Original paper

Soft Computing 3 (1999) 53–62 © Springer-Verlag 1999

Modeling the expectations of inflation in the OLG model with genetic programming

S.H. Chen, C.H. Yeh

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, GP control parameters, and the selection mechanism. We find that as long as the survival-of-the-fittest principle is maintained, the evolutionary operators are only secondarily important. However, once the survival-of-the-fittest principle is absent, the well-coordinated economy is also gone and the inflation rate can jump quite wildly. To some extent, these results shed light on the biological foundations of economics.

Key words Genetic programming, overlapping generations models, bounded rationality, agent-based computational economics, Pareto-superior equilibrium

1 Introduction

While there are several approaches to introducing *dynamic general equilibrium structures* to economics, the *overlapping generations model* (hereafter, OLG), proposed by Allais [1] and Samuelson [19] and developed by Diamond [10], Shell [21], Lucas [13], and Gale [11], may be regarded as the most popular in current macroeconomics. Unlike *infinite-horizon models*, the OLG does not assume that agents can live forever. But, before they all pass away, the new generations are born, which overlap with the existing generations for a number of periods. The essence of the OLG is that living individuals are prevented from trading with the unborn or the dead. This feature of the OLG enables us to address all kinds of macroeconomic issues with an explicit reference to the *demographic structure*, which is certainly a key element in a real economy. Over the last two decades, the OLG model has been extensively applied to studies of savings, bequests, demanding for assets, prices of assets, inflation, business cycles, economic growth, and the effects of taxes, social security and budget deficits.

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 1980's the assumptions of *perfect foresight* and *rational expectations* were adopted to simplify the analysis. Recent research trends tend to relax these assumptions 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 [16], one of the most promising model classes is from a research group called *agent-based computational economics* (ACE) [22], of which Arifovic [2] 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 [17,18] and she concluded that GAs were superior to other learning schemes, such as the *recursive least squares*.

Given the contribution of Arifovic [2, 3] 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 [9], 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 variable* in the future. For example, in [20], the endogenous state variable is the *inflation rate*; in [13], it is the *exchange rate*, and in [5], the *quality of labor*. 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 [2,3], it is *the demand for money and foreign assets*; in [4], it is the *time spent on training*. There is nothing wrong with these application however, in terms of the distinction made by

S.-H. Chen ✉[1], C.-H. Yeh[2]
AI-ECON Research Group, Department of Economics, National Chengchi University, Taipei 11623 Taiwan
[1]E-mail: chchen@nccu.edu.tw
[2]E-mail: g3258501@grad.cc.nccu.edu.tw

Research support from NSC grant NSC. 86-2415-H-004-022 is gratefully acknowledged. An early version of this paper has been presented in the Evolutionary Programming 1998 Conference. We are grateful to D. Fogel, L. Tesfatsion, and two anonymous referees for their helpful suggestions.

Marimon and Sunder [18], 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 [7]. 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. By Grandmont [12], if the equilibrium of an OLG is characterized by *limit cycles* or *strange attractors* rather than by fixed points, then what agents need to learn is not just a number but a functional relationship, such as $x_t = f(x_{t-1}, x_{t-2}, \ldots)$. Moreover, in this situation, if agents are restricted to learning only a number, then it is likely that the learning process will fail to converge to any Pareto optimum. 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.

To be comparable with the existing literature we employ the inflation model studied in [2]. Briefly speaking, this model has two stationary equilibria associated with two inflation rates. The welfare implication of these two equilibria is also different. The inflation rates observed in the experiment with human subjects converged to the low-inflation stationary equilibrium. We want to check whether GP-based agents can replicate this result and the robustness of this replication.

The rest of this paper is organized as follows. Sect. 2 gives an agent-based version of a simple OLG of a monetary economy. Sect. 3 illustrates how to use GP to model the learning and adaptation in this model. Experimental designs are described in Sect. 4, followed by the analysis of simulation results in Sect. 5 and suggestions of some further explorations in Sect. 6. Concluding remarks are given in Sect. 7.

2 The overlapping generations model

2.1 The main model

The analytical model employed in this paper is based on a very simple OLG of a monetary economy. This model was first introduced by Allais [1] and Samuelson [19], and later on was used extensively to study issues of *inflationary finance*, e.g., Bryant and Wallace [6]. Arifovic [2] and Bullard and Duffy [7, 8] applied GAs to extending this simple model into an agent-based version and simulated plausible processes of adapting to out-of-equilibrium behaviour. The model presented in this paper is very similar to Arifovic [2] and Bullard and Duffy [7,8]. However, in order to facilitate computer simulation, we also make some slight modifications, and we will come back to this point later.

Our model can be described as follows:

- It consists of overlapping generations of two-period-lived agents.
- 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 $t+1$ $(e^1 > e^2 > 0)$. Presumably e^1 is assumed to be greater than e^2 in order to increase the likelihood (not ensure) that agents will choose to hold money from period 1 to 2 to push value forward.
- 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

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

  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 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 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:

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

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

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 $E(\cdot)$ 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):

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

where $\pi_{i,t+1}^e$ is agent i's expectation of the *inflation rate* $\pi_{t+1} (\equiv P_{t+1}/P_t)$. 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 Eq. (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 Eq. (3), we know that

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

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 the 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 Eqs. (6)–(8), we get

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

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.,

$$\pi^e_{i,t} = \pi_t, \quad \forall i, t. \tag{11}$$

Then by substituting the first-order condition (4) into Eq. (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 Eq. (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 ($\pi_{t+1} = \pi_t$), Eq. (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 $A = (1+e^2)N + (1-\rho)(1+e^1)N + 2k$.

