Genetic Programming in the Overlapping Generations Model: An Illustration with the Dynamics of the Inflation Rate*

Shu-Heng Chen† and **Chia-Hsuan Yeh**‡

AI-ECON Research Group
Department of Economics
National Chengchi University
Taipei, Taiwan 11623
†E-mail: chchen@cc.nccu.edu.tw
‡E-mail: g3258501@grad.cc.nccu.edu.tw

Abstract. In this paper, genetic programming (GP) is employed to model learning and adaptation in the overlapping generations model, one of the most popular dynamic economic models. Using a model of inflation with multiple equilibria as an illustrative example, we show that our GP-based agents are able to coordinate their actions to achieve the Pareto-superior equilibrium (the low-inflation steady state) rather than the Pareto-inferior equilibrium (the high-inflation steady state). We also test the robustness of this result with different initial conditions, economic parameters, and GP control parameters.

1 Introduction

While there are several approaches to introducing *dynamic general equilibrium structures* to economics, the *overlapping generations model* (hereafter, **OLG**), proposed by Allais (1947) and Samuelson (1958), may be regarded as the most popular in current macroeconomics.

Despite its popularity, one of the technical issues which remain unsolved in the **OLG** is *how expectations and learning take place* in this overlapping-generations structure. In the early 80's, the assumptions of *perfect foresight* and *rational expectations* were adopted to simplify the analysis. Recent research trends tend to relax these assumptions and have contributed to the literature of *bounded rationality.* While models of bounded rationality abound, they are not equally promising in accounting for real observations. By Lucas' criterion

* This paper is an abbreviated version of Chen and Yeh (1998). Research support from NSC grant NSC. 86-2415-H-004-022 is gratefully acknowledged. The authors are grateful to David Fogel and one anonymous referee for painstaking reviews and many helpful suggestions. This paper is devoted to the memory of Mr. Paul Lin with Sun Fast International Corp., who had been a great supporter for our research for many years. To many people's grief, he died at the age of 38 on September 30, 1997 of liver cancer.

(Lucas, 1986), one of the most promising model classes is from a research group called *agent-based computational economics* (**ACE**) (Tesfatsion, 1996)[1], of which Arifovic (1995) is a typical example. In her studies, Arifovic applied *genetic algorithms* (**GAs**) to modeling the learning and adaptation in the **OLG**. She further compared the simulation results based on GAs with those from laboratories with human subjects, and she concluded that GAs were superior to other learning schemes, such as the *recursive least squares*[2].

Given the contribution of Arifovic (1995, 1996), our purpose is to move one step further, i.e., within the framework **ACE**, we attempt to use a more general version of **GAs** to model learning and adaptation in the **OLG**. The technique we use is *genetic programming* (**GP**). The significance of replacing **GAs** with **GP** in the economic context has been documented in Chen and Yeh (1996), but we would like to review it here in the specific context of the **OLG** model.

In many interesting **OLGs**, "expectations" refer to the expectations (forecasts) of *endogenous state variables* in the future. For example, in Sargent and Wallace (1982), the endogenous state variable is the *inflation rate*. Call these variables *expectations variables*. Then a model of bounded rationality should make expectations of these state variables *explicit*. However, to our best knowledge, in almost all applications of **GAs** to the **OLG**, this part is completely missing. Instead, it is *other endogenous variables* on which adaptive models are built. For example, in Arifovic (1995, 1996), it is *the demand for money and foreign assets*. There is nothing wrong with these applications; however, in terms of the distinction made by Marimon and Sunder (1994), what we learned from these studies is, at best, *learning how to optimize*, not *learning how to forecast*. If we want to know how agents' expectations evolve when the assumption of perfect foresight or rational expectations is relaxed, then the above works certainly fail to serve this purpose.

