

Convergence to Rational Expectations in Learning Models: A Note of Caution

YiLi Chien, In-Koo Cho, and B. Ravikumar

We show in a simple monetary model that the learning dynamics do not converge to the rational expectations monetary steady state. We then show it is necessary to restrict the learning rule to obtain convergence. We derive an upper bound on the gain parameter in the learning rule, based on economic fundamentals in the monetary model, such that gain parameters above the upper bound would imply that the learning dynamics would diverge from the rational expectations monetary steady state. (JEL C60, D84)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2021, 103(3), pp. 351-65.
<https://doi.org/10.20955/r.103.351-65>

1 INTRODUCTION

Macroeconomic variables, such as inflation and interest rates, have been important objects of investigation in economics. A common assumption in economic models is that agents (households, firms, and governments) have rational expectations (RE) about these variables and thus accurately forecast the dynamics of the variables. An alternative approach limits the forecasting ability of the agents: They *learn* from the history of the variables. The central question of a learning model is whether the agents learn to behave over time in such a manner that the economy converges to the RE equilibrium in the long run. Thus, the convergence property of learning dynamics has been a key issue. Early examples that examine convergence to a RE equilibrium include Bray (1982), Bray and Savin (1986), Lucas (1986), Marcet and Sargent (1989a), and Woodford (1990).

We revisit the convergence issue using a simple two-period overlapping generations model of inflation—Example 1 in Bullard (1994). Ours is an endowment economy where money is the only store of value and monetary policy follows a constant money growth rule.

YiLi Chien is a research officer and economist at the Federal Reserve Bank of St. Louis. In-Koo Cho is a professor of economics at Emory University and Hayang University and a research fellow at the Federal Reserve Bank of St. Louis. B. Ravikumar is a senior vice president, the deputy director of research, and an economist at the Federal Reserve Bank of St. Louis. The authors are grateful for helpful conversations with James Bullard, George Evans, Thomas J. Sargent, and participants of “Learning Week” at the Federal Reserve Bank of St. Louis in 2017. Financial support from the National Science Foundation is gratefully acknowledged.

© 2021, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

When we do not impose RE, agents forecast inflation (or, equivalently, the rate of return on money) using a learning rule that is a convex combination of past expected inflation and actual inflation (a constant-gain algorithm), so the data on inflation affects learning. Based on the forecast, they choose consumption and real balances. These choices, in turn, affect the path of prices and generate feedback from learning to actual inflation. Other examples of such feedback include Bray (1982) on prices, Marcet and Sargent (1989b) on hyperinflation, and Evans and Honkapohja (1995) on business cycles.

Our results are as follows. First, we show numerically that, for some parameter configurations and initial conditions, the dynamic system produces cycles, nonmonotonic convergence, and a nonzero forecast error in the limit; that is, agents never learn the actual inflation. These simulation results show that the learning dynamics do not necessarily converge to the RE monetary steady state. (Besides the economic fundamentals such as endowments, preferences, and money growth, our simulations require initial conditions and one parameter of the learning rule—the constant gain or the weight that the learning rule places on the difference between actual and expected inflations, i.e., the forecast error.)

Second, we show that the gain parameter affects the convergence properties of the learning rule. We demonstrate that the convergence region around the monetary steady state varies with the gain. So, for some values of the gain parameter, the dynamic system monotonically converges to the steady state, while for other values it displays cycles. We derive a necessary condition that the gain has to satisfy in order to guarantee convergence to the monetary steady state. This condition depends on economic fundamentals. Alternative learning rules with a different information set break the dependence of the convergence property on the gain parameter.

Section 2 describes the learning model of Bray (1982). Section 3 sets up the learning model of Bullard (1994) and derives the difference equation that drives the learning dynamics. Section 4 provides simulation results that show nonconvergence and illustrate the dependence of convergence on the gain parameter in the learning rule. Section 5 contains concluding remarks.

2 A LEARNING MODEL OF PRICES

In this section, we present a version of the learning model of Bray (1992). This model has two equations. The first is a forecasting equation that describes the learning rule of an agent, and the second is an equation that describes the actual law of motion of the economic variable of interest, price p_t . The main purpose here is to illustrate the convergence property of the learning rule.

The agent forecasts the next period's price as the sample average of past prices p_t, p_{t-1}, \dots, p_1 , so the data affects the learning rule. Let p_{t+1}^e denote the time- t forecast of next period's price:

$$(1) \quad \begin{aligned} p_{t+1}^e &= \frac{1}{t} [p_t + p_{t-1} + \dots + p_1] \\ &= \frac{1}{t} p_t + \frac{t-1}{t} p_t^e. \end{aligned}$$

The evolution of p_t^e , or the forecasting equation, could then be expressed as

$$(2) \quad p_{t+1}^e = p_t^e + \frac{1}{t}(p_t - p_t^e).$$

In (2), the gain function—the weight placed by the learning rule on the forecast error—is $\frac{1}{t}$. The learning rule is thus a decreasing-gain algorithm.