2.2 Discussions

Before proceeding further, we would like to make comments on parameters ρ and k appearing in the model presented above. Instead of the *nominal deficit*, this paper chooses the *deficit ratio* ($\rho \equiv$ deficits/aggregate savings) as a policy parameter. This choice distinguishes this paper from Arifovic [2] and Bullard and Duffy [7, 8], where the nominal deficit is chosen to be the policy parameter. The reason for this different choice is two-fold. Firstly, from the perspective of computation efficiency, using the deficit as the policy parameter may cause the run-off of the simulation when the deficit is greater than aggregate savings. When this happens, both Arifovic [2] and Bullard and Duffy [7,8] forced the algorithm to be reinitialized and the simulation to begin anew. We consider this design quite inefficient, particularly, when the chance of reinitialization can be even higher in using GP. This is because GP is searching for a larger space than the GA used by Arifovic, Bullard and Duffy. Secondly, as far as economic interpretation is concerned, these two different parameters make no difference given that the economy can converge to the same equilibrium, and as we shall see later, this is indeed the case. Moreover, in terms of policy making, the *deficit ratio* can be more practical and more informative. For example, one of the theories to explain the recent crisis of Thai Baha is that Thai's deficit ratio is higher than 8%. Due to these considerations, the deficit ratio ρ is chosen to be the policy parameter in this paper.

Also, by Eq. (10), the reason to add a parameter k in Eq. (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) and Bullard and Duffy [7] 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? Put in a more general way, are decentralized agents able to coordinate intelligently to single out the best outcome? Issues like this have motivated Arifovic [2], Bullard and Duffy [7,9], and this paper.

In those earlier studies, GAs-based agents are shown to be able to select the *Pareto-superior equilibrium*. The only difference between Arifovic [2] and Bullard and Duffy [7], is the speed of convergence. As Bullard and Duffy [2] asserted, "Our initial impression is that the *learning how to forecast* version of genetic algorithm learning converges faster than the *learning how to optimize* implementation studied by Arifovic"

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

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	$\{\pi_{t-1}, \pi_{t-2}, \dots, \pi_{t-10}, R\}$
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 [14]

(p. 21). In this paper, we shall conduct three series of experiments to answer the following questions:

1. Are GP-based agents able to coordinate well enough to select the Pareto-superior equilibrium?
2. Is the result obtained robust?

Before discussing the design of experiments, let us first illustrate how to modeling agents' adaptive expectation in the OLG with GP.

3 GP-based Agents in the OLG

3.1 Coding and decoding of expectations

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, $i=1, \dots, N$, 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 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., $\pi_{t-1}, \dots, \pi_{t-10}$. Thus, the forecasting functions that agents may use are the linear and nonlinear functions of $\pi_{t-1}, \dots, \pi_{t-10}$. Clearly, this specification of terminal set limits the GP to fining forecasting rules that are functions only of a fixed finite number of lagged inflation terms. But, extending this terminal set by allowing agents to make use of current and past values of the money supply is straightforward and is left for future study.

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 Eq. (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 Eq. (16).

$$\lambda_{i,t}=\frac{U_{i,t}}{\sum_{i=1}^N U_{i,t}}. \qquad (16)$$

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*. These three operators are described below.

3.2 Genetic updating of expectations

(1) *Reproduction*: Reproduction makes copies of individual parse trees. The criterion used in copying is the normalized fitness value $\lambda_{i,t}$. If $gp_{i,t}$ is an individual in the population GP_t with normalized fitness value $\lambda_{i,t}$, then in each run of the reproduction operator, $gp_{i,t}$ will be copied into the GP_{t+2} with probability $\lambda_{i,t}$. Reproduction is performed on only a specified share, p_r, of the population. The rest of the offspring are generated by the other two operators, *crossover* and *mutation*.

(2) *Crossover*: The crossover operation for the genetic programming paradigm is a sexual operation that starts with two parental parse trees which are randomly selected from population GP_t in accordance with the normalized fitness described above. Next, by exchanging the parts of these parents, two off spring are produced. This exchange begins by randomly and independently selecting one point in each parental parse tree using a uniform distribution described below.
By the syntax of LISP, each point (atom) of a parse tree could be either for that matter, a *leaf* (terminal) or an *inner code* (function). Thus, the point (atom) selected could either be a terminal or a function. As specified in Table 2, the probability of the crossover point being a terminal, or a function is 50–50. Given that a terminal or function is to

be the point chosen for crossover, the probability of any terminal or function being chosen as the crossover point is uniformly distributed. For example, if the crossover point is to be a terminal, and there are three terminals in the parse tree, the probability of any one of the three terminals being chosen for the crossover point is one-third (1/3). $(100p_c)\%$ of the new generation is created in this way.

(3) *Mutation*: The operation of mutation also allows new individuals to be created. It begins by selecting a parse tree $gp_{i,t}$ from the population GP_t based on $\lambda_{i,t}$. Once a particular $gp_{i,t}$ is selected, mutation is a process of a random change of the value of a point (atom) within $gp_{i,t}$. Each point (atom) has a small probability of being altered by mutation, which is independent of other point (atoms). As specified in Table 2, the probability used throughout this paper is 0.0033. To be a syntactically and semantically valid LISP S-expression, terminals can only be altered by the member from the terminal set and functions can only be altered by the member with the same number of arguments from the function set. The altered individual is then copied into GP_{t+2}. $(100p_m)\%$ of the new generation is created in this way.