The only exception known to us is Bullard and Duffy (1994). In that paper, GAs were applied to modeling *the expectations of the inflation rate*. However, in their model what learning agents learn is a just a *number* of the inflation rate rather than *a regularity about the motion of the inflation rate*, which is a *function*. We consider it too restrictive to learn just a number. Therefore, in this paper, we would like to generalize Bullard and Duffy's evolution of "beliefs" from a sequence of populations of *numbers* to a sequence of populations of *functions*, and genetic programming serves as a convenient tool to make this extension.

2 The Overlapping Generations Model

2.1 The Main Model

Our model can be described as follows:

– It consists of overlapping generations of two-period-lived agents.

[1] There is a web page on **ACE** maintained by Leigh Tesfatsion. See http://www.econ.iastate.edu/tesfatsi/abe.htm.

[2] Laboratory studies on this issue can be found in Marimon and Sunder (1994).

- At time t, N young agents are born. Each of them lives for two periods $(t, t+1)$. At time t, each of them is endowed with e^1 units of a perishable consumption good, and with e^2 units at time $t+1$ $(e^1 > e^2 > 0)$.
- An agent born at time t consumes in both periods. c_t^1 is the consumption in the first period (t), and c_t^2 the second period $(t+1)$.
- All agents have identical preference given by[3]

$$U(c_t^1, c_t^2) = \ln(c_t^1 + 1) + \ln(c_t^2 + 1) \tag{1}$$

- In addition to the perishable consumption good, there is an asset called *money* circulated in the society. The nominal money supply at time t, denoted by H_t, is exogenously determined by the government and is held distributively by the old generation at time t. For convenience, we shall define h_t to be $\frac{H_t}{N}$, i.e., the nominal per capita money supply.

This simple **OLG** gives rise to the following *agent's maximization problem* at time t:

$$\max_{(c_{i,t}^1, c_{i,t}^2)} \quad \ln(c_{i,t}^1 + 1) + \ln(c_{i,t}^2 + 1) \tag{2}$$

$$s.t. \;\; c_{i,t}^1 + \frac{m_{i,t}}{P_t} = e^1, \quad c_{i,t}^2 = e^2 + \frac{m_{i,t}}{P_{t+1}}, \tag{3}$$

where $m_{i,t}$ represents the nominal money balances that agent i acquires at time period t and spends in the time period $t+1$, and P_t denotes the nominal price level at time period t. Since P_{t+1} is not available at period t, what agents actually can do is to maximize their *expected utility* $E(U(c_t^1, c_t^2))$ by regarding P_{t+1} as a random variable, where $E(.)$ is the expectation operator. Because of the special nature of the utility function and budget constraints, the first-order conditions for this *expected utility maximization problem* reduce to the certainty equivalence form (4):

$$c_{i,t}^1 = \frac{1}{2}(e^1 + e^2 \pi_{i,t+1}^e + \pi_{i,t+1}^e - 1) \tag{4}$$

where $\pi_{i,t+1}^e$ is agent i's expectation of the *inflation rate* $\pi_{t+1} (\equiv \frac{P_{t+1}}{P_t})$. This solution tells us the optimal decision of savings for agent i given her expectation of the inflation rate, $\pi_{i,t+1}^e$.

Suppose the government deficit is all financed through seignorage, then we can derive the dynamics (time series) of nominal price $\{P_t\}$ and inflation rate $\{\pi_t\}$ from Equation (4). To see this, let us denote the savings of agent i at time t by $s_{i,t}$. Clearly,

$$s_{i,t} = e^1 - c_{i,t} \tag{5}$$

By Equation (3), we know that

$$m_{i,t} = s_{i,t} P_t, \quad \forall i, t. \tag{6}$$

[3] The reason to add the constant here is to avoid evaluating $\ln(0)$, which can happen in the first period when agents choose to save all e^1.