The law of motion of actual price p_t is assumed to be

$$(3) \quad p_t = A - Bp_t^e,$$

where $A, B > 0$ are parameters of the model. In (3), the realized price is influenced by the agent's forecast. Thus, the learning model has two-way feedback.

Under RE, $p_t = p_t^e$. Equation (3) then implies $p_t = A/(B + 1)$; that is, the price remains constant.

The learning equilibrium, however, is a solution to the system (2) and (3). Substituting for p_t from (3) into (2), we get the learning dynamics

$$(4) \quad p_{t+1}^e = p_t^e + \frac{1}{t}(A - (B+1)p_t^e).$$

It is easy to see from (4) that as $t \rightarrow \infty$, the price forecast converges to a constant. In the long run, the *change* in forecasts is zero and p_t^e converges to $A/(B + 1)$, the RE equilibrium.

Some remarks are in order at this stage. Most learning models in macroeconomics create forecasts using least-squares estimation, which is a generalized form of the sample average of past observations; see, for instance, Evans and Honkapohja (2001) for examples of least-squares learning. By the strong law of large numbers we know that the sample average in the forecasting equation (1) eventually converges to the population mean. However, as we obtain more data over time, the weight $\frac{1}{t}$ on the forecast error decreases and the speed of convergence to the population mean decreases. One could speed up the convergence with a discounted average:

$$(5) \quad p_{t+1}^e = p_t^e + \gamma(p_t - p_t^e),$$

where $\gamma \in (0,1)$. The learning rule (5) is a constant-gain algorithm. This algorithm ensures that the speed of convergence remains the same over time, but the downside is the accuracy: Instead of the strong law of large numbers, that is, convergence in probability, we obtain weak convergence of the forecast to the population mean as $t \rightarrow \infty$. Note, however, that both algorithms imply convergence to the population mean. The learning literature, for the most part, treats the gain as a free parameter that the modeler can choose, without altering the convergence properties of the algorithm. See, for instance, Sargent (1999). In the next section, we use a constant-gain learning algorithm to illustrate that the gain parameter does indeed affect the limit properties of the learning rule.

3 A LEARNING MODEL OF INFLATION

The setup below is a constant-gain version of Example 1 in Bullard (1994). Consider an overlapping generations endowment economy where each generation lives two periods (young and old). We denote the generation born in period t as generation t . Each agent in generation $t = 0, 1, 2, \dots$ has a logarithmic utility function with no discounting:

$$U_t = \ln c_{1,t} + \ln c_{2,t},$$

where $c_{i,t}$ is consumption of the generation- t agent in $i = 1, 2$ period of the agent's life. Each generation- t agent is endowed with 2 and 2λ , $\lambda \in (0, 1)$, units of perishable consumption goods when young and old, respectively. The population size of each generation is normalized to 1.

Fiat money is the only store of value. The government finances its expenses by issuing fiat money, which affects the price level every period and therefore the inflation rate. Monetary policy is described by an exogenous constant growth rate of money:

$$(6) \quad M_t = \theta M_{t-1}, \theta \in (1, \lambda^{-1}).$$

The timing is as follows. In each period, the old agents enter with the nominal balances from the previous period. The young agents make their consumption and saving decisions. The government purchases goods by injecting money. Finally, consumption takes place based on realized prices at the end of the period.

Our focus here is on the monetary steady state. Under RE, we show that there is a unique monetary steady state. We then examine a learning model's convergence properties.

3.1 Rational Expectations

Given the deterministic setup, agents of each generation know the entire sequence of prices under RE. Given the prices in t and $t+1$, the lifetime budget constraint of a generation- t agent is

$$(7) \quad c_{1,t} + \frac{P_{t+1}}{P_t} c_{2,t} \leq 2 + 2\lambda \frac{P_{t+1}}{P_t}.$$

The problem of generation t is

$$\max_{\{c_{1,t}, c_{2,t}\}} \ln c_{1,t} + \ln c_{2,t}$$

subject to (7). Combining the first-order conditions with respect to $c_{1,t}$ and $c_{2,t}$, we get

$\frac{c_{1,t}}{c_{2,t}} = \frac{P_{t+1}}{P_t}$. Together with (7), the first-order conditions imply the optimal (interior) choices are

$$c_{1,t} = 1 + \lambda \frac{P_{t+1}}{P_t} \text{ and } c_{2,t} = \frac{P_t}{P_{t+1}} + \lambda.$$

Therefore, the saving of generation t is $2 - c_{1,t} = 1 - \lambda \frac{P_{t+1}}{P_t}$. Since fiat money is the only store of value, the saving must be in the form of real money balances:

$$\frac{M_t}{P_t} = 1 - \lambda \frac{P_{t+1}}{P_t}.$$

Note that the prices in t and $t+1$ affect the agent's optimal choices of consumption and real balances only through the ratio $\frac{p_{t+1}}{p_t}$.

