ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS

WORKING PAPER SERIES

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

Authors	YiLi Chien, In-Koo Cho, and B. Ravikumar
Working Paper Number	2020-027B
Revision Date	September 2020
Citable Link	https://doi.org/10.20955/wp.2020.027
Suggested Citation	Chien, Y., Cho, I.-K., Ravikumar, B., 2020; Convergence to Rational Expectations in Learning Models: A Note of Caution, Federal Reserve Bank of St. Louis Working Paper 2020-027. URL https://doi.org/10.20955/wp.2020.027

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

YiLi Chien [†] In-Koo Cho [‡] B. Ravikumar[§]

September 19, 2020

Abstract

This paper illustrates a challenge in analyzing the learning algorithms resulting in second-order difference equations. We show in a simple monetary model that the learning dynamics do not converge to the rational expectations monetary steady state. We then show that to guarantee convergence, the gain parameter used in the learning rule has to be restricted based on economic fundamentals in the monetary model.

JEL codes: C60, D84

Keywords: learning algorithm, gain function, convergence, rational expectations equilibrium.

[*]We 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. The views in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

[†]Federal Reserve Bank of St. Louis; Email: yilichien@gmail.com.

[‡]Department of Economics, Emory University, Hanyang University, and Federal Reserve Bank of St. Louis; Email: in-koo.cho@emory.edu; Web: https://sites.google.com/site/inkoocho.

[§]Federal Reserve Bank of St. Louis; Email: b.ravikumar@wustl.edu.

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 the variables, so they 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 RE equilibrium include Bray (1982), Bray and Savin (1986), Lucas (1986), Marcet and Sargent (1989a), and Woodford (1990).

The data did not influence the agents' learning rule in the first generation of learning models, although the learning rule influenced the data. See, for instance, Cagan (1956). In more recent formulations, the learning rule depends on the data, so there is feedback in both directions– from data to the learning rule and vice versa. Examples include Bray (1982) on prices, Marcet and Sargent (1989b) on unpleasant monetarist arithmetic, and Evans and Honkapohja (1995) on business cycles.

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. When we do not impose rational expectations (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. The two-way feedback results in a second-order nonlinear difference equation in expected inflation. This complicates the analysis of the dynamic system since second-order nonlinear difference equations do not have known solutions.

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 non-zero forecast error in the limit, i.e., 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 two 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. Writing the second-order difference equation as a system of two first-order difference equations, 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 restrictions that the gain has to satisfy in order to guarantee convergence to the monetary steady state. These restrictions depend on economic fundamentals. Learning rules that result in dynamics described by a first-order difference equation 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 second-order difference equation that drives the learning dynamics. Section 4 provides simulation results that show non-convergence 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 (1982). 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 $p_t, p_{t-1}, \ldots, p_1$, so the data affects the learning rule. Let p^e_{t+1} denote the time-t forecast of next period's price:

$$\begin{aligned} p^e_{t+1} &= \frac{1}{t}\left[p_t + p_{t-1} + \ldots + p_1\right] \\ &= \frac{1}{t}p_t + \frac{t-1}{t}p^e_t. \end{aligned} \tag{1}$$

The evolution of p^e_t, or the forecasting equation, could then be expressed as:

$$p^e_{t+1} = p^e_t + \frac{1}{t}\left(p_t - p^e_t\right). \tag{2}$$

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:

$$p_t = A - Bp^e_t, \tag{3}$$

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 $p_t = A/(B+1)$, i.e., 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

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

It is easy to see from (4) that as $t \to \infty$ 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 forecast using least squares estimation, which is a generalized form of the sample average of past observations; see, for instance, Evans and Honkapohja (2001). 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:

$$p_{t+1}^e = p_t^e + \delta\left(p_t - p_t^e\right) \tag{5}$$

where $\delta \in (0, 1)$. 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, i.e., convergence in probability, we obtain weak convergence of the forecast to the population mean as $t \to \infty$. 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 to obtain accuracy or convergence speed, without altering the convergence properties of the algorithm. See, for instance, Sargent (1999). In the next section, we illustrate using a constant-gain learning algorithm 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 simpler version of Example 1 in Bullard (1994) that illustrates the convergence properties of the learning rule.