In equilibrium, the nominal aggregate money demand must equal nominal money supply, i.e.,

$$\sum_{i=1}^{N} m_{i,t} = H_t = H_{t-1} + d_t P_t, \quad \forall t. \tag{7}$$

The second equality says that the money supply at period t is the sum of the money supply at period $t-1$ and the nominal deficit at period t, $d_t P_t$. This equality holds because we assume the government deficits are all financed by seignorage. Furthermore, let us assume that government spending is a fixed proportion ρ of aggregate savings and, for reasons clarified below, the government is assumed to have a constant revenue k, or simply

$$d_t = \rho \sum_{i=1}^{N} s_{i,t} - k. \tag{8}$$

Summarizing Equations (6)-(8), we get

$$\sum_{i=1}^{N} s_{i,t} P_t = \sum_{i=1}^{N} s_{i,t-1} P_{t-1} + P_t(\rho \sum_{i=1}^{N} s_{i,t} - k) \tag{9}$$

Hence, the price dynamics are governed by the following equation:

$$\pi_t = \frac{P_t}{P_{t-1}} = \frac{\sum_{i=1}^{N} s_{i,t-1}}{(1-\rho)\sum_{i=1}^{N} s_{i,t} + k} \tag{10}$$

Now suppose that each agent has perfect foresight, i.e.,

$$\pi_{i,t}^{e} = \pi_t, \quad \forall i, t. \tag{11}$$

Then by substituting the first-order condition (4) into Equation (9), we can have

$$(1-\rho)\frac{N}{2} P_t((e^1 - \pi_{t+1}e^2) + 1 - \pi_{t+1}) = \frac{N}{2} P_{t-1}((e^1 - \pi_t e^2) + 1 - \pi_t) - P_t k, \tag{12}$$

With Equation (12) rearranged, the paths of equilibrium inflation rates under perfect foresight dynamics are

$$(1-\rho)N((e^1 - \pi_{t+1}e^2) + 1 - \pi_{t+1}) = \frac{N}{\pi_t}((e^1 - \pi_t e^2) + 1 - \pi_t) - k. \tag{13}$$

At steady state ($\pi_{t+1} = \pi_t$), Equation (13) has two real stationary solutions (fixed points), a low-inflation stationary equilibrium, π_L^*, and a high-inflation one, π_H^*, given by

$$\pi_L^* = \frac{A - \sqrt{A^2 - 4(1-\rho)(1+e^1)(1+e^2)N^2}}{2(1-\rho)(1+e^2)N}, \tag{14}$$

$$\pi_H^* = \frac{A + \sqrt{A^2 - 4(1-\rho)N^2(1+e^1)(1+e^2)N^2}}{2(1-\rho)(1+e^2)N}, \tag{15}$$

where $A = (1+e^2)N + (1-\rho)(1+e^1)N + 2k$.

Table 1. Stationary Inflation Rates and Utilities under Different Values of k

k	0.1	1.0	10	20	50
π_L^*	1.2495	1.2450	1.2036	1.1630	1.0646
π_H^*	2.5010	2.5100	2.5964	2.6870	2.9354
U_L^*	2.4205	2.4217	2.4333	2.4456	2.4795
U_H^*	2.3026	2.3026	2.3029	2.3039	2.3090

For all cases of ks, $e^1 = 4$, $e^2 = 1$, and $\rho = 0.2$. U_L^* refers to the utility of agents in the low-inflation steady state, whereas U_H^* refers to the utility of agents in the high-inflation steady state.

2.2 Discussion

By Equation (10), the reason to add a parameter k in Equation (8) is quite clear. Without a positive k, it is possible that π_t can go to infinity if aggregate savings are 0. However, adding the constant k is harmless. First, k has a simple economic interpretation, i.e., the non-tax revenue. Second, if π_t can converge to either one of these two equilibria for all values of k, then we can well approximate the economy studied by Arifovic (1994) by choosing a sufficiently small k. For example, k is set to be 0.1 in this paper for most simulations.