Let $\Pi_{t+1} \equiv \frac{p_{t+1}}{p_t}$ denote the *actual* inflation rate between periods t and $t+1$. Then, the demand for real balances can be written as

$$\frac{M_t}{p_t} = 1 - \lambda \Pi_{t+1}.$$

Note that if inflation exceeds $\frac{1}{\lambda}$, the real rate of return on money, $\frac{p_t}{p_{t+1}}$, is “too low” and the young agent would like to borrow, not save. However, in a two-period overlapping generations setup, this is impossible. Thus, for inflation rates greater than or equal to $\frac{1}{\lambda}$, the young agent would just consume his endowment. Hence, for money to be held (i.e., for real balances to be positive) the inflation rate must be less than $\frac{1}{\lambda}$. Thus,

$$(8) \quad \frac{M_t}{p_t} = \max(0, 1 - \lambda \Pi_{t+1}).$$

Monetary policy (6) implies

$$\frac{M_t}{p_t} \Pi_t = \theta \frac{M_{t-1}}{p_{t-1}}.$$

The asset market clearing condition implies that money supplied in each period must equal money demand in that period. We can substitute money demand into the above equation and get

$$\max(0, 1 - \lambda \Pi_{t+1}) \Pi_t = \theta \max(0, 1 - \lambda \Pi_t).$$

Thus, the law of motion for the inflation rate is

$$(9) \quad \max(0, 1 - \lambda \Pi_{t+1}) = \theta \frac{\max(0, 1 - \lambda \Pi_t)}{\Pi_t}.$$

Given perfect foresight, a *RE equilibrium* is a sequence of quantities and prices— $c_{1,t}, c_{2,t-1}, \frac{M_t}{p_t}, p_t, \Pi_t$ from $t = 0, 1, \dots, \infty$ —such that agents in each generation choose consumption and real balances optimally, the asset market clears in every period, and the evolution of inflation satisfies equation (9). (Recall that in Section 2, the RE equilibrium is unique.)

Equation (9) clearly admits two steady states $\Pi_t = \Pi_{t+1} = \Pi^*$: Either $\Pi^* = \theta$ or $\Pi^* = \frac{1}{\lambda}$, where θ is the monetary steady state. It is easy to see that for $\Pi_t < \frac{1}{\lambda}$, the real balances are positive, so equation (9) can be simplified as

$$\Pi_{t+1} = \theta + \frac{1}{\lambda} - \left(\frac{\theta}{\lambda}\right) \frac{1}{\Pi_t},$$

and the mapping from Π_t to Π_{t+1} is increasing and concave. Furthermore, around θ , the slope is greater than 1, so the monetary steady state is unstable.

3.2 Learning

Suppose that instead of perfect foresight on prices, generation- t agents have to make their optimal choices according to their expectation of p_{t+1} . In other words,

$$\frac{M_t}{p_t} = 1 - \lambda \frac{p_{t+1}^e}{p_t}$$

where the superscript e denotes expected value. Let $\Pi_{t+1}^e \equiv \frac{p_{t+1}^e}{p_t}$ denote the *expected* inflation rate between periods t and $t+1$. Again, if expected inflation exceeds $\frac{1}{\lambda}$, then the expected real return on money is too low and the young agents will not hold any money. Thus, similar to (8), the demand for real balances under learning is

$$(10) \quad \frac{M_t}{p_t} = \max\left(0, 1 - \lambda \Pi_{t+1}^e\right).$$

As in Section 3.1, using the money supply at time t , we get $\frac{M_t}{p_t} \Pi_t = \theta \frac{M_{t-1}}{p_{t-1}}$. Substituting for money demand in this equation, we get the asset market clearing condition:

$$\max\left(0, 1 - \lambda \Pi_{t+1}^e\right) \frac{p_t}{p_{t-1}} = \theta \max\left(0, 1 - \lambda \Pi_t^e\right).$$

The above relationship, in turn, yields the law of motion for inflation under learning:

$$(11) \quad \Pi_t = \theta \frac{\max\left(0, 1 - \lambda \Pi_t^e\right)}{\max\left(0, 1 - \lambda \Pi_{t+1}^e\right)}.$$

Note that unlike in Section 3.1, the actual inflation in the case of learning is influenced by the inflation forecast in period $t-1$ as well as the forecast in period t . What remains to be specified is how agents forecast inflation.

Suppose agents at time t form their expectations as follows:

$$(12) \quad \Pi_{t+1}^e = \Pi_t^e + \gamma \left[\Pi_{t-1} - \Pi_t^e \right],$$

where the gain $\gamma \in (0, 1)$ is a constant.¹ Equation (12) is the counterpart to learning rule (5) in Section 2.

