Chapter 11

Learning

11.1 Motivation

In all of the lectures so far we have assumed that agents know everything it is possible to know about the structure of the economy. Even in the models of monetary policy under uncertainty, agents knew the final distribution of the uncertain parameters. In this lecture we relax this assumption and look at models where agents are able to learn the key parameters of the model. We begin by examining how the nature of monetary policy is affected by learning considerations and then continue in the second part of the lecture to examine how the equilibrium which emerges in the economy can be different under learning.

11.2 Key readings

There is no required reading for today's lecture. The examples we will discuss are simplified versions of the following two papers.

Ellison and Valla (2001) "Learning, uncertainty and central bank activism in an economy with strategic interactions", *Journal of Monetary Economics*, 48, 153-171

Marcet and Nicolini (1998) "Recurrent hyperinflations and learning", *CEPR Discussion Paper*, No. 1875

11.3 Related reading

Balvers and Cosimano (1994) "Inflation variability and gradualist monetary policy", *Review of Economic Studies*, 61, 721-738

Bertocchi and Spagat (1993) "Learning, experimentation and monetary policy", *Journal of Monetary Economics*, 32, 169-183

Evans and Honkapohja (2001) *Learning and Expectations in Macroeconomics*, Princeton University Press.

Marimon and Sunder (1993) "Indeterminacy of equilibria in a hyperinflationary world: experimental evidence", Econometrica, 61, 1073-1108

Wieland (2000) "Learning by doing and the value of optimal experimentation", *Journal of Economic Dynamics and Control*, 24, 501-534

11.4 Learning and monetary policy

In the previous lecture we saw that the presence of uncertainty lead to a cautious policy by the central bank. This call for Brainard conservatism has been challenged in a number of papers, originally by Bertocchi and Spagat (1993) and Balvers and Cosimano (1994). They argue that a cautious policy is suboptimal because it is very poor from a learning point of view. If the central bank is cautious in its use of policy then it will be very difficult to learn what the effects of monetary policy are. In terms of learning, the central bank should be more aggressive in policy since that way it learns the key parameters about how the policy works. To see how this argument works in practice we will look at a stylised example, which is a simplification of Ellison and Valla (2001).

The structure of the economy is characterised by the Phillips curve (11.1), where output y_t depends on inflation π_t and two supply shocks ε_{1t} and ε_{2t}. Only the first supply shock is observable. Both shocks have the same variance σ_ε.

$$y_t = \beta\pi_t + \varepsilon_{1t} + \varepsilon_{2t} \tag{11.1}$$

The Phillips curve parameter is not known with certainty by the central bank. For simplicity, we will assume it can take one of two possible values, β_H or β_L. The central bank believes that it is high with probability p_t and low with probability $1-p_t$. The loss function of the central bank (11.2) is assumed to punish the expected present discounted value of deviations of output from natural rate and inflation from zero.

$$\mathcal{L} = E_t \sum_{i=0}^{\infty} \delta^n \left[(y_{t+i} - y^*)^2 + \chi\pi_{t+i}^2 \right] \tag{11.2}$$

δ is the subjective discount factor and χ is the weight placed upon inflation as against output deviations from target. We begin by solving the model for the case where the central bank does not internalise its learning. In this passive learning case the maximisation problem of the central bank reduces to the static problem (11.3). There are no intertemporal linkages in the model and the central bank just minimises losses period by period.

$$\min_{\pi_t} E_t \left\{ (\beta\pi_t + \varepsilon_{1t} + \varepsilon_{2t} - y^*)^2 + \chi\pi_t^2 \right\} \tag{11.3}$$

To calculate the expected loss the central bank uses its beliefs about the parameter β so the problem becomes (11.4).

$$\min_{\pi_t} \left\{ p_t E_t(\beta_H \pi_t + \varepsilon_{1t} + \varepsilon_{2t} - y^*)^2 + (1 - p_t) E_t(\beta_L \pi_t + \varepsilon_{1t} + \varepsilon_{2t} - y^*)^2 + \chi \pi_t^2 \right\} \tag{11.4}$$

The associated optimal policy is (11.5).

$$\pi_t = -\frac{p_t \beta_H^2 + (1 - p_t)\beta_L^2}{p_t \beta_H^2 + (1 - p_t)\beta_L^2 + \chi} \varepsilon_{1t} \tag{11.5}$$

This policy is equivalent to the Brainard policy derived in the previous lecture. The central bank internalises the uncertainty about β in its policy and adjusts its reaction to the observed supply shock accordingly. Figure 11.1 shows this policy for a numerical example based on Ellison and Valla (2001). The upper dashed line shows the inflation choice of the central bank after a one standard deviation negative ε_{1t} supply shock, as a function of beliefs. Whatever the central bank's beliefs about the parameter β it always increases inflation to reflate the economy after a negative supply shock. The lower dashed line shows the corresponding inflation choice for a positive ε_{1t} supply shock.

