Learning and Macroeconomics

George W. Evans[1*] and Seppo Honkapohja[2]

[1]Department of Economics, University of Oregon, Eugene, Oregon 97403-1285, and School of Economics and Finance, University of St. Andrews, KY16 9AL Scotland; email: gevans@uoregon.edu

[2]Bank of Finland, PO Box 160, FI00101 Helsinki, Finland; email: seppo.honkapohja@bof.fi

Annu. Rev. Econ. 2009. 1:421–49

First published online as a Review in Advance on June 8, 2009

The *Annual Review of Economics* is online at econ.annualreviews.org

This article's doi:
10.1146/annurev.economics.050708.142927

1941-1383/09/0904-0421$20.00

*Corresponding author.

Key Words

E-stability, stochastic approximation, persistent learning dynamics, business cycles, monetary policy, asset prices

Abstract

Expectations play a central role in modern macroeconomic theories. The econometric learning approach models economic agents as forming expectations by estimating and updating forecasting models in real time. The learning approach provides a stability test for rational expectations and a selection criterion in models with multiple equilibria. In addition, learning provides new dynamics if older data are discounted, if models are misspecified, or if agents choose between competing models. This paper describes the expectational stability (E-stability) principle and the stochastic approximation tools used to assess equilibria under learning. Applications of learning to a number of areas are reviewed, including the design of monetary and fiscal policy, business cycles, self-fulfilling prophecies, hyperinflation, liquidity traps, and asset prices.

MN argue that under RE the model cannot properly explain the main stylized facts of hyperinflation, and that a learning formulation is more successful. They use a variation of learning rule (Equation 20) in which t is replaced by α_t, where $\alpha_t = \alpha_{t-1} + 1$ if $|\left(\frac{P_{t-1}}{P_{t-2}} - \beta_{t-1}\right)/\beta_{t-1}|$ falls below some bound, and otherwise $\alpha_t = \bar{\alpha}$; that is, a constant gain is used. The qualitative features of the model are approximated by the system $\frac{P_t}{P_{t-1}} = h(\beta_{t-1}, d_t)$, where

$$h(\beta, d) = \begin{cases} T(\beta; d) & \text{if } 0 < T(\beta; d) < \beta_U \\ \bar{\beta} & \text{otherwise.} \end{cases}$$

Figure 3 describes the dynamics of system. There is a stable region, consisting of values of β below the unstable high-inflation steady state β_H, and an unstable region that lies above it. This gives rise to very natural recurring hyperinflation dynamics: Starting from β_L, occasionally a sequence of random shocks may push β_t into the unstable region, at which point the gain is revised upward to $1/\bar{\alpha}$, and inflation follows an explosive path until it is stabilized by ERR. Then the process begins again.

The model with learning has useful policy implications. ERR is valuable as a way of ending hyperinflation if the economy enters the explosive regime. However, a higher $E(d_t)$ makes average inflation higher and the frequency of hyperinflations greater. This indicates the importance of the orthodox policy of reducing deficits as a way of minimizing the likelihood of hyperinflation paths.

3.3.2. Liquidity traps and deflationary spirals. Deflation and liquidity traps have at times been a concern. As we have seen, in contemporaneous Taylor rules, interest rates should respond to the inflation rate more than one-for-one in order to ensure determinacy and stability under learning near the target inflation rate. However, as emphasized by Benhabib et al. (2001), if one considers the interest-rate rule globally, the requirement that net nominal interest rates must be nonnegative implies that the rule must be nonlinear and also, for any continuous rule, the existence of a second steady state at a lower (possibly negative) inflation rate. This is illustrated in **Figure 4**. The Fisher equation $R = \pi/\beta$, depicted in the figure, is obtained from the usual Euler equation for consumption in a steady state. Here R stands for the interest rate factor (the net interest rate is $R - 1$), $\pi_t = P_t/P_{t-1}$ stands for the inflation factor ($\pi - 1$ is the net inflation rate), π^* denotes the intended steady state, π_L is the unintended steady state, and π_L may correspond to either a very low positive or a negative net inflation rate, i.e., deflation. Benhabib et al. (2001) show that under RE, there is a continuum of "liquidity trap" paths that converge on π_L. The pure RE analysis thus suggests a serious risk of the economy following these liquidity trap paths.

What happens under learning? Evans & Honkapohja (2005) analyzed a flexible-price perfect competition model. We showed that deflationary paths are possible, but that the real risks, under learning, were paths in which inflation slipped below π_L and then continued to fall further. For this flexible-price model, we showed that this could be avoided by a switch to an aggressive money supply rule at low inflation rates.