When the agents forecast the inflation Π_{t+1}^e they do not know the price p_t and, hence, do not know the actual inflation Π_t . Lack of knowledge of p_t does not affect the demand for real balances, since equation (10) implies that the demand depends on the *ratio* of prices.

Using equation (11) to substitute for Π_{t-1} in the constant-gain learning rule, we get the law of motion for expected inflation Π_{t+1}^e under learning:

$$(13) \quad \Pi_{t+1}^e = \Pi_t^e + \gamma \left[\theta \frac{\max(0, 1 - \lambda \Pi_{t-1}^e)}{\max(0, 1 - \lambda \Pi_t^e)} - \Pi_t^e \right]$$

$$(14) \quad = \Pi_t^e + \gamma \left[\theta \frac{\max(0, 1 - \lambda \Pi_t^e + \lambda \Delta \Pi_t^e)}{\max(0, 1 - \lambda \Pi_t^e)} - \Pi_t^e \right],$$

where $\Delta \Pi_t^e \equiv \Pi_t^e - \Pi_{t-1}^e$.

A *learning equilibrium* is a sequence of quantities, prices, and forecasts— $c_{1,t}, c_{2,t-1}, \frac{M_t}{P_t}, p_t, \Pi_{t+1}^e$ from $t = 0, 1, \dots, \infty$ —such that agents in each generation choose consumption and real balances optimally based on their forecast of inflation, the asset market clears in every period, and the two-way feedback from expected inflation to actual inflation satisfies equation (14).

Remark 1. For equation (14) to describe a learning equilibrium, we have to impose an additional restriction that $\Pi_{t+1}^e < \frac{1}{\lambda}$. If any element in the sequence of Π_t^e 's exceeds $\frac{1}{\lambda}$, then (14) cannot be used to recover future expected inflations.

Several features of equation (14) are worth noting. First, (14) has a steady state: $\Pi_{t-1}^{*e} = \Pi_t^{*e} = \Pi_{t+1}^{*e} = \Pi^{*e} = \theta$. Second, when the steady-state expected inflation equals θ , equation (11) implies actual inflation is also equal to θ . Third, equation (14) is a second-order nonlinear difference equation, which is not tractable. Second-order difference equations are common in other models in the learning literature. See, for instance, Marcet and Sargent (1989b); Bullard (1991); Bullard (1994); Evans, Honkapohja, and Marimon (2001); Marcet and Nicolini (2003); Adam, Marcet, and Nicolini (2016); and Adam, Marcet, and Beutel (2017). As we illustrate in the appendix, second-order nonlinear difference equations are not germane to the monetary model described here; a change in the information set and the learning rule results in a first-order difference equation.

4 RESULTS

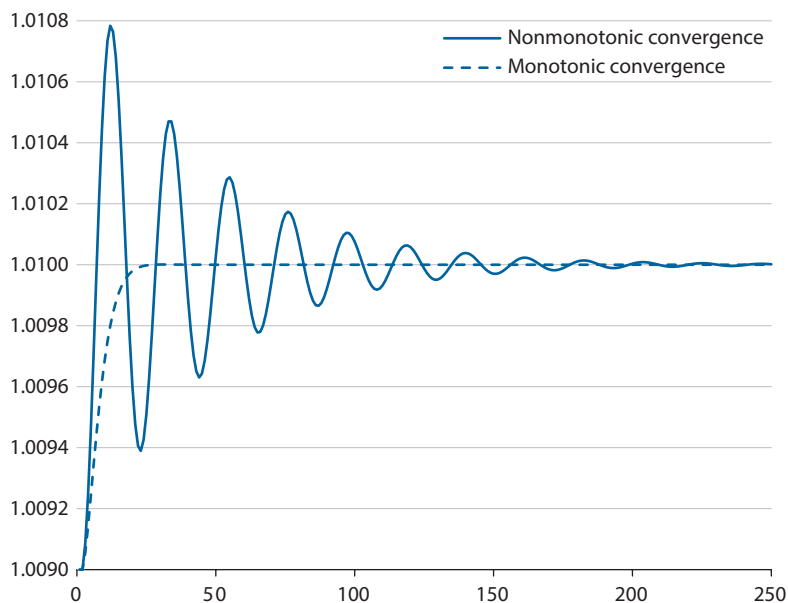
In this section, we first simulate the dynamics of expected inflation in (14) and examine its convergence properties. We then show the dependence of the convergence property on the gain parameter. For some values of the gain parameter, the learning dynamics converge to the RE monetary steady state and for others the convergence does not occur.