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 $t = 0, 1, 2, ...$ 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 his 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:

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

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

$$c_{1,t} + \frac{p_{t+1}}{p_t} c_{2,t} \leq 2 + 2\lambda \frac{p_{t+1}}{p_t}. \tag{7}$$

The problem of generation t is therefore

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

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,

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

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

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

Thus, the law of motion for the inflation rate is:

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

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,

$$\frac{M_t}{p_t} = 1 - \lambda\frac{p^e_{t+1}}{p_t},$$

where the superscript "e" denotes expected value. Let $\Pi^e_{t+1} \equiv \frac{p^e_{t+1}}{p_t}$ 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

$$\frac{M_t}{p_t} = \max\left(0, 1 - \lambda\Pi^e_{t+1}\right). \tag{10}$$

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:

$$\max\left(0, 1 - \lambda\Pi^e_{t+1}\right)\frac{p_t}{p_{t-1}} = \theta\max\left(0, 1 - \lambda\Pi^e_t\right).$$

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

$$\Pi_t = \theta\,\frac{\max\left(0, 1 - \lambda\Pi^e_t\right)}{\max\left(0, 1 - \lambda\Pi^e_{t+1}\right)}. \tag{11}$$

Note that unlike 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.

Recursive learning Suppose agents at time t form their expectations as follows:

$$\Pi_{t+1}^e = \Pi_t^e + \delta\left[\Pi_{t-1} - \Pi_t^e\right], \tag{12}$$

where the gain $\delta \in (0,1)$ is a constant.[1] 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 and is invariant to price p_t so long as the ratio $\frac{p_{t+1}^e}{p_t}$ remains the same.

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:

$$\begin{aligned}
\Pi_{t+1}^e &= \Pi_t^e + \delta\left[\theta\,\frac{\max\left(0, 1-\lambda\Pi_{t-1}^e\right)}{\max\left(0, 1-\lambda\Pi_t^e\right)} - \Pi_t^e\right] && (13)\\
&= \Pi_t^e + \delta\left[\theta\,\frac{\max\left(0, 1-\lambda\Pi_t^e + \lambda\Delta\Pi_t^e\right)}{\max\left(0, 1-\lambda\Pi_t^e\right)} - \Pi_t^e\right], && (14)
\end{aligned}$$

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 $t = 0, 1, \ldots, \infty$—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.

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.

[1] In Appendix A, 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).

4.1 Numerical results

In the numerical exercises below, we simulate (14) for different values of $\delta \in (0, 1)$. 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 δ and π.

Nonmonotonic convergence For $\delta = 0.0850$ 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 $\delta = 0.055$, the convergence is monotonic starting from the same initial conditions.

Instability For $\delta = 0.0891$ 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 3 illustrates the expected and actual inflations. As noted earlier, expected inflation settles down to a stable orbit for some parameter configurations. However, the actual inflation settles down to a different stable orbit for the same parameter configuration, i.e., 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 δ 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.

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 δ. To do this, we transform the second-order difference equation into two first-order equations. Note that we don't need the max operator in (14) in the derivation below since we are examining the dynamic system around θ. Define

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

Figure 1: Convergence to the monetary steady state.