2.3 Multiple Equilibria in the Model

The result of multiple equilibria, the existence of two stationary solutions, in this class of models is well known. These two stationary solutions differ not only in the inflation rate but also in the welfare implication. Agents' welfare under the high inflation rate π_H^* is inferior to that under the low inflation rate π_L^*, i.e., $U_H^* < U_L^*$. To see the difference, pairs of (π_L^*, U_L^*) and (π_H^*, U_H^*) are listed in Table 1 with respect to different values of k. Due to this difference, the steady state corresponding to the high inflation rate is called the *Pareto-inferior equilibrium*, and the steady state corresponding to the low inflation rate is called the *Pareto-superior equilibrium*. Given these equilibria with different welfare implications, will learning agents be able to pick up the good one rather than be trapped in the bad one? In this paper, we shall conduct three series of experiments to answer this question.

3 GP-based Agents in the OLG

This section provides a brief description of the way we apply genetic programming to modeling the expectations of the inflation rate in the **OLG** model. Let GP_t, a population of trees, represent a collection of agents' expectations of the inflation rate at time period t. The agent i born at time t, $i = 1, \ldots, N$, makes a decision about savings using the forecasting function, $gp_{i,t}$ ($gp_{i,t} \in GP_t$), a *parse tree* written over the *function set* and *terminal set* given in Table 2. In this

Table 2. Tableau of GP-Based Adaptation

Number of agents born in each period	250
Number of trees created by the full method	25 (Y), 25 (O)
Number of trees created by the grow method	25 (Y), 25 (O)
Function set	$\{+, -, \times, \%, Exp, Rlog, sin, cos\}$
Terminal set	$\{\pi_{t-1}, \pi_{t-2}, \cdots, \pi_{t-10}, R\}$
Number of trees created by reproduction	$p_r \times 250$
Number of trees created by crossover	$p_c \times 250$
Number of trees created by mutation	$p_m \times 250$
Probability of mutation	0.0033
Maximum depth of tree	17
Probability of leaf selection under crossover	0.5
Number of generations	1000
Maximum number in the domain of Exp	1700
Criterion of fitness	Utilities

"Y" stands for the initial young generation and "O" stands for the initial old generation. The number of trees created by full method or grow method is the number of trees initialized in Generation 0 in cases where the depth of tree is 2, 3, 4, 5, or 6. For details, see Koza (1992).

paper, all simulations conducted are based on the terminal set which includes the ephemeral random floating-point constant R ranging over the interval [-9.99, 9.99] and the inflation rate lagged up to 10 periods, i.e., $\pi_{t-1},\ldots,\pi_{t-10}$. Thus, the forecasting functions that agents may use are the linear and nonlinear functions of $\pi_{t-1},\ldots,\pi_{t-10}$.

The decoding of a parse tree $gp_{i,t}$ gives the forecasting function used by agent i at time period t, i.e., $\pi^e_{i,t+1}(\Omega_{t-1})$ where Ω_{t-1} is the information of the past inflation rates up to π_{t-1}. Evaluating $\pi^e_{i,t+1}(\Omega_{t-1})$ at the realization of Ω_{t-1} will give the inflation rate predicted by agent i at time period $t+1$, i.e., $\pi^e_{i,t+1}$. The *fitness* of a parse tree $gp_{i,t}$ is determined by the value of the agent's utilities gained at the end of her life based on Equation (1), i.e., $U_{i,t} = U(c^1_{i,t}, c^2_{i,t})$.

Each fitness value $U_{i,t}$ is then normalized. The *normalized fitness* value $\lambda_{i,t}$ is given in Equation (16).

$$\lambda_{i,t} = \frac{U_{i,t}}{\sum_{i=1}^{N} U_{i,t}} \tag{16}$$

It is clear that the normalized fitness is a *probability measure*. Moreover, $\lambda_{i,t}$ is greater for a better parse tree $gp_{i,t}$. Once $\lambda_{i,t}$ is determined, GP_{t+2} is generated from GP_t by three primary genetic operators, i.e., *reproduction, crossover*, and *mutation*. The implementation of these three operators is detailed in Chen and Yeh (1998).

