ELSEVIER Journal of Monetary Economics 36 (1995) 219–243

JOURNAL OF Monetary ECONOMICS

Genetic algorithms and inflationary economies

Jasmina Arifovic

Department of Economics, Simon Fraser University, Burnaby, B.C. V5A 1S6, Canada

(Received February 1993; final version received November 1994)

Abstract

This paper studies overlapping generations economies in which agents use genetic algorithms to learn correct decision rules. The results of computer simulations show that a genetic algorithm converges to the unique monetary steady state in case of a constant money supply policy and to the low-inflation stationary equilibrium in case of a constant real deficit financed through seignorage. Features of the genetic algorithm adaptation are compared to the performance of other learning algorithms and to the behavior observed in experiments with human subjects in the same OLG environments.

Key words: Learning; Genetic algorithms; Monetary policy

JEL classification: D83; C63

1. Introduction

Economic models with incomplete market structure, where money helps to overcome the constraints imposed by limited exchange possibilities, usually have a continuum of equilibria. The overlapping generations (OLG) model (Samuelson, 1958; Wallace, 1980) of a monetary economy is one such example that involves multiplicity of perfect foresight or rational expectations equilibrium

I am grateful to my thesis advisors Robert Lucas, Thomas Sargent, and Michael Woodford for their valuable advice and continuing support. I would also like to thank Susan Bennet, Jim Bullard, John Duffy, Curtis Eaton, and Ramon Marimon for their useful comments. Financial support from the Social Sciences Research Council of Canada is gratefully acknowledged.

SSDI 0304393295012O3 Z

paths. Whereas the rational expectations hypothesis does not characterize behavior outside equilibrium paths, the hypothesis of adaptive behavior does provide guidance on how agents may behave under any observed history. As Lucas (1986) has suggested, the stability results obtained through the analysis of learning dynamics in these models can be used to single out more likely equilibria since these dynamics may be a plausible conjecture about actual human behavior. OLG economies with fiat money have already been the subject of studies of learning (Lucas, 1986; Marcet and Sargent, 1989; Woodford, 1990; Grandmont and Laroque, 1991; Evans and Honkapohja, 1994, 1995; Bullard, 1994; Duffy, 1994). They have also been simulated in experiments with human subjects (Lim, Prescott, and Sunder, 1994; Marimon and Sunder, 1993; Marimon, Spear, and Sunder, 1993; Arifovic, 1992).

This paper studies a two-period OLG model with fiat money in which agents' learning is modeled using a *genetic algorithm* (Holland, 1975). Learning by genetic algorithm (GA) is examined in the context of two OLG environments, one with the policy of a constant money supply and the other with the policy of a constant real deficit financed through seignorage.

The behavior of GAs and other computer-based adaptive algorithms has been studied in a number of different economic models (for example, Miller, 1989; Marimon, McGrattan, and Sargent, 1989; Rust, Palmer, and Miller, 1994; Arifovic, 1994a, 1994b; Bullard and Duffy, 1994).[1] Economic agents' ability to learn Nash equilibrium behavior, equilibrium selection, and the usefulness of these algorithms in the computation of equilibria are some of the issues examined in these studies. Further, the behavior of these algorithms has been compared to the behavior observed in laboratory experiments with human subjects (Crawford, 1991; Miller and Andreoni, 1990a, 1990b; Arthur, 1991; Arifovic, 1992). The results of these studies suggest that computer-based adaptive algorithms can perform better than models with rational economic agents in explaining some of the regularities observed in experimental economics.[2]

In economic modeling, GAs describe the evolution of a population of agents' decision rules which are represented by *chromosomes*, strings of finite length, written over a binary alphabet $\{0, 1\}$.The performance of each chromosome in a given environment is evaluated through its *fitness function* which measures the value of profit or utility resulting from the behavior prescribed by the chromosome. The rules are updated using a set of four genetic operators: reproduction, crossover, mutation, and election. Reproduction makes copies of individual

[1]Arifovic (1994b) studies a two-country OLG model in which agents use the GA to update their consumption and portfolio decisions. Bullard and Duffy (1994) examine a sequence of N-period OLG environments in which agents use the GA to update their price expectations.