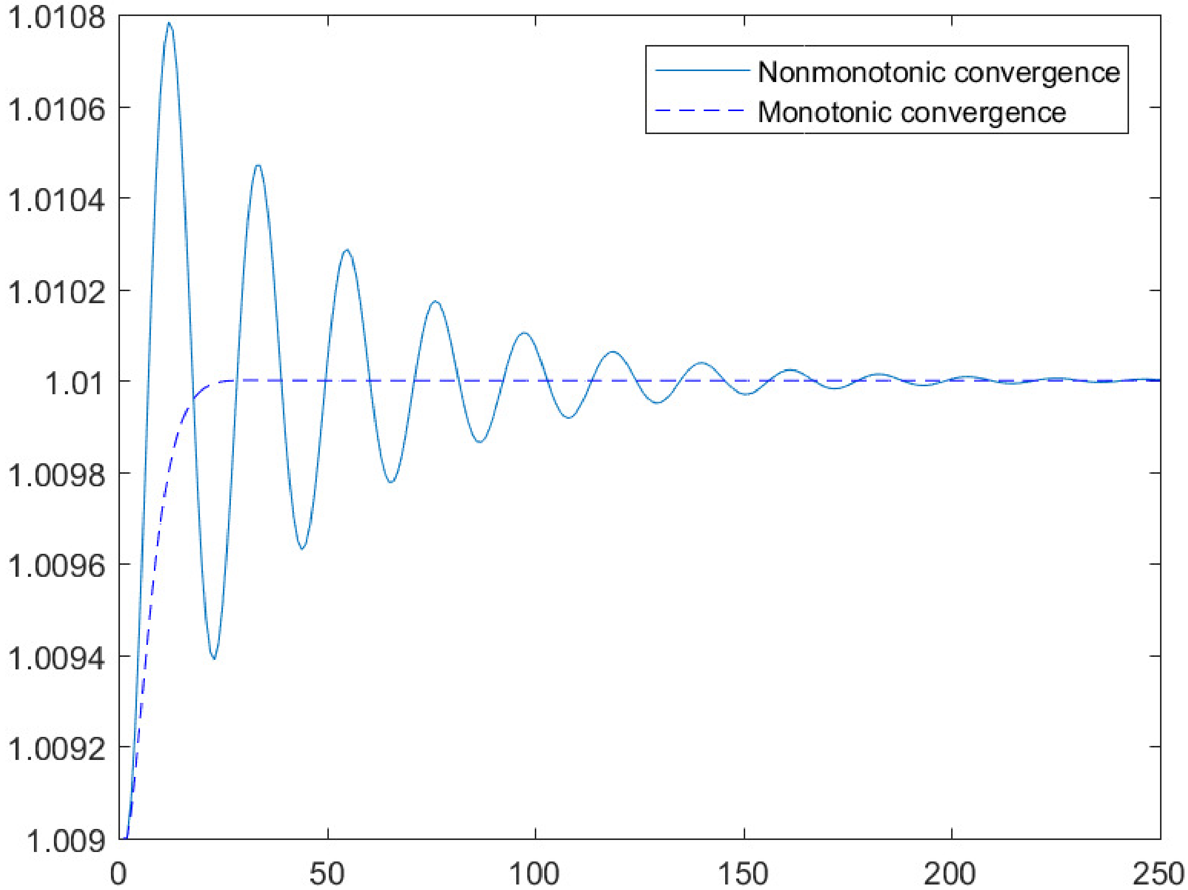

Notes: The horizontal axis is t and the vertical axis is Π^e_{t+1}. The parameters are $\theta = 1.01$ and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi^e_0 = \Pi^e_1 = \pi = 1.009$. For nonmonotonic convergence, $\delta = 0.085$ and for monotonic convergence $\delta = 0.055$.

Figure 2: Stable orbit under learning.