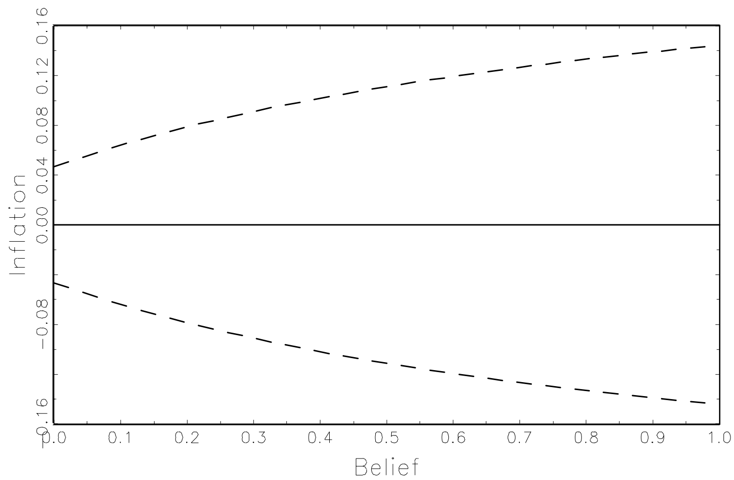

Figure 11.1: Passive learning policy

We now calculate the policy which explicitly takes learning into account. The central bank has several pieces of information which it can use to improve its inference about the actual value of the parameter β. At the end of the period it knows the observable supply shock ε_{1t}, its own inflation choice π_t and the outcome in terms of output y_t. The question is whether y_t has most likely been generated by a model with β_H or β_L. The answer to this lies in comparing the likelihood of outcome y_t in the β_H case (where the distribution of y_t is given by equation (11.6)) to the likelihood in the β_L case (where the distribution of y_t is given by equation (11.7)).

$$y_t \,|\beta_H \sim N\left[\beta_H \pi_t + \varepsilon_{1t}; \sigma_\varepsilon\right] \tag{11.6}$$

$$y_t \,|\beta_L \sim N\left[\beta_L \pi_t + \varepsilon_{1t}; \sigma_\varepsilon\right] \tag{11.7}$$

Comparing the distributions (11.6) and (11.7) to the actual outcome gives information on whether β really is high or low. To see how this feeds into beliefs we use Bayes rule. For this reason the learning in this section is often referred to as Bayesian learning. Equation (11.8) shows how Bayes rule can be applied to our problem.

$$\begin{aligned} P(\beta_H \,|y_t) &= \frac{P(\beta_H \cap y_t)}{P(y_t)} = \frac{P(\beta_H)P(y_t \,|\beta_H)}{P(y_t)} \\ p_{t+1} &= \frac{p_t P(y_t \,|\beta_H)}{p_t P(y_t \,|\beta_H) + (1-p_t)P(y_t \,|\beta_L)} \\ p_{t+1} &= \mathcal{B}(p_t, \pi_t, \varepsilon_{1t}, y_t) \end{aligned} \tag{11.8}$$

$\mathcal{B}$ is the Bayesian operator showing how beliefs p_t are updated on the basis of all relevant information. If the central bank internalises the learning mechanism (11.8) then it is said to be active learning. The problem of the active learning central bank is shown in (11.9).

$$\begin{aligned} \underset{\pi_t}{\mathcal{L}} &= E_t \sum_{i=0}^{\infty} \delta^n \left[(y_{t+i} - y^*)^2 + \chi \pi_{t+i}^2\right] \\ & \quad \text{s.t.} \\ y_t &= \beta \pi_t + \varepsilon_{1t} + \varepsilon_{2t} \\ p_{t+1} &= \mathcal{B}(p_t, \pi_t, \varepsilon_{1t}, y_t) \\ p_t &= p_0 \end{aligned} \tag{11.9}$$

This is an example of a non-linear-quadratic control problem and so will not have an optimal solution that is linear. For such a problem it is not possible to obtain closed-form analytical solutions but it is possible to use numerical techniques to derive an approximation to the optimal policy. We follow the approach of Wieland (2000) and show our numerical results in Figure 11.2.