[2]Sargent (1993) contains descriptions of several economic models in which GAs were used.

chromosomes: chromosomes with higher fitness values have a higher probability of being reproduced. The creation of new rules is accomplished through the application of crossover and mutation. Crossover randomly exchanges parts of chromosomes, while mutation changes a bit value on a randomly chosen position in a binary string. Finally, the election operator tests newly generated rules. They are placed in a new population of chromosomes only if their performance on past data is better than their parents'.

GA economies can be viewed as models of decentralized learning. This fact allows for the examination of the evolution of heterogeneous beliefs and their convergence towards common beliefs. Convergence towards homogeneous beliefs represents, at the same time, the GA economy's convergence to a rational expectations equilibrium. Furthermore, the GA imposes a low requirement on the computational ability of economic agents. Agents that use GAs for the updating of their beliefs do not derive their decision rules from first-order conditions of their optimization problem. Yet, the algorithm exhibits relatively fast convergence to a rational expectations equilibrium.

This paper compares the results of an OLG model with GA learning to the results of the same model where the agents form expectations via either the sample average of past prices or least squares adaptive algorithms. One of the objectives of the paper is to examine whether the GA selects the same equilibria as do these alternative learning algorithms and whether it is sensitive to the conditions which result in the least squares algorithm's divergence. Another objective is to compare these algorithms in their ability to capture features of the behavior exhibited in experiments with human subjects.

The OLG economy with a constant supply of money has two stationary equilibria, autarky, in which fiat money has no value, and a stationary monetary equilibrium, which is unstable under perfect foresight dynamics. The autarkic equilibrium is the attractor of all equilibrium price paths with an initial price that is greater than the stationary monetary equilibrium price. Under the perfect foresight dynamics, the stationary monetary equilibrium is attainable only if the initial price is equal to the stationary monetary equilibrium price level. GA simulations of this OLG environment converge to the stationary equilibrium in which fiat money is valued. This equilibrium is also the point of convergence of the adaptive algorithm which uses the sample average of past price levels for the price forecasting (Lucas, 1986). Likewise, the experimental OLG economies simulated by Lim, Prescott, and Sunder (1994) exhibited price paths close to the stationary monetary equilibrium.

The model with a constant deficit financed via seignorage has two stationary equilibria in which money is valued, a low-inflation and a high-inflation stationary equilibrium. The system attains the low-inflation stationary equilibrium only if the initial inflation rate is equal to the inflation rate of this stationary equilibrium. All perfect foresight paths whose initial inflation rate is greater than the low stationary inflation rate converge to the high-inflation, Pareto-inferior stationary equilibrium. These stability conditions imply that an increase in the deficit results in a decrease in the inflation rate of a stable stationary equilibrium. The least squares learning algorithm (Marcet and Sargent, 1989) converges to the low-inflation stationary equilibrium provided the deficit is low enough.[3] The GA converges to the low-inflation stationary equilibrium as well. Results of simulations show that the GA also converges for deficit values and initial conditions for which least squares exhibited divergent behavior. This is important because inflationary paths observed in experiments with human subjects conducted by Marimon and Sunder (1993) were close to the low-inflation stationary equilibrium. Moreover, the experimental economies did not exhibit divergent inflationary paths in cases of deficit values and initial conditions for which least squares did not converge. These results were confirmed by Arifovic's (1992) laboratory experiments. Based on this evidence, the GA economies behave more like the economies with human subjects than the least squares economies or the economies in which agents update their beliefs using a sample average of past prices.

The rest of the paper is organized in four sections. The description of the OLG model and of the rational expectations equilibrium paths under both policies is given in Section 2. The GA and its application to an OLG economic environment are presented in the third section, while Section 4 contains the results of GA computer simulations. Comparisons to the behavior of other learning algorithms and to the features of the experimental OLG economies are presented in the fifth section. Concluding remarks are given in Section 6.