4 Experimental Designs

The parameters of most of the **OLGs** simulated in this paper are based on the setup: $e^1 = 4, e^2 = 1, \rho = 0.2, k = 0.1$. By this set of parameters, $\pi_L^* = 1.2495$, and $\pi_H^* = 2.5010$. These values are also shown in the second column of Table 1.

To see whether or not GP-based agents can coordinate to converge to the Pareto-superior equilibrium, the **OLG** is simulated by feeding it with four sets of initial values of $\{\pi_{-1}, ..., \pi_{-10}\}$. These four sets of initial values are chosen so that $\{\pi_{-1}, ..., \pi_{-10}\}$ are randomly distributed over the ranges (1.10, 1.20), (1.25, 1.35), (3.30, 4.30) and (4.50, 5.50). Clearly, the first two sets, Sets 1 and 2, are chosen to be the neighborhoods of π_L^*; Set 1 is below π_L^*, and Set 2 above it. Sets 3 and 4 are chosen far higher than π_H^*. This design enables us to check both the local and global stability of π_L^*. We number the experiments corresponding to these four different sets of initial values as Experiments 1, 2, 3 and 4.

A related test for the global stability is to conduct a perturbation test. We first run the **OLG** model under the original chosen parameters. If it converges to π_L^*, we shall perturb π_L^* by changing the values of some parameters, and see whether or not the new π_L^* will be selected again. In this paper, we consider the case in which k is changed from 0.1 to 10. This perturbation test shall be numbered as Experiment 5.

Experiments 1-5 are designed to test the global stability of π_L^*. However, a typical question frequently raised is whether or not these results are sensitive to the genetic operators used. To answer this question, we consider four sets of (p_r, p_c, p_m). For Experiments 1-5, $p_r = 0.12, p_c = 0.68$, and $p_m = 0.20$. We then consider the significance of each genetic operator in Experiments 6-8. In Experiment 6, only reproduction is used, i.e., $p_r = 1, p_c = 0$, and $p_m = 0$. Similarly, $p_c = 1$ in Experiment 7, and $p_m = 1$ in Experiment 8.

5 Simulations Results

For each design, five runs were implemented. Since results are quite similar among the five simulations for each design, we only report one of the results for each design here and the full details can be found in Chen and Yeh (1998). The basic statistics of each simulation are summarized in Table 3 and the plot of the whole time series of π_t is exhibited in Figures 1-8.

From Table 3, we can make the following conclusions.

- GP-based agents are able to coordinate with each other to converge to the low-inflationary stationary equilibrium. The evidence shows that in all the simulations, except the one with structural change, π_t converges to a small neighborhood of π_L^* ($\overline{\pi} = 1.2495$).
  - As a corollary, the evidence also shows that the convergence to π_L^* is insensitive to the initial condition. The initial rates of inflation in Experiments 3 and 4 are quite far away from π_L^* and are closer to π_H^*. However, this does not make π_t converge to π_H^*, and π_H^*, being one of

Table 3. Results of Experiments 1-8:

Experiment	$\overline{\pi}$	δ_π	$\overline{U}$	δ_U
1	1.2495	0.0018	2.4205	0.0006
2	1.2495	0.0020	2.4204	0.0008
3	1.2495	0.0021	2.4204	0.0008
4	1.2495	0.0018	2.4205	0.0006
5	1.2036	0.0047	2.4332	0.0012
6	1.2495	0.0000	2.4205	0.0000
7	1.2495	0.0000	2.4205	0.0000
8	1.2495	0.0093	2.4201	0.0033

$\overline{\pi}$ = the mean inflation rate of a simulation (from Generation 501 to 1000).
δ_π = the standard deviation of π_t of a simulation (from Generation 501 to 1000).
$\overline{U}$ = the mean welfare of a simulation (from Generation 501 to 1000), where
$U_t = \frac{\sum_i U_{i,t}}{250}$.
δ_U = the standard deviation of U_t of a simulation (from Generation 501 to 1000).

the stationary equilibria in the perfect-foresight setup, can hardly be reached in this ACE setup. **Bounds on rationality do change equilibria in economic systems.** In our case, the equilibrium with high inflation and low utilities is eliminated.