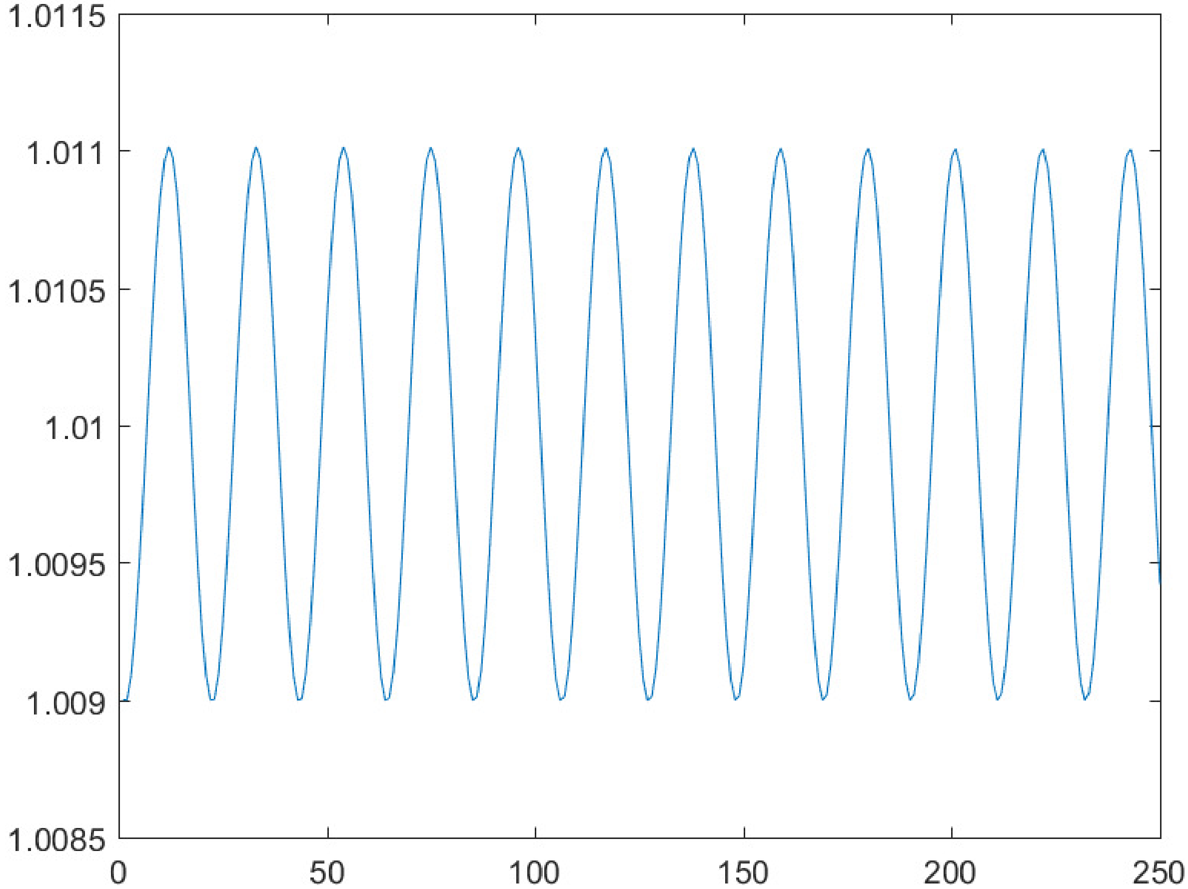

Notes: The horizontal axis is t and the vertical axis is Π^e_{t+1}. The parameters are $\delta = 0.0891$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi^e_0 = \Pi^e_1 = \pi = 1.009$.

Figure 3: Agents don't learn the actual inflation.