2. The economic model

The economy consists of overlapping generations of two-period-lived agents. Each generation consists of an equal number, N, of agents. Every agent of generation t lives over two consecutive periods, t and $t + 1$, and consumes c_t^1 in the first period (youth) and c_t^2 in the second period (old age). When young, each agent is endowed with w^1 units of a perishable consumption good, and with w^2 units when old ($w^1 > w^2$). The amount of fiat money that government supplies at time t is given by Nh_t, where h_t is the nominal per capita money supply.

[3] The distinction should be made between the two uses of the term convergence. When considering the convergence of the sample average of past prices and the least squares, it is used in an analytical sense, while when analyzing the GA, it is used in a computational sense, i.e., it refers to the convergence obtained in computer simulations.

All agents have identical preferences given by $U(c_t^1, c_t^2) = c_t^1 c_t^2$. Each young individual faces the following maximization problem:

$$\max c_t^1 c_t^2,$$

$$\text{s.t.} \quad m_t = (w^1 - c_t^1)p_t, \quad c_t^2 = w^2 + (m_t/p_{t+1}),$$

where m_t represents the nominal money balances that an agent acquires in the first period and spends in the second period of his life and p_t is the nominal price level at time period t. The first-order conditions of a young agent's maximization problem are simplified to give

$$\frac{w^2 + (m_t/p_{t+1})}{p_t} = \frac{w^1 - (m_t/p_t)}{p_{t+1}}.$$

In equilibrium, the nominal money demand per capita, m_t, must equal nominal per capita money supply, h_t, in each period. Eq. (1), which describes the behavior of the equilibrium nominal price level, is derived from the above first-order conditions after substituting h_t for m_t:

$$p_{t+1} = (w^1/w^2)p_t - (2/w^2)h_t. \tag{1}$$

2.1. Constant money supply

If the government keeps the amount of money constant, $h_t = h$ for all t, the difference equation (1) has a stationary solution with valued fiat money. In fact, it is the unique equilibrium in which a version of the quantity theory of money holds. It is given by $p_t = p^*$ for all t, where

$$p^* = 2h/(w^1 - w^2). \tag{2}$$

This stationary competitive equilibrium price system exists for $w^1/w^2 > 1$. It is also Pareto-optimal, with first-period consumption equal to second-period consumption ($c^{1,*} = c^{2,*}$). Since this equilibrium is unstable, the economy attains it only if $p_0 = p^*$.

There is also a continuum of monetary equilibria indexed by the initial price level p_0 in the interval (p^*, ∞). All of the equilibria with an initial price greater than p^* converge to the stationary equilibrium in which money has no value. This has led some to suppose that the stationary equilibrium with valued money is unlikely to be reached. The results of the learning analysis provide an interesting contrast to this view.

Lucas (1986) studied an infinite-horizon OLG economy with a constant money supply in which agents use the sample average of past price levels to form expectations about next period's price level. The rule that agents use to update their expectations is given by

$$p_{t+1}^e = \frac{t}{t+1} p_t^e + \frac{1}{t+1} p_{t-1}, \tag{3}$$

where p^e_{t+1} is the point expectation formed at t about the price in $t+1$. The rule defined by Eq. (3) is equivalent to

$$p^e_{t+1} = \frac{1}{t+1}[p_t + \cdots + p_0]. \tag{4}$$

Given some $p^e_1 > p^*$, the system with this adaptive scheme converges to the stationary monetary equilibrium.

2.2. Constant deficit financed through seignorage

If the government follows a policy of financing the constant real per capita deficit, d, through seignorage, the monetary rule that would implement this policy is given by

$$d = (h_t - h_{t-1})/p_t. \tag{5}$$

Thus the money supply in period t, h_t units per head, is no longer constant. As a result, we have

$$h_t = h_{t-1} + dp_t. \tag{6}$$

From the first-order conditions of the consumer maximization problem, nominal money balances that an agent of generation t carries from time t to time period $t+1$ are given by

$$m_t = \tfrac{1}{2}p_t(w^1 - \pi_{t+1}w^2), \tag{7}$$

where $\pi_{t+1} = p_{t+1}/p_t$ is the inflation rate between period t and $t+1$. Thus, using Eqs. (5) and (7), and the equilibrium condition that $m_t = h_t$, the deficit d is given by