4.1 Numerical Results

In the numerical exercises below, we simulate (14) for different values of $\gamma \in (0, 1)$. We set the monetary steady state $\theta = 1.01$; we set $\frac{1}{\lambda} = 1.1$, so the assumptions $\lambda < 1$ and $\theta < \frac{1}{\lambda}$ are satisfied. Since (14) is a second-order difference equation, we need two initial conditions to simulate the dynamics. We set them to be the same: $\Pi_0^e = \Pi_1^e = \pi$. We explore the dynamic properties by altering γ and π .

Figure 1


Convergence to the Monetary Steady State



NOTE: The x-axis is t and the y-axis is Π_{t+1}^e . The parameters are $\theta = 1.01$ and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi_0^e = \Pi_1^e = \pi = 1.009$. For nonmonotonic convergence, $\gamma = 0.085$ and for monotonic convergence $\gamma = 0.055$.

Figure 2

Agents Don't Learn the Actual Inflation



NOTE: The x-axis is t and the y-axis is Π_t^e for expected inflation and Π_t for actual inflation. The parameters are $\gamma = 0.0891$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi_0^e = \Pi_1^e = \pi = 1.009$.

Nonmonotonic Convergence. For $\gamma = 0.0850$ and for the initial condition $\pi = 1.009$, the trajectory of Π_t^e converges to RE monetary steady state θ nonmonotonically. It follows a damped-oscillation path as shown in Figure 1. For $\gamma = 0.055$, the convergence is monotonic starting from the same initial conditions.

Instability. For $\gamma = 0.0891$ and the same value of π , the expected-inflation trajectory settles to a stable orbit; see Figure 2.

Agents Never Learn the Actual Inflation. Figure 2 also illustrates the actual inflation, which settles down to a different stable orbit; that is, the agents never learn the actual inflation.

In sum, the learning dynamics neither display convergence to the RE monetary steady state nor exhibit zero forecast errors in the long run. In the next section, we examine the role of the gain parameter γ in convergence to the monetary steady state. If the convergence is attained, then the forecast errors will automatically go to zero in the long run since, as noted earlier, if expected inflation equals θ , then actual inflation equals θ as well. We show that if the gain parameter exceeds a critical value, then the learning dynamics do not converge to the RE monetary steady state.

4.2 Convergence Depends on the Gain Parameter

In this section, we focus on the local neighborhood around θ to illustrate the dependence of convergence properties of the learning dynamics (14) on γ . To begin, rewrite (14) as

$$\Delta\Pi_{t+1}^e = \gamma \left[\theta \frac{1 - \lambda\Pi_t^e + \lambda\Delta\Pi_t^e}{1 - \lambda\Pi_t^e} - \Pi_t^e \right].$$

Note that this rewritten equation does not have the max operator in (14), since we are examining the dynamic system locally around θ . For the dynamics to converge to θ , both $\Delta\Pi_{t+1}^e$ and $\Delta\Pi_t^e$ have to converge to zero. Suppose we start from the initial conditions $\Pi_{t-1}^e = \Pi_t^e \neq \theta$. Then, $\Delta\Pi_t^e = 0$, so $\Delta\Pi_{t+1}^e = \gamma[\theta - \Pi_t^e]$. Clearly, $\Delta\Pi_{t+1}^e \neq 0$ in the neighborhood of θ . So, the initial condition $\Delta\Pi_t^e = 0$ does not guarantee $\Delta\Pi_{t+1}^e = 0$ unless we start exactly at the monetary steady state. (Note the reverse as well: Setting $\Delta\Pi_{t+1}^e = 0$ does not guarantee $\Delta\Pi_t^e = 0$ unless $\Pi_{t+1}^e = \Pi_t^e = \theta$.) Figure 3 illustrates this point for the case where the expected inflation converges to θ nonmonotonically, and Figure 4 illustrates it for the case where the expected inflation converges to θ monotonically.

To examine the conditions under which both $\Delta\Pi_t^e$ and $\Delta\Pi_{t+1}^e$ go to zero, we transform the second-order difference equation into two first-order equations. Define


$$z_t \equiv \frac{1}{\gamma} (\Pi_t^e - \Pi_{t-1}^e).$$

We can then write (14) as a system of *first-order* difference equations in Π_t^e and z_t :

$$\Pi_{t+1}^e = \Pi_t^e + \gamma \left[\theta \frac{1 - \lambda\Pi_t^e + \lambda\gamma z_t}{1 - \lambda\Pi_t^e} - \Pi_t^e \right],$$

Figure 3

Expected Inflation Converges to the Monetary Steady State Nonmonotonically



NOTE: For Panel A, the x-axis is t and the y-axis is Π_t^e for expected inflation between periods t and $t + 1$. For Panel B, the x-axis t and the y-axis is $\Delta\Pi_t^e = \Pi_t^e - \Pi_{t-1}^e$ and $\Delta\Pi_{t+1}^e = \Pi_{t+1}^e - \Pi_t^e$. The parameters are $\gamma = 0.085$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi_0^e = \Pi_1^e = \pi = 1.0$.