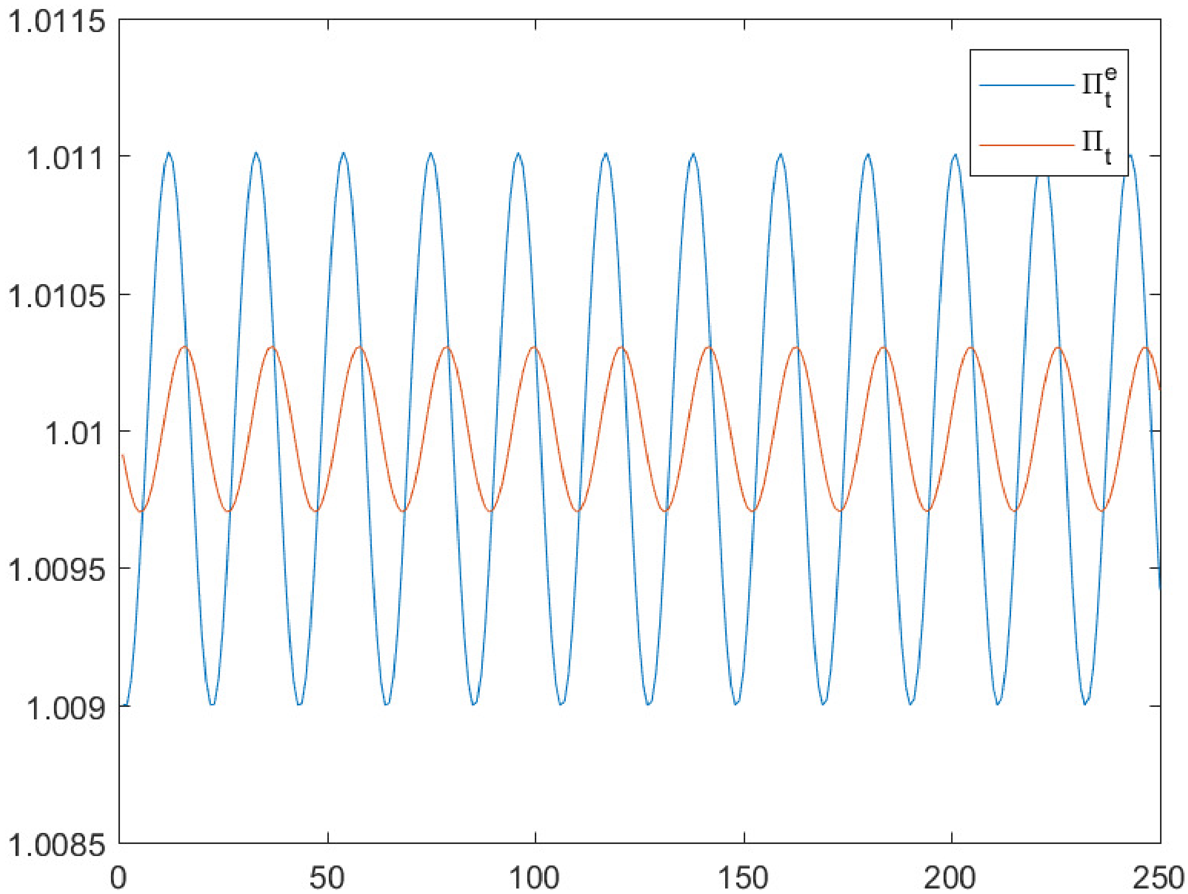

Notes: The horizontal axis is t and the vertical axis is Π_t^e for expected inflation and Π_t for actual inflation. The parameters are $\delta = 0.0891$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. The initial conditions are $\Pi_0^e = \Pi_1^e = \pi = 1.009$.

We can then write (14) as

$$z_{t+1} = z_t + \left[\theta \, \frac{1 - \lambda \Pi_t^e + \lambda \delta z_t}{1 - \lambda \Pi_t^e} - \Pi_t^e - z_t\right]. \tag{15}$$

We then consider the dynamics of the vector (Π_t^e, z_t), which is dictated by a *first-order* difference equation system of (14) and (15).

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

$$\begin{aligned} \dot{\Pi}^e &= \delta\left[\theta \, \frac{1 - \lambda \Pi^e + \lambda \delta z}{1 - \lambda \Pi^e} - \Pi^e\right] \\ \dot{z} &= \left[\theta \, \frac{1 - \lambda \Pi^e + \lambda \delta z}{1 - \lambda \Pi^e} - \Pi^e\right] - z \end{aligned}$$

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

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

A tedious calculation checking the eigenvalues of the Hessian around $(\Pi^e, z) = (\theta, 0)$ shows that (Π^e, z) is stable at $(\theta, 0)$ if and only if

$$\frac{\lambda\theta}{1 - \lambda\theta} - 1 < \frac{1}{\delta}.$$

In our numerical simulations, $\theta = 1.01$ and $\lambda = 0.9091$ so that the left-hand side equals 10.22. If δ 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.

First-order 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):

$$\Pi_{t+1}^e = \Pi_t^e + \delta\left[\theta - \Pi_t^e\right]. \tag{16}$$

Figure 4: Current and future expected inflations: Learning Dynamics vs Local Approximation.