3.3 Discussion

The genetic updating procedure described above is quite standard except that newborn agents (GP_t) inherit ideas from their grandparents (GP_{t-2}). This special feature, called the *non-overlapping information structure* in [8], is due to the fact that parent's fitness values are not available when newborn agents are young. Therefore, in the OLG, learning occurs between *non-overlapping* generations of grandparents to grandchildren.

In the 2-period OLG, this *non-overlapping* genetic updating procedure may not be too troublesome. But, as noticed by Sargent (1993, p. 102), "When we extend the horizon beyond two periods, it becomes increasingly inconvenient to model learning in this way because we have to wait longer for the consequences of life-time savings behaviour to be known". Nevertheless, the long waiting problem that Sargent refers to is not difficult to solve. To do so, simply replace the *utility function* by some *statistical loss functions*, such as the *sum of square errors*. Unlike the utility function, which can be evaluated only at the end of agent's life statistical, loss function can be evaluated at each period of the agent's life span, and this information can be immediately used to form next generation's expectations. However, the premise to follow this line is to model agents as *forecasters* not *optimizers*. Since most GA applications fail to address the forecasting aspect of agents' behaviour, it is not surprising that this line has not been taken. Interestingly enough, Bullard and Duffy [8] addressed this problem by augmenting the genetic algorithm with an *emulation* procedure. This procedure, to some extent, attempts to incorporate the forecasting aspect. However, since the forecasting aspect is not directly coded in their GAs, this procedure is not adaptive and, in fact, is limited to only three prespecified forecasting functions.

While, at this stage, we only restrict our attention to the 2-period OLG and still take the utility function as the fitness function, we would like to modify our GP in the direction suggested above when we extend our simulations to the N-period OLG in future research.

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 $\{\pi_{-1}, \ldots, \pi_{-10}\}$. These four sets of initial values are chosen so that $\{\pi_{-1}, \ldots, \pi_{-10}\}$ 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 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 1–4.

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 leave the full details in the appendix (Tables 5–7). 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 Fig. 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^* ($\bar{\pi}=1.2495$).

– 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 the stationary equilibria in the perfect foresight setup, can hardly be reached in this ACE setup. *Bounds on rationality do change*