- Furthermore, GP-based agents are capable of converging to the new low-inflationary stationary equilibrium after the perturbation. In Experiment 5, $\overline{\pi} = 1.2036$, which is exactly the π_L^* under $k = 10$. From Figure 5, we can also see that the transition speed from the old equilibrium to the new one is very fast.

– Whether or not π_t will converge to a niche of π_L^* does not depend on the choice of the pair (p_r, p_c, p_m) $(p_r + p_c + p_m = 1)$, as can be seen from Experiments 6-8. Using only one of these genetic operators is sufficient to achieve the same result (See Table 3). Nevertheless, there is a difference between the reproduction and crossover operators and the mutation operator. From Figures 6-7 or the corresponding δ_π in Table 3, we can see that if only the reproduction or crossover operator is employed, then the convergence to π_L^* is *strict* in the sense that $\pi_t = \pi_L^*$ $\forall t$ as t is large enough. But, this strict convergence result disappears when only the mutation operator is applied. This can also be seen from Figure 8 and the corresponding δ_π in Table 3. In fact, setting $p_m = 1$ results in the highest value of δ_π among other setups. Therefore, the fluctuations of π_t observed in most of the experiments are due to the mutation operator.

6 Concluding Remarks

In this paper, we provide a concrete example to demonstrate how genetic programming can be applied to modeling learning and expectations in the **OLG**. Our simulations indicate that the main feature observed in the laboratory with human subjects, namely, agents being able to coordinate their actions to achieve the Pareto-superior equilibrium, can be replicated by these GP-based agents. The agent-based approach suggested here is more general than those used in the earlier studies and may be considered as a basis for studying other **OLGs** where learning and adaptation play a crucial role for the determination of equilibrium.

References

1. Allais, M. (1947), "Economie et Interet," Imprimerie Nationale, Paris.
2. Arifovic, J. (1995), "Genetic Algorithms and Inflationary Economies," *Journal of Monetary Economies*, 36, pp. 219-243.
3. Arifovic, J. (1996), "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies," *Journal of Political Economy*, Vol. 104, No. 3, pp. 510-541.
4. Bullard, J. and J. Duffy (1994), "Using Genetic Algorithms to Model the Evolution of Heterogeneous Beliefs," mimeo, Federal Reserve Bank of St. Louis and University of Pittsburgh.
5. Chen, S.-H. and C.-H. Yeh (1996), "Genetic Programming Learning in the Cobweb Model," in P. J. Angeline and K. E. Kinnear (eds.), *Advances in Genetic Programming*, Vol. II, MIT Press. pp. 443-466.
6. Chen, S.-H. and C.-H. Yeh (1998), "Modeling the Expectations of Inflation in the OLG model with Genetic Programming," *AI-ECON Research Group Working Paper Series No. 9801*, Department of Economics, National Chengchi University.
7. Koza, J. R. (1992), *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Cambridge: MIT Press.
8. Lucas, R. E., Jr., (1986), "Adaptive Behavior and Economic Theory," *Journal of Business*, 59, pp. 401-426.
9. Marimon, R. and S. Sunder (1994), "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach," *Economic Theory*, 4, pp. 131-162.
10. Samuelson, P. A. (1958), "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, Vol. 66, No. 6, pp. 467-482.
11. Tesfatsion, L. (1996), "How Economists Can Get Alife," in B, Arthur, S. Durlauf and D. Lane (eds.), *The Economy as an Evolving Complex Systems, II,* Santa Fe Institute in the Science of Complexity, Vol. XXVII, Addison-Wesley, Reading, MA. pp. 533-564.