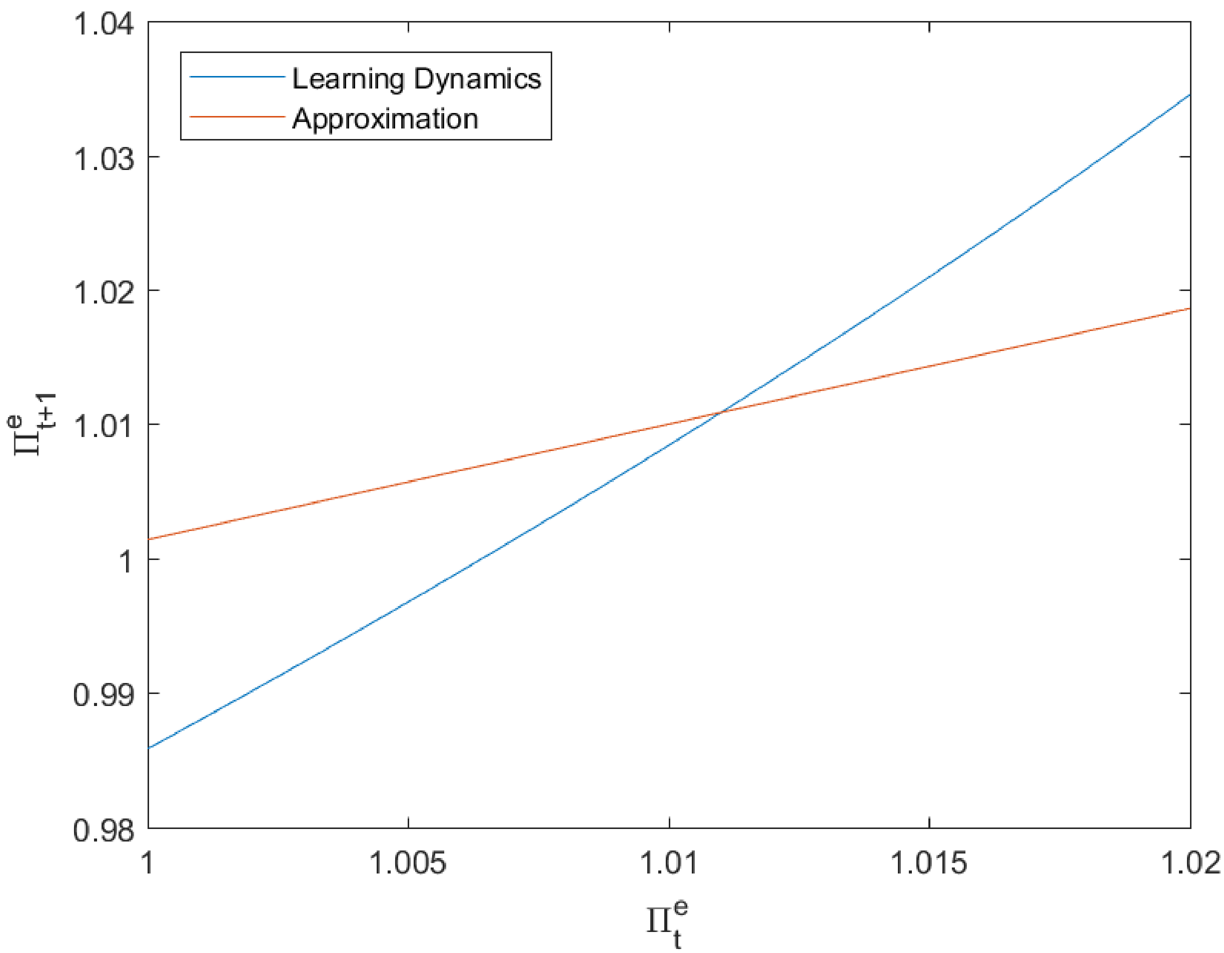

Notes: The horizontal axis is Π^e_t and the vertical axis is Π^e_{t+1}. The parameters are $\delta = 0.14$, $\theta = 1.01$, and $\frac{1}{\lambda} = 1.1$. For the learning dynamics, $\Pi^e_{t-1} = 1.011$.

This requires only one initial condition. It is easy to see that as $t \to \infty$, the expected inflation $\Pi^e_t \to \theta$, independent of δ. Thus, the first-order approximation could mistakenly suggest convergence to the RE monetary steady state even when the second-order learning dynamics imply non-convergence. Figure 4 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. Our specification results in a second-order difference equation in expected inflation. 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 that do not lead to second-order difference equations. One such learning rule is described in Appendix B. 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.

References

Bray, Margaret. 1982. "Learning, Estimation and the Stability of Rational Expectations." *Journal of Economic Theory* 26:318–339.

Bray, Margaret and N. Eugene Savin. 1986. "Rational Expectations Equilibria, Learning and Model Specification." *Econometrica* :1129–1160.

Bullard, James B. 1994. "Learning Equilibria." *Journal of Economic Theory* 64 (2):468–485.

Cagan, Phillip. 1956. "The Monetary Dynamics of Hyperinflation." In *Studies in the Quantity Theory of Money*, edited by Milton Friedman. University of Chicago Press, 25–117.