$$d = (w^1/2) - \pi_{t+1}(w^2/2) - (w^1/2\pi_t) + (w^2/2). \tag{8}$$

Rearranging (8), the paths of equilibrium inflation rates under perfect foresight dynamics are

$$\pi_{t+1} = (w^1/w^2) + 1 - (2d/w^2) - (w^1/w^2)(1/\pi_t). \tag{9}$$

Provided that the deficit satisfies $d < d_{\max} = (w^2/2)[1 + (w^1/w^2) - 2(w^1/w^2)^{1/2}]$, Eq. (9) has two real stationary solutions, a low-inflation stationary equilibrium, π^*_1, and a high-inflation stationary equilibrium, π^*_2, given by

$$\pi^*_{1,2} = \frac{(w^1/w^2) + 1 - (2d/w^2) \pm \sqrt{((w^1/w^2) + 1 - (2d/w^2))^2 - 4(w^1/w^2)}}{2}.$$

The high-inflation stationary equilibrium is the stable solution, being the attractor for a continuum of rational expectations equilibrium paths, starting from $\pi_0 \in (\pi_1^*, w^1/w^2)$. For the value of π_0 equal to w^1/w^2, zero money balances are demanded. If the initial inflation rate π_0 is equal to π_1^*, the system attains a low-inflation stationary equilibrium. In a stationary equilibrium, the ratio between first- and second-period consumption is equal to the stationary inflation rate. The low-inflation stationary equilibrium is Pareto-superior to the high-inflation one.

Marcet and Sargent (1989) analyze a least squares learning scheme in the context of this model. If agents learn about the price level using the least squares algorithm, their expectation of the price in $t + 1$ is

$$p_{t+1}^e = \beta_t p_t, \tag{10}$$

where β_t is the estimate of the inflation rate obtained through the regression on past values of prices:

$$\beta_t = \left[\sum_{s=1}^{t-1} p_{s-1}^2\right]^{-1} \left[\sum_{s=1}^{t-1} p_s p_{s-1}\right]. \tag{11}$$

Marcet and Sargent show that under least squares learning the model either converges to the low-inflation stationary equilibrium or no equilibrium exists.[4] Note that under the rational expectations hypothesis the stable stationary equilibrium is a high-inflation one. This result of least squares learning is also classical in the sense that a higher deficit is associated with a higher stable stationary inflation rate. Under the rational expectations, an increase in the deficit results in a lower stationary inflation rate. Further, for some cases of high deficit for which there exists an equilibrium under the rational expectations hypothesis, there is no equilibrium under least squares learning.

3. Genetic algorithm application

At each integer point in time $t \geq 1$, there are two populations of chromosomes, one being the new population of generation t, the young, the other being the population of generation $t - 1$, the old. A population of generation t, $A(t)$, consists of N chromosomes $A_{i,t}$, $i \in [1, \ldots, N]$, which represent decision rules about first-period consumption for N agents. A *chromosome* is a string of finite length ℓ, written over the binary alphabet $\{0, 1\}$.

Decoding and normalization of a chromosome yields a real number that represents the value of first-period consumption. For a chromosome i of length

[4]With an alternative preference map, it can happen that the system under least squares would converge to a periodic equilibrium. See Bullard (1994).

ℓ the decoding works in the following way:

$$x_{i,t} = \sum_{k=1}^{\ell} a_{i,t}^{k} 2^{k-1},$$

where $a_{i,t}^{k}$ is the value (0, 1) taken at the kth position in the ith string. After decoding, the integer $x_{i,t}$ is normalized in order to obtain a value of the first-period consumption $c_{i,t}^{1}$, $c_{i,t}^{1} \in [0, w^{1}]$ of agent i of generation t. Thus,

$$c_{i,t}^{1} = x_{i,t}/\bar{K}, \tag{12}$$

where $\bar{K}$ is a coefficient chosen to normalize the value of $x_{i,t}$. Once $c_{i,t}^{1}$ is determined, savings of a chromosome i of generation t, $s_{i,t}$, are given as