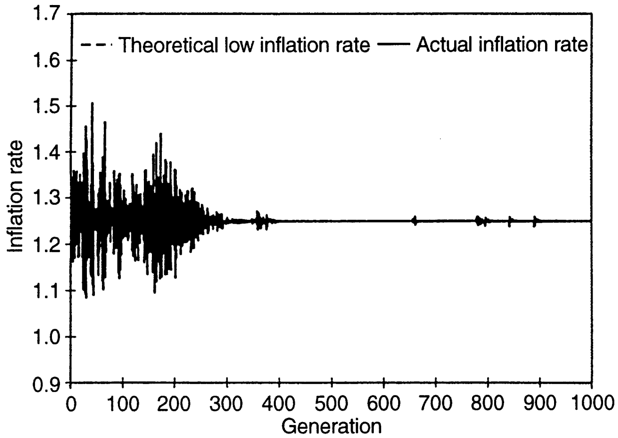

Fig. 1. Equilibrium inflation rate in each generation

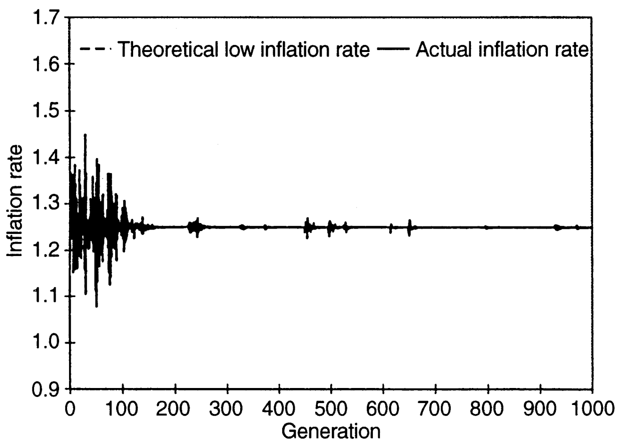

Fig. 2. Equilibrium inflation rate in each generation

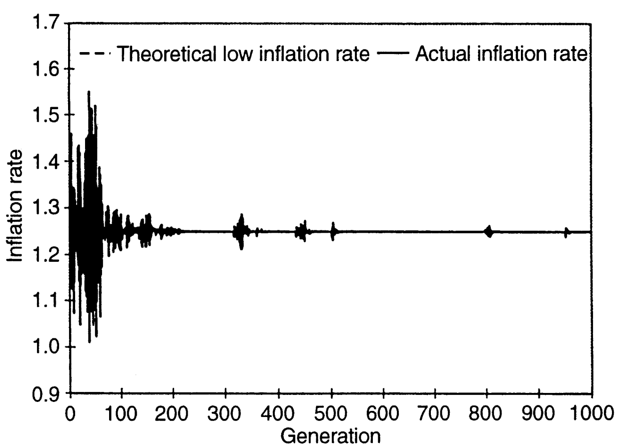

Fig. 3. Equilibrium inflation rate in each generation

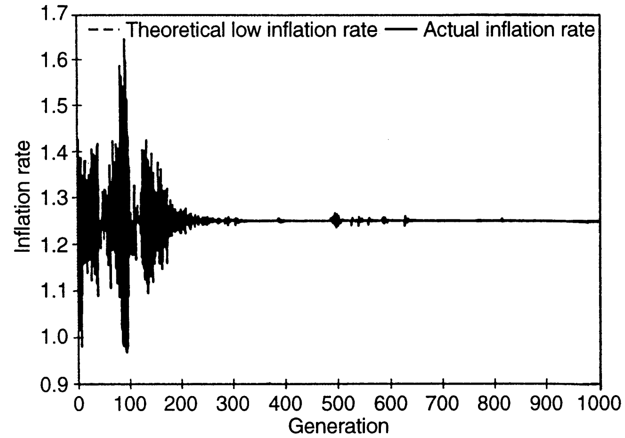

Fig. 4. Equilibrium inflation rate in each generation

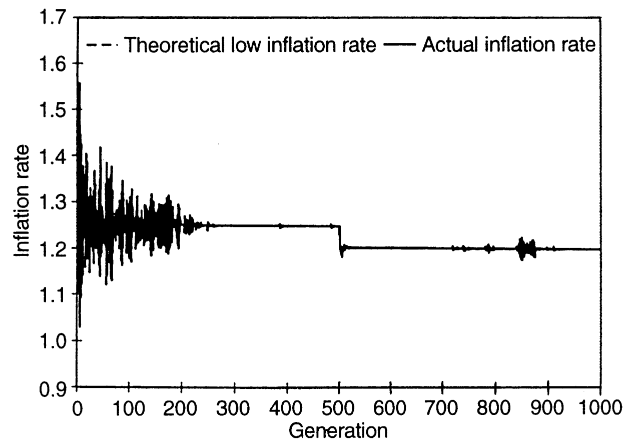

Fig. 5. Equilibrium inflation rate in each generation

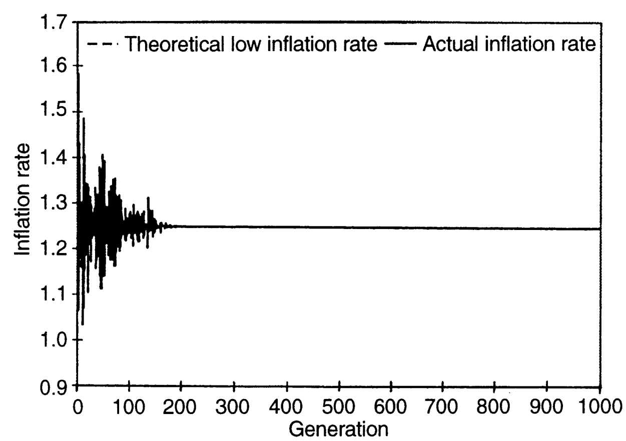

Fig. 6. Equilibrium inflation rate in each generation

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, $\bar{\pi} = 1.2036$, which is exactly the π_L^* under $k = 10$. From Fig. 5, we can also see that the transition speed from the old equilibrium to the new one is very fast.

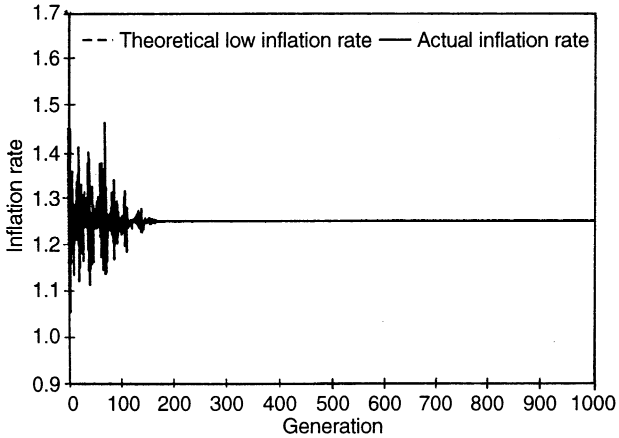

Fig. 7. Equilibrium inflation rate in each generation

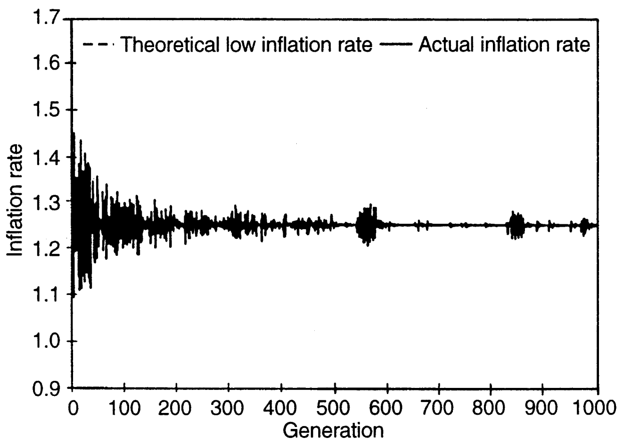

Fig. 8. Equilibrium inflation rate in each generation

- 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 Fig. 6 and 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 Fig. 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.

  This result was also noticed by Arifovic [2], who introduced an operator, called *election*, to offset the effect of mutation on population diversity and the resulting fluctuations. Nevertheless, unlike us, she did not go further to separate the effect of each operator and to check whether it is necessary to introduce additional operator to obtain the convergence result.