Evans, George W. and Seppo Honkapohja. 1995. "Local Convergence of Recursive Learning to Steady States and Cycles in Stochastic Nonlinear Models." *Econometrica* 63 (1):195–206.

———. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.

Kushner, Harold J. and G. George Yin. 1997. *Stochastic Approximation Algorithms and Applications*. Springer-Verlag.

Lucas, Robert. 1986. "Adaptive Behavior and Economic Theory." *The Journal of Business* 59 (4):S401–26.

Marcet, Albert and Thomas J. Sargent. 1989a. "Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models." *Journal of Economic Theory* 48:337–368.

———. 1989b. "Convergence of Least Squares Learning and the Dynamics of Hyperinflation." In *Economic Complexity: Chaos, Sunspots, Bubbles and Nonlinearity*, edited by William Barnett, John Geweke, and Karl Shell. Cambridge University Press, 119–137.

Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton University Press.

Woodford, Michael D. 1990. "Learning to Believe in Sunspots." *Econometrica* 58:277–307.

A Decreasing-gain learning rule

Suppose agents at time t form their expectations by estimating a first-order autoregression on price level, $p_{t-1} = \beta_t p_{t-2} + \varepsilon_t$, where ε is white noise. When they estimate the coefficient β_t they do not know the price p_t, so β_t is estimated using least squares and the data on p_s, $s = 0, 1, \ldots, t-1$. For the derivation below, we assume for simplicity that the expected inflation is always less than $\frac{1}{\lambda}$.

The estimate of Π^e_{t+1} at time t is given by

$$\Pi^e_{t+1} = \beta_t = \frac{\sum_{s=1}^{t-1} p_{s-1} p_s}{\sum_{s=1}^{t-1} p_{s-1}^2}.$$

The above equation can be written recursively as

$$\begin{aligned}
\Pi^e_{t+1} &= \frac{\sum_{s=1}^{t-2} p_{s-1} p_s}{\sum_{s=1}^{t-2} p_{s-1}^2} \frac{\sum_{s=1}^{t-2} p_{s-1}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} + \frac{p_{t-2} p_{t-1}}{\sum_{s=1}^{t-1} p_{s-1}^2} \\
&= \Pi^e_t \frac{\sum_{s=1}^{t-2} p_{s-1}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} + \frac{p_{t-2} p_{t-1}}{\sum_{s=1}^{t-1} p_{s-1}^2} \\
&= \Pi^e_t \frac{\sum_{s=1}^{t-2} p_{s-1}^2 + p_{t-2}^2 - p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} + \frac{p_{t-2} p_{t-1}}{\sum_{s=1}^{t-1} p_{s-1}^2} \\
&= \Pi^e_t \frac{\sum_{s=1}^{t-1} p_{s-1}^2 - p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} + \frac{p_{t-2} p_{t-1}}{\sum_{s=1}^{t-1} p_{s-1}^2} \\
&= \Pi^e_t - \Pi^e_t \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} + \frac{p_{t-2} p_{t-1}}{\sum_{s=1}^{t-1} p_{s-1}^2} \\
&= \Pi^e_t + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \left[\frac{p_{t-1}}{p_{t-2}} - \Pi^e_t \right] \\
&= \Pi^e_t + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \left[\Pi_{t-1} - \Pi^e_t \right]
\end{aligned}$$

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

$$\begin{aligned}
\Pi^e_{t+1} &= \Pi^e_t + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \left[\theta \frac{1 - \lambda \Pi^e_{t-1}}{1 - \lambda \Pi^e_t} - \Pi^e_t \right] && (17) \\
&= \Pi^e_t + \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \left[\theta \frac{1 - \lambda \Pi^e_t + \lambda \Pi^e_t - \lambda \Pi^e_{t-1}}{1 - \lambda \Pi^e_t} - \Pi^e_t \right], \\
&= \Pi^e_t + \delta_t \left[\theta \frac{1 - \lambda \Pi^e_t + \lambda \Delta \Pi^e_t}{1 - \lambda \Pi^e_t} - \Pi^e_t \right], && (18)
\end{aligned}$$

where δ_t and $\Delta \Pi^e_t$ are defined, respectively, as

$$\delta_t \equiv \frac{p_{t-2}^2}{\sum_{s=1}^{t-1} p_{s-1}^2} \qquad (19)$$

and

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

As noted in the text, (17) has a steady state: $\Pi_{t-1}^{*e} = \Pi_t^{*e} = \Pi_{t+1}^{*e} = \Pi^{*e} = \theta$, provided δ_t is finite in the long run. The behavior of δ_t is dictated by actual prices. From equation (11), we can see that if expected inflation reaches the steady state θ, so does actual inflation, which means prices grow at the rate θ. Then, we can write

$$\delta_t = \frac{\theta^{t-2} p_0}{p_0^2 + (\theta p_0)^2 + \ldots + (\theta^{t-2} p_0)^2}.$$