$$s_{i,t} = w^{1} - c_{i,t}^{1}. \tag{13}$$

The price of the consumption good at time t is then given by

$$p_t = Nh \Big/ \sum_{i=1}^{N} s_{i,t}. \tag{14}$$

The nominal price p_t and individual savings $s_{i,t}$ are used to compute nominal money balances, $m_{i,t}$, that agent i ($i \in [1, N]$) of generation t carries from period t to period $t+1$,

$$m_{i,t} = p_t s_{i,t}. \tag{15}$$

At time $t+1$, the second-period consumption of member i ($i \in [1, N]$) of generation t is determined as

$$c_{i,t}^{2} = (m_{i,t}/p_{t+1}) + w^{2}. \tag{16}$$

The fitness of a chromosome i of generation t is given by the value of agent i's utility at the end of period $t+1$ (the second period of life):

$$\mu_{i,t} = U_i(c_{i,t}^{1}, c_{i,t}^{2}) = c_{i,t}^{1} c_{i,t}^{2}.$$

Beliefs about first-period consumption of members of generation t are updated using four operators: reproduction, crossover, mutation, and election. The population of updated rules is then used by members of generation $t+2$.

Reproduction makes copies of individual strings. The probability that a string will be copied is proportional to its fitness value. Thus a probability that a string $A_{i,t}$ will get a copy $C_{i,t}$ is given by

$$P(C_{i,t}) = \mu_{i,t} \Big/ \sum_{i=1}^{N} \mu_{i,t}, \quad i \in [1, N].$$

Reproduction operates like a biased roulette wheel. Each string is allocated a slot sized in proportion to its fitness. The number of spins of the wheel is equal to the number of chromosomes in a population and each spin yields a reproduction candidate. When a chromosome is selected, its exact copy is made. Once N copies are made (the number of strings in a population is kept constant), they enter into a *mating pool* to undergo application of other genetic operators.

Crossover operates on members of the mating pool. First, two strings are selected from the mating pool at random. Then, an integer number k is selected from $(1, \dots, \ell - 1)$ again at random. Two new strings are formed by swapping the set of values to the right of the position k. The total number of pairs that is selected is $N/2$ (where N is an even integer). Crossover takes place on each pair with probability *pcross*. An example of the crossover between two chromosomes for $\ell = 8$ and $k = 4$ is given below:

$$\begin{cases} 1 \quad 1 \quad 0 \quad 0 \mid 1 \quad 0 \quad 1 \quad 0 \\ 0 \quad 0 \quad 0 \quad 1 \mid 0 \quad 1 \quad 0 \quad 1 \end{cases}$$

After the application of crossover, two resulting strings are

$$\begin{cases} 1 \quad 1 \quad 0 \quad 0 \quad 0 \quad 1 \quad 0 \quad 1 \\ 0 \quad 0 \quad 0 \quad 1 \quad 1 \quad 0 \quad 1 \quad 0 \end{cases}$$

Mutation randomly changes the value of a position within a string. The value at each position within a string is exposed to a chance of being altered to the other value taken by the binary alphabet. The probability of mutation, *pmut*, is independent and identical across positions.

The election operator (Arifovic, 1991, 1994a)[5] tests newly generated offspring before they are allowed to enter into a new population of generation $t + 2$. The string value of each new offspring is decoded in order to obtain the value of first-period consumption that an offspring would represent were it used as an actual decision rule. Utility associated with that decision rule is computed using the inflation rate of period t. The utility obtained in this manner represents the offspring's *potential* fitness value. This potential fitness of an offspring is compared to the *actual* fitness values of its parents (i.e., the fitness values of the two parent strings that were evaluated at the end of period t). If potential fitness is higher than the fitness of one or both of parents, then the offspring enters into the population of a new generation. The possible results of the election operator test are the following: If only one offspring (out of two offspring for each parent's pair) has a fitness higher than both of its parents, it replaces the parent with a lower fitness while the parent with a higher fitness remains in the population. In the case that both offspring have fitnesses higher than a fitness value of each

[5] I developed the election operator in my dissertation at the University of Chicago. It is a form of an elitist selection procedure which reduces the effects of the mutation operator over the course of a simulation. For more detail, see Arifovic (1991, 1994a).

parent, they replace both parents as new members of the population. If both parents have fitnesses higher than their offspring, they remain in the population of the new generation, while their offspring do not enter.