$$z_{t+1} = z_t + \left[\theta \frac{1 - \lambda\Pi_t^e + \lambda\gamma z_t}{1 - \lambda\Pi_t^e} - \Pi_t^e - z_t \right].$$

Or,

$$\Delta\Pi_{t+1}^e = \gamma \left[\theta \frac{1 - \lambda\Pi_t^e + \lambda\gamma z_t}{1 - \lambda\Pi_t^e} - \Pi_t^e \right],$$

$$\Delta z_{t+1} = \left[\theta \frac{1 - \lambda\Pi_t^e + \lambda\gamma z_t}{1 - \lambda\Pi_t^e} - \Pi_t^e - z_t \right].$$


In the above first-order system, the first equation is of order γ , while the second is of order 1. For small γ , z evolves faster than Π^e , so z arrives at the stable point before Π^e does.

To analyze the dynamics of this vector, (Π_t^e, z_t) , continuous-time methods for Markov chains are useful. The first-order system can be written as

$$(15) \quad \dot{\Pi}^e = \gamma \left[\theta \frac{1 - \lambda\Pi^e + \lambda\gamma z}{1 - \lambda\Pi^e} - \Pi^e \right]$$

Figure 4

Expected Inflation Converges to the Monetary Steady State Monotonically



NOTE: For Panel A, the x-axis is t and the y-axis is Π_t^e for expected inflation between periods t and $t+1$. For Panel B, the x-axis t and the y-axis is $\Delta\Pi_t^e = \Pi_t^e - \Pi_{t-1}^e$ and $\Delta\Pi_{t+1}^e = \Pi_{t+1}^e - \Pi_t^e$. The parameters are $\gamma = 0.055$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi_0^e = \Pi_1^e = \pi = 1.0$.

$$(16) \quad \dot{z} = \left[\theta \frac{1 - \lambda\Pi^e + \lambda\gamma z}{1 - \lambda\Pi^e} - \Pi^e \right] - z$$

or, more concisely,

$$\begin{bmatrix} \dot{\Pi}^e \\ \dot{z} \end{bmatrix} = \Psi \left(\begin{bmatrix} \Pi^e \\ z \end{bmatrix} \right).$$

If $\Pi^e = \theta$ is a stable point of the learning dynamics, then $\dot{\Pi}^e = 0$ and $\dot{z} = 0$. It is easy to see

$$\Psi \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix} \right) = 0.$$

To account for the difference in speeds noted earlier, we solve the system sequentially by first holding $\Pi^e < \frac{1}{\lambda}$ fixed and solving (16) for $\dot{z} = 0$ to derive a relationship between z and Π^e .

We then use this relationship to derive a condition for $\dot{\Pi}^e = 0$ using (15). This sequential

approach yields the familiar result that for γ close to 0, $\Pi^e \rightarrow \theta$ as $t \rightarrow \infty$. This result, however, does not answer how “small” γ should be in simulations to guarantee convergence to RE.

We answer the question partially by examining local stability of the first-order system around $(\Pi^e, z) = (\theta, 0)$. Consider the Hessian of the right-hand sides of (15) and (16):

$$\begin{bmatrix} -\theta\gamma^2 \frac{\lambda^2 z}{(1-\lambda\Pi^e)^2} - \gamma & \frac{\lambda\theta\gamma^2}{1-\lambda\Pi^e} \\ -\theta\gamma \frac{\lambda^2 z}{(1-\lambda\Pi^e)^2} - 1 & \frac{\lambda\theta\gamma}{1-\lambda\Pi^e} - 1 \end{bmatrix}.$$

At $(\Pi^e, z) = (\theta, 0)$, the above matrix becomes

$$\begin{bmatrix} -\gamma & \frac{\lambda\theta\gamma^2}{1-\lambda\theta} \\ -1 & \frac{\lambda\theta\gamma}{1-\lambda\theta} - 1 \end{bmatrix}.$$

The eigenvalues for the dynamic system, denoted by μ , then solve

$$\mu^2 + \left(1 + \gamma - \frac{\lambda\theta\gamma}{1-\lambda\theta}\right)\mu + \gamma = 0.$$