6
Further exploration

That the GP-based agents were remarkably able to select the Pareto-superior equilibrium inevitably leads to the question: where does the "magic" come from? Is it easy to get this convergence result? To solve this puzzle, we conducted more simulations using the same GP-based agents described above with the exception that the *proportionate selection mechanism* was replaced by *uniform selection* (selection with the uniform distribution). The purpose of doing so is to see the role of *the survival-of-the-fittest principle* implemented in these simulations, we reran all the simulations in the same order, and the results are presented in Table 4 and Fig. 9–16.

From Fig. 9–16, we can see that π_t in all the simulations fluctuates quite dramatically. In many cases, we see the appearance of *super inflation*, which is not observed in Experiments 1–8. From Table 4, we can also see that $\bar{\pi}$s are generally biased upwards. Furthermore, the agents are worse off in that they have lower utility. In sum, the GP-based agents with uniform selection perform much worse than those with proportionate selection in terms of converging to the Pareto-superior equilibrium.

Table 3. Results of Experiments 1–8

Experiment	$\bar{\pi}$	δ_π	$\bar{U}$	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

$\bar{\pi}$ = 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)
$\bar{U}$ = the mean welfare of a simulation (from Generation 501 to 1000), where $U_t=\sum_i U_{i,t}/250$
δ_U = the standard deviation of U_t of a simulation (from Generation 501 to 1000)