Evans et al. (2008) reconsider the issues in an NK model with sticky prices, due to adjustment costs, and deviations of output from flexible-price levels. Monetary policy follows a global Taylor rule, as above. Fiscal policy is standard, including exogenous government purchases g_t and a Ricardian tax policy that depends on real debt level. The

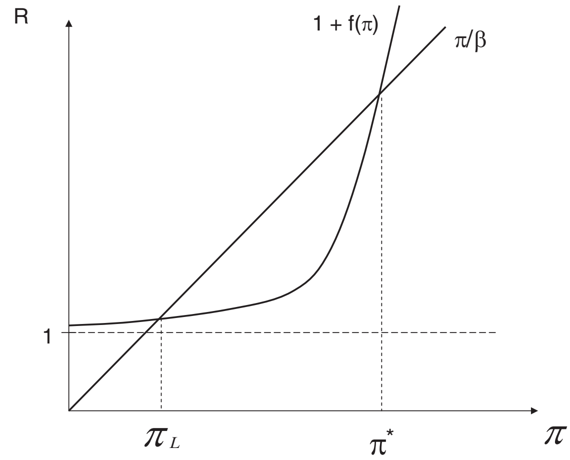

Figure 4

Multiple steady states with a global Taylor rule. Shown is the interest-rate policy $R = 1 + f(\pi)$ as a function of π (a dependence on aggregate output is omitted for simplicity). The straight line represents the Fisher equation $R = \pi/\beta$. R stands for the interest rate factor and $\pi_t = P_t/P_{t-1}$ for the inflation factor. π^* denotes the intended steady state, at which the Taylor principle of a more than one-for-one response is satisfied, and π_L is the unintended steady state.

model equations are nonlinear, and the nonlinearity in its analysis under learning is retained. The key equations are

$$\frac{\alpha\gamma}{v}(\pi_t - 1)\pi_t = \beta\frac{\alpha\gamma}{v}(\pi^e_{t+1} - 1)\pi^e_{t+1} + (c_t + g_t)^{(1+\varepsilon)/\alpha} - \alpha\left(1 - \frac{1}{v}\right)(c_t + g_t)c_t^{-\sigma_1} \text{ and}$$
$$c_t = c^e_{t+1}(\pi^e_{t+1}/\beta R_t)^{\sigma_1}.$$

The first equation is the nonlinear NK Phillips curve, and the second equation is the IS curve. There are also money- and debt-evolution equations.

There are two stochastic steady states at π_L and π_H. If the random shocks are i.i.d., then steady-state learning is appropriate for both c^e and π^e, specifically

$$\pi^e_{t+1} = \pi^e_t + \phi_t(\pi_{t-1} - \pi^e_t) \text{ and}$$
$$c^e_{t+1} = c^e_t + \phi_t(c_{t-1} - c^e_t),$$

where ϕ_t is the gain sequence. The intended steady state π^* is locally stable under learning, whereas the unintended steady state π_L is unstable. The key observation is that π_L is a saddle point, which implies the existence of deflationary spirals under learning. In particular, after a sufficiently pessimistic expectational shock, c^e and π^e will follow paths leading to deflation and stagnation. This is illustrated in **Figure 5** (see color insert), which gives the E-stability dynamics.

For the intuition, suppose that we are initially near the π_L steady state and that we consider a small drop in π^e. With fixed R this would lead through the IS curve to lower c and thus, through the Phillips curve, to lower π. Because only small reductions in R are possible given the global Taylor rule, the reduction in c and π cannot be offset. The falls in realized c and π lead, under learning, to reductions in c^e and π^e, and this sets in motion the deflationary spiral.

Thus, large adverse shocks to expectations or structural changes can set in motion unstable downward paths. Can policy be altered to avoid deflationary spiral? Evans et al. (2008) show that it can. The recommended policy is to set a minimum inflation threshold $\tilde{\pi}$, where $\pi_L < \tilde{\pi} < \pi^*$. The authorities would follow normal monetary and fiscal policy, provided that it delivers $\pi_t > \tilde{\pi}$. However, if π_t threatens to fall below $\tilde{\pi}$, then aggressive policies would be implemented to ensure that $\pi_t = \tilde{\pi}$: Interest rates would be reduced, if necessary to near the zero lower bound $R = 1$, and if this is not sufficient, then government purchases g_t would be increased as required. It can be shown that these policies can indeed ensure $\pi_t \geq \tilde{\pi}$ always under learning and that they can lead to global stability of the intended steady state at π^*. Perhaps surprisingly, it is essential to have an inflation threshold. Use of an output threshold to trigger aggressive polices will not always avoid deflationary spirals.