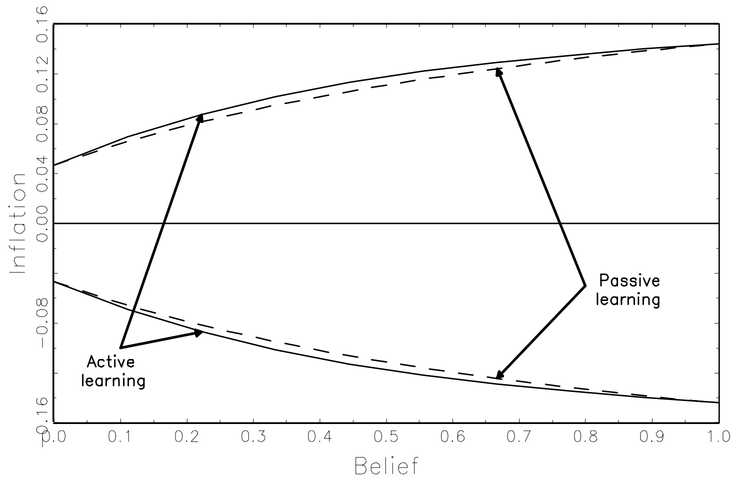

Figure 11.2: Optimal and passive learning policies

Compared with the passive learning policy, the reaction under the active learning policy is stronger, being further away from zero. The active learning policy is more aggressive than the passive learning policy, precisely because this has informational gains which help learning. This result has been used to suggest that central banks should engage in probing or experimentation - adjusting monetary policy to find out how the economy works.

11.5 Learning as an equilibrium selection device

The paper by Marcet and Nicolini (1998) shows how learning has important implications for the equilibrium outcome in the economy. They begin with a simple overlapping generations model in which each generation lives for two periods. When the agent is young (period 0) they receive an endowment of real value w_t^0, save M_{t+1}^d in the form of cash for old age and therefore consume $C_t^0 = w_t^0 - M_{t+1}^d/p_t$. When the agent is old (period 1) they receive endowment w_t^1 and use their money savings to give consumption $C_t^1 = w_t^1 + M_{t+1}^d/p_{t+1}$. With a logarithmic utility function, the decision problem faced by the agent when young is shown by (11.10).

$$\begin{aligned}
&\max E\left[\log(C_t^0) + \log(C_t^1)\right] \\
&\text{s.t.} \\
&C_t^0 = w_t^0 - M_{t+1}^d/p_t \\
&C_t^1 = w_t^1 + M_{t+1}^d/p_{t+1}
\end{aligned} \tag{11.10}$$

After substituting in for C_t^0 and C_t^1 and taking the first order condition we can derive the demand for money function (11.11), where $a = w_t^0/2$ and $b = w_t^1/2$.

$$\frac{M_{t+1}^d}{p_t} = a - b\pi_{t+1}^e \tag{11.11}$$

The amount of money demanded by the young in this overlapping generations model depends on the pattern of endowments and the expectation of future inflation. In the remainder of this section the focus will be on how the expectation of future inflation is formed - whether everything is already known or does the agent learn things useful for calculating the expectation. We complete the model by specifying a process for the money supply. For simplicity it is assumed that the real government deficit d is completely financed by changes in the money supply so the government budget constraint is given by equation (11.12).

$$M_{t+1}^s - M_t^s = d \tag{11.12}$$

In equilibrium, money demand (11.11) equals money supply (11.12) so we can write the determination of inflation as equation (11.13).

$$\pi_t = \frac{a - b\pi_t^e}{a - b\pi_{t+1}^e - d} \tag{11.13}$$

This is the key equation of the model. It shows how inflation is determined by expectations of current and future inflation. We now assume perfect-foresight rational expectations so $\pi_t^e = \pi_t \forall t$, in which case equation (11.13) becomes an expression showing how inflation evolves. In other words, we see how π_{t+1} depends on π_t as shown in equation (11.14).

$$\pi_{t+1} = R(\pi_t) = \frac{1}{b}\left[a - d + b - \frac{a}{\pi_t}\right] \tag{11.14}$$

There are two things to note about equation (11.14). Firstly, $R'(\pi_t) > 0$ and $R''(\pi_t) < 0$, implying the function is concave. Secondly, the equation has two stationary fixed points at which $\pi^* = R(\pi^*)$. These two properties imply that the mapping $R(\cdot)$ from π_t to π_{t+1} can be described as in Figure 11.3.