The two eigenvalues have negative real parts if and only if

$$-\left(1 + \gamma - \frac{\lambda\theta\gamma}{1-\lambda\theta}\right) < 0;$$

that is,

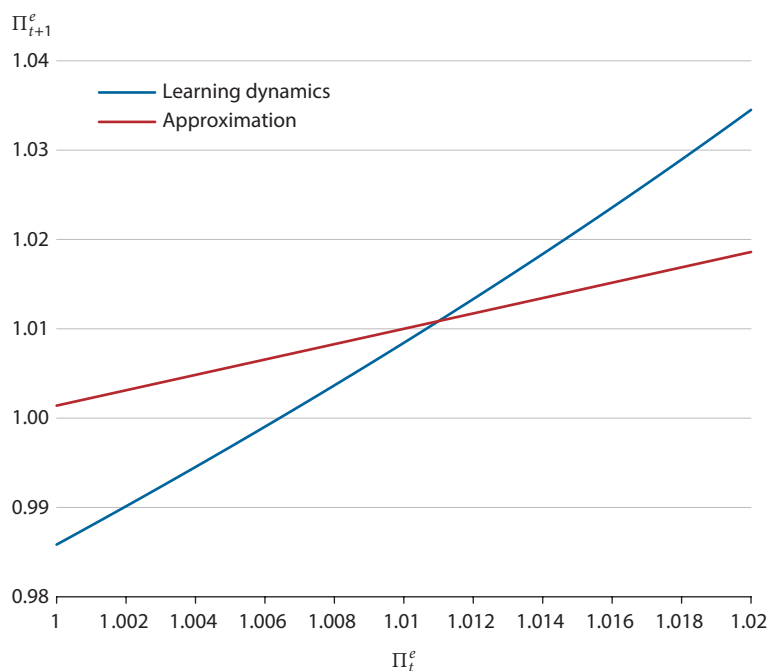
$$\begin{aligned} \frac{1}{\gamma} &> -1 + \frac{\lambda\theta}{1-\lambda\theta} \\ &= -1 + \frac{\theta}{\frac{1}{\lambda} - \theta}. \end{aligned}$$

In our numerical simulations, $\theta = 1.01$ and $\frac{1}{\lambda} = 1.1$ so that the right-hand side equals 10.22.

Thus, starting from an initial condition in the local neighborhood of θ , if γ is greater than 0.0978, then the learning dynamics will not converge to the RE monetary steady state. This is a sufficient condition for *lack of convergence*. In practice, the dynamics may not converge for some γ less than 0.0978, as demonstrated in Figure 2.

Approximation. One could adopt a different approach to study the local behavior of (14) around θ by using a first-order difference equation to approximate (14), as in Evan, Honkapohja, and Marimom (2001), for instance:

$$(17) \quad \Pi_{t+1}^e = \Pi_t^e + \gamma[\theta - \Pi_t^e].$$

Figure 5**Current and Future Expected Inflation: Learning Dynamics Versus Local Approximation**

NOTE: The parameters are $\gamma = 0.14$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. For the learning dynamics, $\Pi_{t-1}^e = 1.011$.

This requires only one initial condition. It is easy to see that as $t \rightarrow \infty$, the expected inflation $\Pi_t^e \rightarrow \theta$, independent of γ . Thus, the first-order approximation could mistakenly suggest convergence to the RE monetary steady state even when the learning dynamics in (14) imply nonconvergence. Figure 5 illustrates the result.

To summarize, for the learning dynamics to converge to the RE monetary steady state, the learning rule has to be tightly parameterized based on economic fundamentals.

5 CONCLUDING REMARKS

We used a simple overlapping generations model to study whether the learning dynamics converge to the RE monetary steady state. We showed via numerical experiments that the learning dynamics can produce cycles and display paths where agents never learn the actual inflation. To guarantee the convergence to RE, the learning rule has to be tightly parameterized based on economic fundamentals. We show that the gain parameter used in the learning rule affects whether the learning dynamics converge to the RE monetary steady state.

An alternative is to consider learning rules with different information sets. One such learning rule is described in the appendix. Such rules result in learning dynamics that are described by a first-order difference equation. For first-order difference equations, Kushner

and Yin (1997) demonstrate that the asymptotic properties of the learning algorithm are determined by the mathematical properties of the forecast error, not by the decreasing- or constant-gain functions. Thus, such learning rules can be used to study convergence to the RE monetary steady state without restricting the gain parameter. ■

APPENDIX: LEARNING RULE USING CURRENT PRICE

Recall that in the learning rule (12), agents in period t do not use the current price p_t to forecast the inflation Π_{t+1}^e . That is, they are forecasting the ratio $\frac{p_{t+1}^e}{p_t}$, but it is assumed that they do not know the denominator p_t . So, the actual inflation, Π_t , is not part of the learning rule (12).

Constant-Gain Learning Rule

Consider an alternative information set and the associated learning rule that uses the current price p_t :

$$(18) \quad \Pi_{t+1}^e = \Pi_t^e + \gamma [\Pi_t - \Pi_t^e].$$

The change in the learning rule affects the forecast of inflation, but not the feedback from expected inflation to actual inflation. That is, equation (11) continues to hold.