Thus, $\lim_{t\to\infty} \delta_t = 1 - \frac{1}{\theta^2} \in (0, 1-\lambda^2) \subseteq (0,1)$.

B 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:

$$\Pi_{t+1}^e = \Pi_t^e + \delta \left[\Pi_t - \Pi_t^e \right]. \tag{20}$$

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:

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

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.

Decreasing-Gain Learning Rule Suppose agents form their expectations by estimating a first-order autoregression on price level, $p_t = \beta_t p_{t-1} + \varepsilon_t$, where ε is white noise. The coefficient β_t is estimated in each period using all of the price information up to period t. That is, generation-t agents use p_s, $s = 0, 1, \ldots, t$ and their estimate of Π_{t+1}^e is given by

$$\Pi_{t+1}^e = \frac{\beta_t p_t}{p_t} = \beta_t = \frac{\sum_{s=1}^t p_{s-1} p_s}{\sum_{s=1}^t p_{s-1}^2}.$$

Again, we will assume for simplicity that the expected inflation is always less than $\frac{1}{\lambda}$.

The above equation can be written recursively as

$$\begin{aligned}
\Pi_{t+1}^e &= \frac{\sum_{s=1}^{t-1} p_{s-1}p_s}{\sum_{s=1}^{t-1} p_{s-1}^2}\frac{\sum_{s=1}^{t-1} p_{s-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} + \frac{p_{t-1}p_t}{\sum_{s=1}^{t} p_{s-1}^2} \\
&= \Pi_t^e \frac{\sum_{s=1}^{t-1} p_{s-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} + \frac{p_{t-1}p_t}{\sum_{s=1}^{t} p_{s-1}^2} \\
&= \Pi_t^e \frac{\sum_{s=1}^{t-1} p_{s-1}^2 + p_{t-1}^2 - p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} + \frac{p_{t-1}p_t}{\sum_{s=1}^{t} p_{s-1}^2} \\
&= \Pi_t^e \frac{\sum_{s=1}^{t} p_{s-1}^2 - p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} + \frac{p_{t-1}p_t}{\sum_{s=1}^{t} p_{s-1}^2} \\
&= \Pi_t^e - \Pi_t^e \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} + \frac{p_{t-1}p_t}{\sum_{s=1}^{t} p_{s-1}^2} \\
&= \Pi_t^e + \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2}\left[\frac{p_t}{p_{t-1}} - \Pi_t^e\right] \\
&= \Pi_t^e + \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2}\left[\Pi_t - \Pi_t^e\right]
\end{aligned}$$

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

$$\begin{aligned}
\Pi_{t+1}^e &= \Pi_t^e + \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2}\left[\theta\frac{1-\lambda\Pi_t^e}{1-\lambda\Pi_{t+1}^e} - \Pi_t^e\right] \\
&= \Pi_t^e + \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2}\left[\theta\frac{1-\lambda\Pi_t^e + \lambda\Pi_{t+1}^e - \lambda\Pi_{t+1}^e}{1-\lambda\Pi_{t+1}^e} - \Pi_t^e\right],
\end{aligned}$$

which can be simplified as

$$\Pi_{t+1}^e = \Pi_t^e + \delta_t\left[\theta\frac{1-\lambda\Pi_{t+1}^e + \lambda\Delta\Pi_{t+1}^e}{1-\lambda\Pi_{t+1}^e} - \Pi_t^e\right], \tag{21}$$

where δ_t and $\Delta\Pi_{t+1}^e$ are defined, respectively, as

$$\delta_t \equiv \frac{p_{t-1}^2}{\sum_{s=1}^{t} p_{s-1}^2} \tag{22}$$

and

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