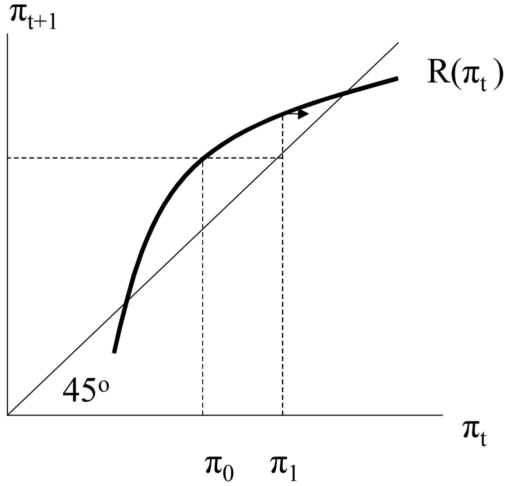

Figure 11.3: Dynamics of inflation under rational expectations

Any point on the $R(\cdot)$ mapping is a rational expectations equilibrium but only two are stationary. To see the dynamic behaviour of this model we can look at the path followed by inflation over time. Suppose that the economy starts out at π_0. Next period's inflation will be given by π_1 and so on. In Figure 11.3 all roads lead to Rome - whatever the level of inflation above the lower equilibrium you start from you will always converge to the high inflation stationary equilibrium.

Convergence to high inflation equilibria is a common feature of rational expectations models. We can say that rational expectations selects the high inflation equilibrium from amongst the set of multiple equilibria available - in this sense it is an equilibrium selection mechanism. However, there are good reasons to believe that in reality we do not see convergence to the high inflation equilibrium. Experimental evidence from Marimon and Sunder (1993) suggests that the economy almost always converges on the low inflation equilibrium. To reconcile these doubts we now turn to learning as an equilibrium selection device. We retain the demand for money function (11.11) and the definition of the money supply (11.12) but relax the assumption of rational expectations and instead propose that agents use an adaptive learning rule (11.15) to form their expectations.

$$\pi^e_{t+1} = \pi^e_t + \alpha(\pi_{t-1} - \pi^e_t) \tag{11.15}$$

Under this rule, agents update their expectations of inflation on the basis of how much inflation in the previous period deviated from that currently expected. If we assume α is small, which will be the case if learning has already proceeded for a reasonable amount of time, $\pi^e_{t-1} \approx \pi^e_t$ in equation (11.13) and we can substitute out for π_{t-1} to arrive at equation (11.16).

$$\pi^e_{t+1} = (1-\alpha)\pi^e_t + \alpha\left[\frac{a - b\pi^e_t}{a - b\pi^e_t - d}\right] \tag{11.16}$$

We now have an equation for the evolution of inflation expectations $\pi^e_{t+1} = G(\pi^e_t)$. Under perfect foresight, we can write this as $\pi_{t+1} = G(\pi_t)$ and note two things about the $G(\cdot)$ mapping (11.17). Firstly, $G'(\pi_t) > 0$ and $G''(\pi_t) > 0$ so the function is convex. Secondly, it has two fixed points $\pi^* = G(\pi^*)$ which coincide with the fixed points of the rational expectations solution. The function is shown in Figure 11.4.

$$\pi_{t+1} = G(\pi_t) = \pi_t + \frac{\alpha d}{a - b\pi_t - d} \tag{11.17}$$

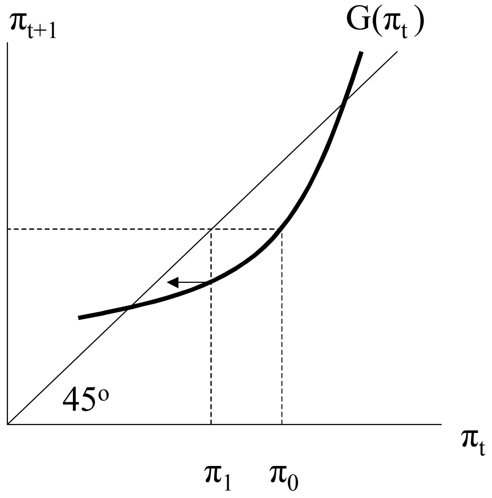

Figure 11.4: Dynamics of inflation under learning

The dynamics of inflation are now quite different from before. Under learning, inflation converges to its low stationary equilibrium instead of the high inflation equilibrium as in rational expectations. The robustness of these type of results has been analysed in many studies. Evans and Honkapohja (2001) is the state of the art textbook on this type of adaptive learning.