3.4. Asset Prices

Asset pricing is another area of recent focus in the learning literature. The potential for adaptive learning to generate new phenomena for asset prices was already apparent in the early work of Timmermann (1993, 1996). Consider the standard risk-neutral asset-pricing framework

$$p_t = \beta E_t^*(p_{t+1} + d_{t+1}), \tag{21}$$

where p_t is the real price of equalities, d_{t+1} is the real dividend paid at the end of period $t+1$, and $0 < \beta \equiv (1+r)^{-1} < 1$ is the discount factor, assumed constant. Assume also that d_t is an exogenous stochastic process, such as

$$\ln(d_t) = \mu + \ln(d_{t-1}) + \varepsilon_t, \tag{22}$$

where ε_t is i.i.d. $N(0, \sigma^2)$ and $\mu + \sigma^2/2 < \ln(1+r)$. Under RE the fundamentals solution is

$$p_t = \frac{1+g}{r-g} d_t, \text{ where } 1 + g = \exp(\mu + \sigma^2/2). \tag{23}$$

Iterating Equation 21 forward, one obtains the present value formula

$$p_t = \sum\nolimits_{j=1}^{\infty} \beta^j E_t^* d_{t+j}, \tag{24}$$

and imposing RE yields Equation 23.

There are a number of empirical puzzles in asset pricing based on this model, including excess volatility of stock prices and predictability of stock returns. A potentially simple and appealing explanation is that traders do not have a priori knowledge of the parameters of the (p_t, d_t) process. The equilibrium price process under learning is generated, as usual, assuming that the parameter estimates are updated over time as new data become available.

There are two natural ways to model stock prices under learning, depending on whether we want to treat Equation 24 or 21 as the key equation that determines p_t, given expectations. If traders are "fundamentalists," then price will be set in accordance with Equation 24, based on forecasts $E_t^* d_{t+j} = d_t \exp(j\hat{\mu}_t + j\hat{\sigma}_t^2/2)$, where $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ are the time t estimates of μ and σ^2, respectively. Such price setting leads to

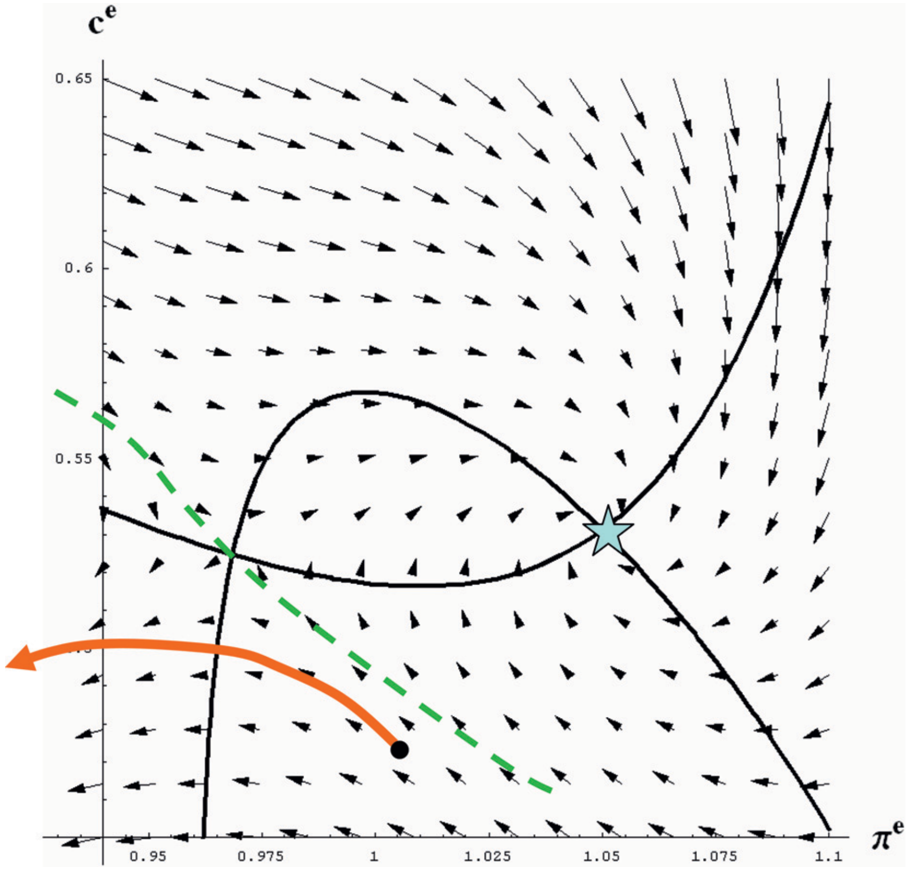

Figure 5

Expectation dynamics under normal policy. The star indicates the intended steady state, which is locally stable under learning. The dashed curve gives the boundary of the stable region, and the arrow shows an unstable path leading to deflation and stagnation.