Figure 1 : Equilibrium Inflation Rate in Each Generation

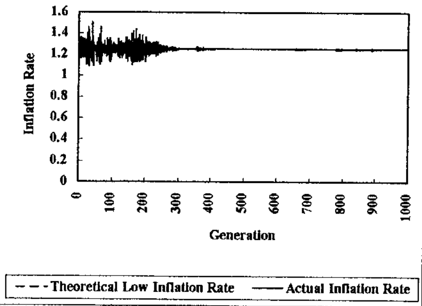

Figure 2 : Equilibrium Inflation Rate in Each Generation

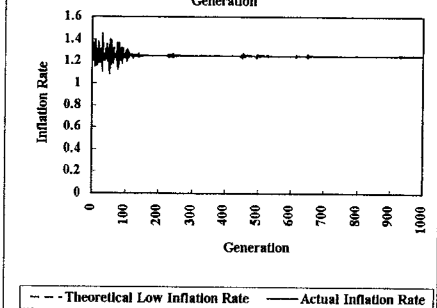

Figure 3 : Equilibrium Inflation Rate in Each Generation

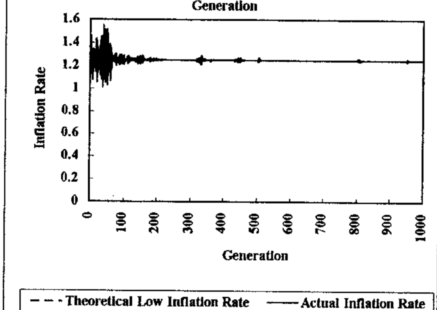

Figure 4 : Equilibrium Inflation Rate in Each Generation

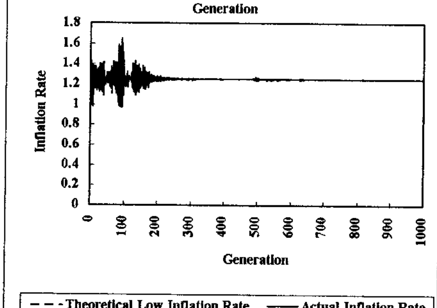

Figure 5 : Equilibrium Inflation Rate in Each Generation

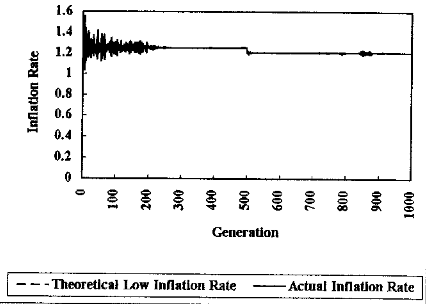

Figure 6 : Equilibrium Inflation Rate in Each Generation

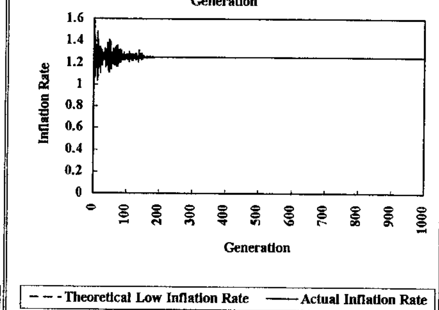

Figure 7 : Equilibrium Inflation Rate in Each Generation

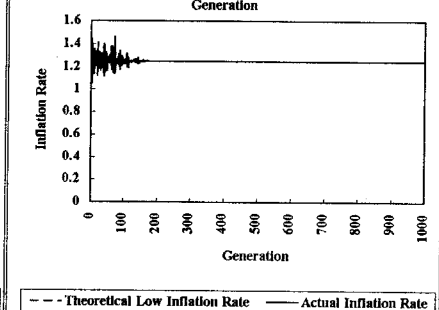

Figure 8 : Equilibrium Inflation Rate in Each Generation