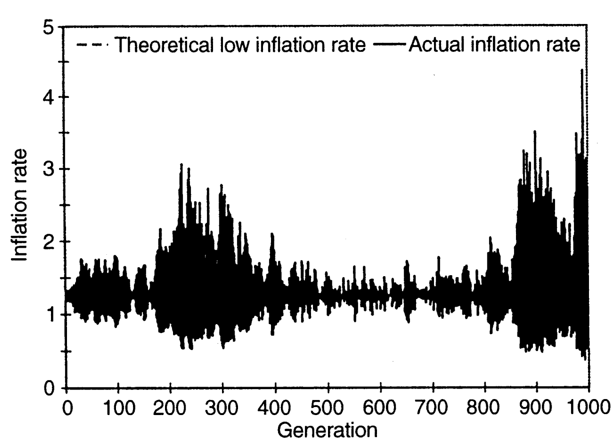

Fig. 9. Equilibrium inflation rate in each generation

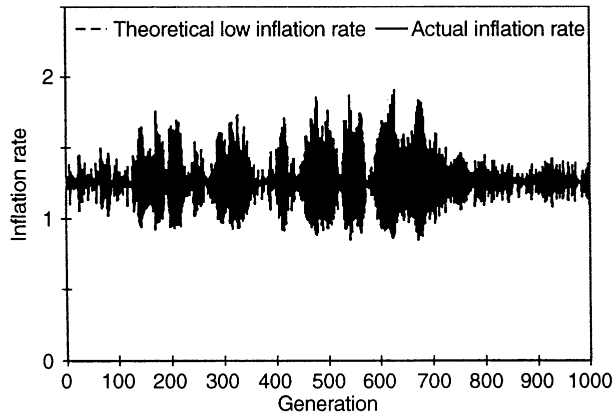

Fig. 10. Equilibrium inflation rate in each generation

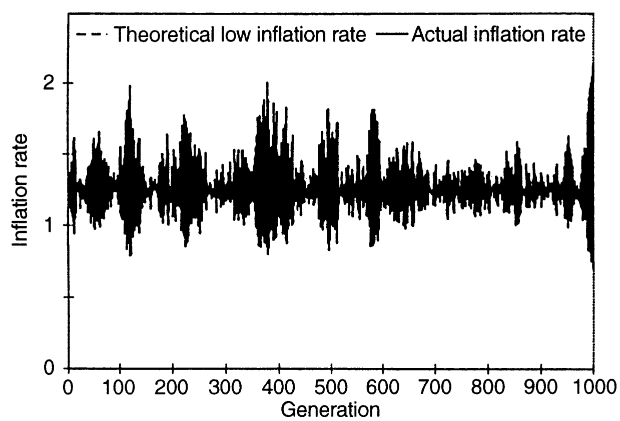

Fig. 11. Equilibrium inflation rate in each generation

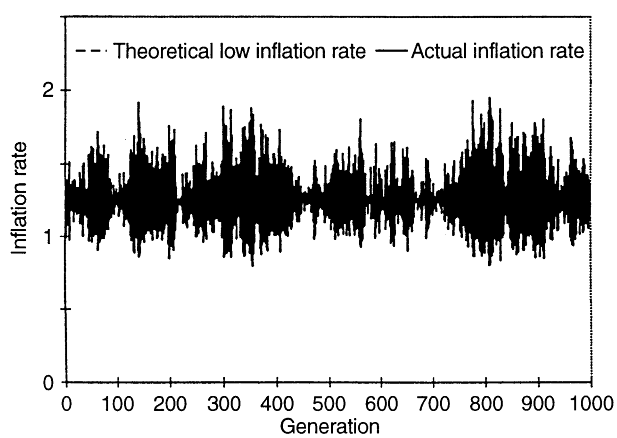

Fig. 12. Equilibrium inflation rate in each generation

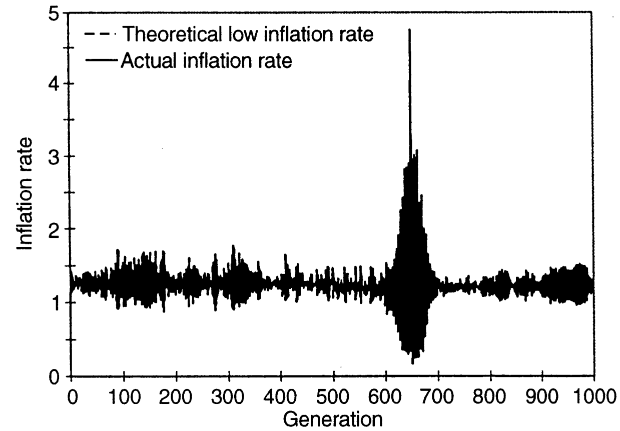

Fig. 13. Equilibrium inflation rate in each generation

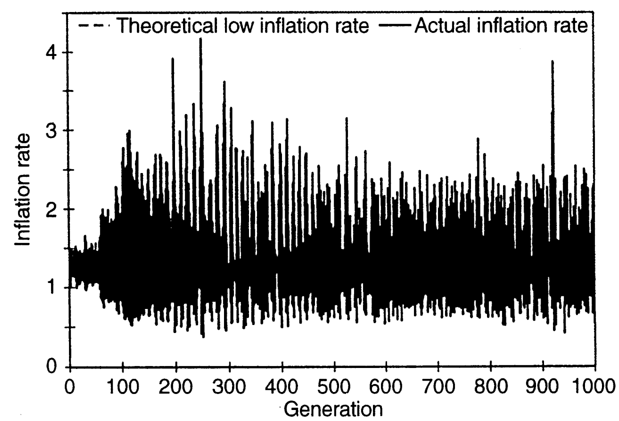

Fig. 14. Equilibrium inflation rate in each generation

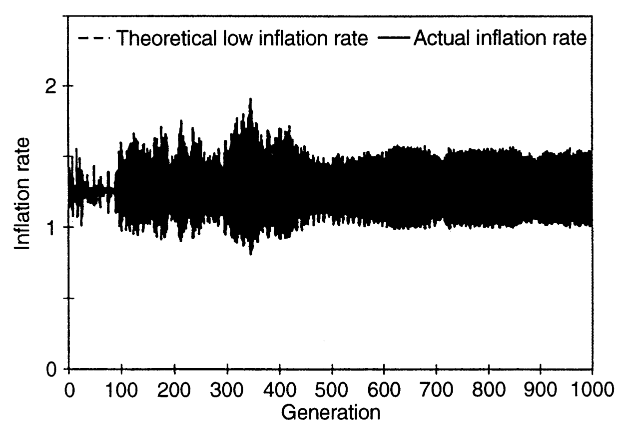

Fig. 15. Equilibrium inflation rate in each generation

From the perspective of *optimization*, such outcomes are not surprising, but from the perspective of *simulating social phenomena*, these outcomes are important in that they show *how human intelligence can be incorporated into the design of artificial societies*. While we do not know how human subjects in the laboratory *actually* learn, there is little doubt that converging to the *Pareto-superior equilibrium* is an *intelligent outcome*. In other words, this result cannot be replicated

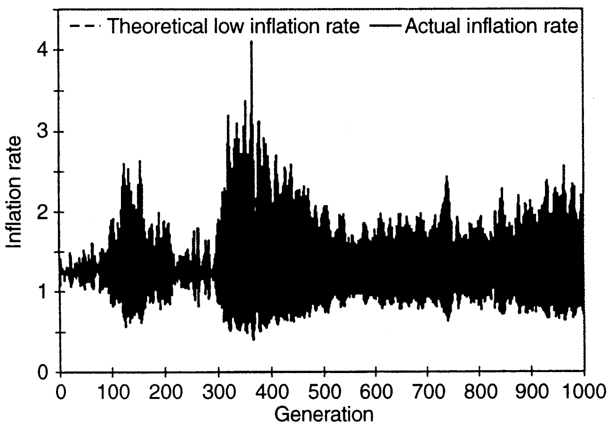

Fig. 16. Equilibrium inflation rate in each generation

Table 4. Results of Experiments 9–16

Experiment	$\bar{\pi}$	δ_π	$\bar{U}$	δ_U
9	1.3640	0.6024	2.3214	0.0817
10	1.2672	0.2148	2.3878	0.0780
11	1.2621	0.1883	2.3495	0.0704
12	1.2738	0.2509	2.2303	0.0778
13	1.2542	0.3791	2.3778	0.1048
14	1.3954	0.6427	2.3028	0.3348
15	1.2739	0.2487	2.4057	0.0895
16	1.3592	0.5365	2.2934	0.1727

$\bar{\pi}$ = 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)
$\bar{U}$ = the mean welfare of a simulation (from Generation 501 to 1000), where $U_t = \sum_i U_{i,t}/250$
δ_U = the standard deviation of the U_t of a simulation (from Generation 501 to 1000)