Using equation (11) to substitute for Π_t in the learning rule above, we get the law of motion for expected inflation Π_{t+1}^e under learning:

$$\begin{aligned} \Pi_{t+1}^e &= \Pi_t^e + \gamma \left[\theta \frac{\max(0, 1 - \lambda \Pi_t^e)}{\max(0, 1 - \lambda \Pi_{t+1}^e)} - \Pi_t^e \right] \\ &= \Pi_t^e + \gamma \left[\theta \frac{\max(0, 1 - \lambda \Pi_{t+1}^e + \lambda \Delta \Pi_{t+1}^e)}{\max(0, 1 - \lambda \Pi_{t+1}^e)} - \Pi_t^e \right], \end{aligned}$$

where $\Delta \Pi_{t+1}^e \equiv \Pi_{t+1}^e - \Pi_t^e$.

The law of motion implies a steady-state $\Pi^e = \theta$, and when expected inflation reaches θ , so does actual inflation. Furthermore, the law of motion is just a first-order difference equation.

NOTES

- ¹ In Chien, Cho, and Ravikumar (2020), we demonstrate that the decreasing-gain learning rule in Bullard (1994), where the gain is a least-squares forecast based on past prices, takes a similar form as the constant-gain rule (12).

REFERENCES

- Adam, Klaus; Marcet, Albert and Beutel, Johannes. "Stock Price Booms and Expected Capital Gains." *American Economic Review*, August 2017, 107(8), pp. 2352-408; <https://doi.org/10.1257/aer.20140205>.
- Adam, Klaus; Marcet, Albert and Nicolini, Juan Pablo. "Stock Market Volatility and Learning." *Journal of Finance*, 2016, 71(1), pp. 33-81; <https://doi.org/10.1111/jofi.12364>.
- Bray, Margaret. "Learning, Estimation and the Stability of Rational Expectations." *Journal of Economic Theory*, 1982, 26, pp. 318-39; [https://doi.org/10.1016/0022-0531\(82\)90007-2](https://doi.org/10.1016/0022-0531(82)90007-2).
- Bray, Margaret M. and Savin, N. Eugene. "Rational Expectations Equilibria, Learning and Model Specification." *Econometrica*, September 1986, 54(5), pp. 1129-160; <https://doi.org/10.2307/1912325>.
- Bullard, James B. "Learning, Rational Expectations and Policy: A Summary of Recent Research." Federal Reserve Bank of St. Louis *Review*, January/February 1991, 73(1), pp. 50-60; <https://doi.org/10.20955/r.73.50-60>.
- Bullard, James B. "Learning Equilibria." *Journal of Economic Theory*, December 1994, 64(2), pp. 468-85; <https://doi.org/10.1006/jeth.1994.1077>.
- Chien, YiLi; Cho, In-Koo and Ravikumar, B. "Convergence to Rational Expectations in Learning Models: A Note of Caution." Working Paper 2020-027A, Federal Reserve Bank of St. Louis, 2020; <https://doi.org/10.20955/wp.2020.027>.
- Evans, George W. and Honkapohja, Seppo. "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models." *Econometrica*, 1995, 63(1), pp. 195-206; <https://doi.org/10.2307/2951703>.
- Evans, George W. and Honkapohja, Seppo. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- Evans, George W.; Honkapohja, Seppo and Marimon, Ramon. "Convergence in Monetary Inflation Models with Heterogeneous Learning Rules." *Macroeconomic Dynamics*, 2001, 5, pp. 1-31; <https://doi.org/10.1017/S1365100501018016>.
- Kushner, Harold J. and Yin, G. George. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, 1997.
- Lucas, Robert E. Jr. "Adaptive Behavior and Economic Theory." *Journal of Business*, 1986, 59(4), pp. S401-26; <https://doi.org/10.1086/296377>.
- Marcet, Albert and Nicolini, Juan P. "Recurrent Hyperinflations and Learning." *American Economic Review*, December 2003, 93(5), pp. 1476-498; <https://doi.org/10.1257/000282803322655400>.
- Marcet, Albert and Sargent, Thomas J. "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models." *Journal of Economic Theory*, 1989a, 48, pp. 337-68; [https://doi.org/10.1016/0022-0531\(89\)90032-X](https://doi.org/10.1016/0022-0531(89)90032-X).
- Marcet, Albert and Sargent, Thomas J. "Convergence of Least-Squares Learning and the Dynamics of Hyperinflation," in William Barnett, John Geweke, and Karl Shell, eds., *Economic Complexity: Chaos, Sunspots, Bubbles and Nonlinearity*. Cambridge University Press, 1989b, pp. 119-137.
- Sargent, Thomas J. *The Conquest of American Inflation*. Princeton University Press, 1999.
- Woodford, Michael D. "Learning to Believe in Sunspots." *Econometrica*, 1990, 58, pp. 277-307; <https://doi.org/10.2307/2938205>.

