chi_\pi < 1$ leads to indeterminacy, *no* REE is stable under learning and aggregate instability would presumably result.

4.1.2 Stable SSEs in other models

Stability under learning is a demanding test for SSEs that is met in only some cases in the NK model. There are, however, other examples of stable SSEs, such as the basic OLG model.

Some nonlinear models can have multiple steady states that are locally stable under RLS learning. In this case there can also be SSEs that take the form of occasional random shifts between neighbourhoods of the distinct stable steady states. Examples of this are the ‘animal spirits’ model of Howitt and McAfee (1992), based on a positive search externality, and the ‘growth cycles’ model of Evans, Honkapohja and Romer (1998) based on monopolistic competition and complementarities between capital goods.

Two stable steady states also play a role in some important policy models. This can arise in a monetary inflation model with a fiscal constraint, developed by Evans, Honkapohja and Marimon (2001), and in the liquidity trap model of Evans and Honkapohja (2005). In these set-ups policy has an important role in eliminating undesirable steady states.

4.2 Dynamics with constant gain learning

An alternative route to explaining economic fluctuations is to modify RLS learning so that more recent observations are given a higher weight. A natural way to motivate this is to assume that agents are concerned about the possibility of structural change. In the RLS formula (2) this can be formally accomplished by replacing t^{-1} with a small ‘constant gain’ $0 < \gamma < 1$, yielding weights that geometrically decline with the age of observations.

This apparently small change leads to ‘boundedly rational’ fluctuations, with sometimes dramatic effects. Three main phenomena have emerged. First, as shown by Sargent (1999) and Cho, Williams and Sargent (2002), even when there is a unique equilibrium, occasional ‘escape paths’ can arise with learning dynamics temporarily driving the economy far from the equilibrium. Sargent shows how the reduction of inflation in the 1982–99 period might be due to such an escape path in which policymakers are led to stop attempting to exploit a perceived (but misspecified) Phillips curve trade-off.

Second, in models with multiple steady states, learning dynamics can take the form of periodic shifts between regimes as a result of intrinsic random shocks interacting with learning dynamics. This is seen in the ‘increasing social returns’ example of Evans and Honkapohja (2001), the hyperinflation model of Marcet and Nicolini (2003), the exchange rate model of Kasa (2004) and the liquidity trap model of Evans and Honkapohja (2005).

Third, even when large escapes do not arise, there can be policy implications, because constant gain learning differs in small but persistent ways from full rationality. Orphanides and Williams (2005) show that policymakers attempting to implement optimal policy should be more hawkish against inflation than under RE.

5. Other developments

There continue to be many new applications of learning dynamics in macroeconomics, with closely related work in asset pricing and game theory.

One recent topic concerns the possibility that agents use a misspecified model. Under RLS learning agents may still converge, but to a restricted perceptions equilibrium, rather than to an REE (see Evans and Honkapohja, 2001). Another recent development is to allow agents to select from alternative predictors. In the Brock and Hommes (1997) model agents choose, based on recent past performance, between a costly sophisticated and a cheap naive predictor. This can lead to complex nonlinear dynamics. Branch and Evans (2006) combine dynamic predictor selection with RLS learning and show the existence of ‘misspecification equilibria’ when all forecasting models are underparameterized.

Other topics and applications include empirical work on expectation formation, calibration and estimation of learning models to data, interaction of policymaker and private-sector learning, learning and robust policy, experimental studies of expectation formation, the role of calculation costs, expectations over long horizons, alternative learning algorithms, expectational and structural heterogeneity, transitional learning dynamics, consistent expectations and near-rationality.

Current interest in learning dynamics is evidenced by five recent Special Issues devoted to learning and bounded rationality, in *Macroeconomic Dynamics* (2003), *Journal of Economic Dynamics and Control* (two in 2005), *Review of Economic Dynamics* (2005), and *Journal of Economic Theory* (2005).

George W. Evans and Seppo Honkapohja

Bibliography

Beck, G. and Wieland, V. 2002. Learning and control in a changing environment. *Journal of Economic Dynamics and Control* 26, 1359–77.

Branch, W. and Evans, G. 2006. Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory* 127, 264–95.

Brock, W. and Hommes, C. 1997. A rational route to randomness. *Econometrica* 65, 1059–95.

Bullard, J. and Mitra, K. 2002. Learning monetary policy rules. *Journal of Monetary Economics* 49, 1105–29.

Cho, I.-K., Williams, N. and Sargent, T. 2002. Escaping Nash inflation. *Review of Economic Studies* 69, 1–40.

Clarida, R., Gali, J. and. Gertler, M. 2000. Monetary policy rules and macroeconomic stability: evidence and some theory. *Quarterly Journal of Economics* 115, 147–80.

Evans, G. and Honkapohja, S. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.

Evans, G. and Honkapohja, S. 2003a. Expectations and the stability problem for optimal monetary policies. *Review of Economic Studies* 70, 807–24.

Evans, G. and Honkapohja, S. 2003b. Adaptive learning and monetary policy design. *Journal of Money Credit and Banking* 35, 1045–72.

Evans, G. and Honkapohja, S. 2005. Policy interaction, expectations and the liquidity trap. *Review of Economic Dynamics* 8, 303–23.

Evans, G. and Honkapohja, S. 2006. Monetary policy, expectations and commitment. *Scandinavian Journal of Economics* 108, 15–38.

Evans, G., Honkapohja, S. and Marimon, R. 2001. Convergence in monetary inflation models with heterogeneous learning rules. *Macroeconomic Dynamics* 5, 1–31.

Evans, G., Honkapohja, S. and Romer, P. 1998. Growth cycles. *American Economic Review* 88, 495–515.

Evans, G. and McGough, B. 2005. Monetary policy, indeterminacy and learning. *Journal of Economic Dynamics and Control* 29, 1809–40.

Guesnerie, R. 2005. *Assessing Rational Expectations 2: 'Eductive' Stability in Economics*. Cambridge, MA: MIT Press.

Honkapohja, S. and Mitra, K. 2004. Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics* 51, 1743–70.

Honkapohja, S. and Mitra, K. 2005. Performance of monetary policy with internal central bank forecasting. *Journal of Economic Dynamics and Control* 29, 627–58.

Howitt, P. and McAfee, R. 1992. Animal spirits. *American Economic Review* 82, 493–507.

Kasa, K. 2004. Learning, large deviations and recurrent currency crises. *International Economic Review* 45, 141–173.

Marcet, A. and Nicolini, J. 2003. Recurrent hyperinflations and learning. *American Economic Review* 93, 1476–98.

Orphanides, A. and Williams, J. 2005. Imperfect knowledge, inflation expectations and monetary policy. In *The Inflation-Targeting Debate*, ed. B. Bernanke and M. Woodford. Chicago: University of Chicago Press.

Sargent, T. 1993. *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.

Sargent, T. 1999. *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Woodford, M. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.