Table 5. Simulation results: Experiments 1–8

Experiment	$\bar{\pi}$	δ_π	$\bar{U}$	δ_U
Sim 1–1	1.2495	0.0018	2.4205	0.0006
Sim 1–2	1.2495	0.0014	2.4205	0.0006
Sim 1–3	1.2495	0.0013	2.4205	0.0005
Sim 1–4	1.2495	0.0019	2.4205	0.0005
Sim 1–5	1.2495	0.0049	2.4204	0.0013
Sim 2–1	1.2495	0.0020	2.4204	0.0008
Sim 2–2	1.2495	0.0031	2.4204	0.0010
Sim 2–3	1.2495	0.0031	2.4204	0.0013
Sim 2–4	1.2495	0.0033	2.4204	0.0009
Sim 2–5	1.2502	0.0417	2.4178	0.0144
Sim 3–1	1.2495	0.0021	2.4204	0.0008
Sim 3–2	1.2495	0.0071	2.4204	0.0020
Sim 3–3	1.2495	0.0017	2.4205	0.0007
Sim 3–4	1.2498	0.0209	2.4198	0.0064
Sim 3–5	1.2495	0.0071	2.4202	0.0016
Sim 4–1	1.2495	0.0018	2.4205	0.0006
Sim 4–2	1.2495	0.0046	2.4204	0.0017
Sim 4–3	1.2495	0.0023	2.4204	0.0009
Sim 4–4	1.2495	0.0052	2.4203	0.0025
Sim 4–5	1.2495	0.0015	2.4205	0.0007
Sim 5–1	1.2036	0.0047	2.4332	0.0012
Sim 5–2	1.2036	0.0028	2.4333	0.0008
Sim 5–3	1.2041	0.0131	2.4328	0.0044
Sim 5–4	1.2036	0.0063	2.4332	0.0021
Sim 5–5	1.2036	0.0049	2.4332	0.0018
Sim 6–1	1.2495	0.0000	2.4205	0.0000
Sim 6–2	1.2495	0.0000	2.4205	0.0000
Sim 6–3	1.2495	0.0000	2.4205	0.0000
Sim 6–4	1.2493	0.0237	2.4205	0.0064
Sim 6–5	1.2495	0.0000	2.4205	0.0000
Sim 7–1	1.2495	0.0000	2.4205	0.0000
Sim 7–2	1.2495	0.0071	2.4204	0.0020
Sim 7–3	1.2495	0.0000	2.4205	0.0000
Sim 7–4	1.2495	0.0000	2.4205	0.0000
Sim 7–5	1.2495	0.0000	2.4205	0.0000
Sim 8–1	1.2495	0.0093	2.4201	0.0033
Sim 8–2	1.2496	0.0084	2.4201	0.0030
Sim 8–3	1.2495	0.0056	2.4202	0.0020
Sim 8–4	1.2496	0.0105	2.4199	0.0030
Sim 8–5	1.2497	0.0116	2.4196	0.0037

without any constraints. Of course, these artificial constraints may have nothing to do with the constraints which make humans behave "intelligently", but they enable us to form hypotheses about what makes humans collectively behave intelligently. In this paper, the contributing factor is found to be *the survival-of-the-fittest principle* which is explicitly implemented through *proportionate selection*. As long as this principle is kept, the importance of the combination of the use of genetic operators is only secondary.

7
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.

Appendix
Softwares

The program to implement the simulations in this paper can be downloaded directly from the website:
http://econo.nccu.edu.tw/ai/staff/csh/Software.htm

Simulations results

In this appendix, the basic statistics of all simulations are reported in Tables 5 and 6. Since five runs are implemented for each design, we code them by "$X-x$" where $X = 1, 2, \ldots, 16$ and $x = 1, 2, \ldots, 5$. For example, "3–5" indicates the fifth run of the third experiment design. In Table 7, we summarize the five runs of each design by averaging all statistics reported in Tables 5 and 6.

Table 6. Simulation results: Experiments 9–16

Experiment	$\bar{\pi}$	δ_π	$\bar{U}$	δ_U
Sim 9–1	1.3640	0.6024	2.3214	0.0817
Sim 9–2	1.3604	0.5808	2.2913	0.1721
Sim 9–3	1.2748	0.2593	2.3885	0.0944
Sim 9–4	1.2730	0.2421	2.3711	0.0682
Sim 9–5	1.3333	0.4713	2.2835	0.1845
Sim 10–1	1.2672	0.2148	2.3878	0.0780
Sim 10–2	3.9493	5.0740	2.1491	0.3282
Sim 10–3	1.2960	0.3543	2.3587	0.1288
Sim 10–4	1.2655	0.1985	2.3880	0.0935
Sim 10–5	1.3461	0.5047	2.3420	0.1723
Sim 11–1	1.2621	0.1883	2.3495	0.0704
Sim 11–2	1.2713	0.2337	2.3252	0.1286
Sim 11–3	1.2828	0.2926	2.2587	0.1163
Sim 11–4	1.2557	0.1192	2.3132	0.0555
Sim 11–5	1.2897	0.3246	2.3571	0.0609
Sim 12–1	1.2738	0.2509	2.2303	0.0778
Sim 12–2	1.2773	0.2687	2.3435	0.1526
Sim 12–3	1.2919	0.3353	2.3816	0.0539
Sim 12–4	1.2727	0.2461	2.3931	0.0735
Sim 12–5	1.2663	0.2040	2.4089	0.0536
Sim 13–1	1.2542	0.3791	2.3778	0.1048
Sim 13–2	1.2047	0.0356	2.4298	0.0127
Sim 13–3	1.2224	0.1605	2.3710	0.0728
Sim 13–4	1.2078	0.0543	2.4258	0.0272
Sim 13–5	1.2177	0.1294	2.3434	0.0810
Sim 14–1	1.3954	0.6427	2.3028	0.3348
Sim 14–2	1.2892	0.3178	2.3830	0.1313
Sim 14–3	1.3819	0.5978	2.4025	0.1582
Sim 14–4	1.2545	0.1073	2.4140	0.0393
Sim 14–5	1.2771	0.2762	2.3648	0.1344
Sim 15–1	1.2739	0.2487	2.4057	0.0895
Sim 15–2	1.2699	0.2277	2.3029	0.1311
Sim 15–3	1.3117	0.4031	2.2898	0.2045
Sim 15–4	1.2847	0.2997	2.3582	0.1566
Sim 15–5	1.2881	0.3172	2.3468	0.1681
Sim 16–1	1.3592	0.5365	2.2934	0.1727
Sim 16–2	1.4549	0.8406	2.2658	0.1925
Sim 16–3	1.3003	0.3668	2.2925	0.2156
Sim 16–4	1.2899	0.3370	2.3030	0.1182
Sim 16–5	1.2771	0.2679	2.2534	0.1404

Table 7. Simulation results: The mean result of five simulations in Each Experiment

Experiment	$\overline{\pi_{All}}$	$\overline{\delta_\pi}$	$\overline{U_{All}}$	$\overline{\delta_U}$
Sim 1	1.2495	0.0022	2.4204	0.0007
Sim 2	1.2497	0.0106	2.4199	0.0037
Sim 3	1.2496	0.0078	2.4203	0.0023
Sim 4	1.2495	0.0031	2.4204	0.0013
Sim 5	1.2037	0.0064	2.4331	0.0021
Sim 6	1.2495	0.0047	2.4205	0.0013
Sim 7	1.2495	0.0014	2.4205	0.0004
Sim 8	1.2496	0.0091	2.4200	0.0030
Sim 9	1.3211	0.4312	2.3312	0.1202
Sim 10	1.8248	1.2693	2.3251	0.1601
Sim 11	1.2723	0.2317	2.3207	0.0863
Sim 12	1.2764	0.2610	2.3515	0.0823
Sim 13	1.2213	0.1518	2.3896	0.0597
Sim 14	1.3196	0.3884	2.3734	0.1596
Sim 15	1.2857	0.2993	2.3407	0.1500
Sim 16	1.3363	0.4698	2.2816	0.1679

References

1. **Allais M** (1947) Economie et Interet, Paris: Imprimerie Nationale
2. **Arifovic J** (1995) Genetic algorithms and inflationary economies, J Monetary Econ, **36**, 219–243
3. **Arifovic J** (1996) The behaviour of the exchange rate in the genetic algorithm and experimental economies, J Political Econ, **104** (3), 510–541
4. **Arifovic J, Bullard J, Duffy J** (1996) Learning in a model of economic growth and development, Working Paper, Simon Fraster University
5. **A. zariadis C, Drazen A** (1990) Threshold externalities in economic development, Quart J Econ, **105**, 501–526
6. **Bryant J, Wallace N** (1984) A price discrimination analysis of monetary policy, Rev Econ Stud, **51** (2), 279–288
7. **Bullard J, Duffy J** (1994) Using genetic algorithms to model the evolution of heterogeneous beliefs, Mimeo, Federal Reserve Bank of St. Louis and University of Pittsburgh
8. **Bullard J, Duffy J** (1998) A model of learning and emulation with artificial adaptive agents, J Econ Dyn Control, **22**, 179–207
9. **Chen SH, Yeh CH** (1996) Genetic programming learning in the cobweb model, In Angeline PJ, Kinnear KE (Eds), Advances in Genetic Programming Vol. II, Cambridge, MA: MIT Press, pp 443–466
10. **Diamond PA** (1965) National debt in a neoclassical growth model Am Econ Rev, **55** (5), 1126–1150
11. **Gale D** (1973) Pure exchange equilibrium of dynamic economic models J Econ Theory, **6**, 12–36
12. **Grandmont JM** (1985) On endogenous competitive business cycles, Econometrica, **22**, 995–1037
13. **Kareken JH, Wallace N** (1981) On the indeterminacy of equilibrium exchange rates, Quart J Econ, **96**, 207–222
14. **Koza JR** (1992) Genetic Programming: On the Programming of Computers by Means of Natural Selection, Cambridge: MIT Press
15. **Lucas RE Jr** (1972) Expectations and the neutrality of money J Econ Theory, **4**, 103–124
16. **Lucas RE Jr** (1986) Adaptive behaviour and economic theory, J Business, **59**, 401–426
17. **Marimon R, Sunder S** (1993) Indeterminacy of equilibria in a hyperinflationary world: experimental evidence, Econometrica, **61**, 1073–1108
18. **Marimon R, Sunder S** (1994) Expectations and learning under alternative monetary regimes: an experimental approach, Econ Theory, **4**, 131–162
19. **Samuelson PA** (1958) An exact consumption-loan model of interest with or without the social contrivance of money, J Political Econ, **66** (6), 467–482
20. **Sargent TJ, Wallace N** (1982) The real bills doctrine vs. the quantity theory: a reconsideration, J Political Econ, **90**, 1212–1236
21. **Shell, K** (1971) Notes on the economics of infinity, J Political Econ, **79**, 1002–1012
22. **Tesfatsion L** (1996) How economists can get Alife, In Arthur B, Durlauf S, Lane D (Eds) The Economy as an Evolving Complex Systems II, Santa Fe Institute in the Science of Complexity, Vol. XXVII, Reading, MA: Addison-Wesley, pp 533–564