A population of chromosomes that will represent decision rules of young agents at $t + 2$ is generated in the following way: First, the application of the reproduction operator to the population of rules of generation t yields a population of N copies. Then, crossover and mutation operators are applied to this population in order to generate new ideas to be tried. Finally, newly generated chromosomes are subjected to the election operator test. Offspring that pass the test, together with parents that are more fit than their offspring, form a new population of decision rules of agents born at time period $t + 2$.

Once the members of the new population are determined, the values for first-period consumption, $c^1_{i,t+2}$, and savings, $s_{i,t+2}$, for each agent of generation $t + 2$ are computed. Aggregate savings, $\sum_{i=1}^{N} s_{i,t+2}$, together with the aggregate nominal money supply Nh (which is equal to the money holdings of agents of generation $t + 1$, $\sum_{i=1}^{N} m_{i,t+1}$), determine the price level of the good that prevails at $t + 2$. The price of the consumption good at time $t + 2$ is thus given by

$$p_{t+2} = Nh \Big/ \sum_{i=1}^{N} s_{i,t+2} \,. \tag{17}$$

Then the second-period consumption of member i ($i \in [1, N]$) of generation $t + 1$ is determined:

$$c^2_{i,t+1} = (m_{i,t+1}/p_{t+2}) + w^2 \,. \tag{18}$$

Finally, fitness values of the members of generation $t + 1$ are computed. The population for generation $t + 3$ is generated from the population of generation $t + 1$, using the genetic operators reproduction, crossover, mutation, and election.

The populations of chromosomes that belong to members of generations 0 and 1 are randomly generated. The system starts off with Nh units of money distributed to the initially old.

The algorithm applied in the case of the OLG model with constant real deficit is identical to the one used in the model with constant money supply, except for the way in which a price is computed in each time period t. Since government finances a constant deficit per head, d, the price in period t is given by

$$p_t = \sum_{i}^{N} s_{i,t-1} p_{t-1} \Big/ \left(\sum_{i}^{N} s_{i,t} - Nd \right), \tag{19}$$

for $\sum_i^N s_{i,t-1} > Nd$. (Note that there is nothing in the algorithm to ensure aggregate savings greater than government total deficit in every time period. Aggregate savings less than Nd is interpreted as a breakdown of the GA monetary economy.)

This GA economy allows for the possibility of savings, but does not give agents a market in which to borrow. There is only a money market with no credit market or storage technology in the model. This GA design closely corresponds to the design of the experimental OLG economies and to the environments in which the behavior of sample average of past prices and least squares learning algorithms were examined. The GA design with a credit market would have prevented the comparison of the GA behavior generated by the two above-mentioned algorithms as well as with the behavior observed in the experiments.[6]

4. Results of simulations

GA simulations for both OLG models, with constant money supply and with constant real deficit, were conducted using populations of thirty strings, with a string length of thirty bits. Initial populations were randomly generated. Each simulation was conducted for 400 periods. For every set of OLG parameter values, simulations with eight different sets of genetic operator rates were conducted in order to examine how these rates affect the GA's behavior. These sets of values are given in Table 1. In addition, every combination of OLG parameter values and genetic operator values was examined using multiple runs with different seed values for the random number generator to ensure the robustness of results to different random number sequences.

4.1. Constant money supply

The results of the computer simulations of GA OLG economies with constant money supply show convergence towards the stationary monetary equilibrium (the same equilibrium to which the Lucas' adaptive scheme converges). The OLG parameter values and the corresponding stationary values of the price level and the first- and second-period consumption values are given in Table 2.

For all sets of OLG parameter values, the fastest convergence was achieved with the second set of genetic operator rates. The price series generated using the GA is shown in Fig. 1.

---

[6]A more general GA design which would include a credit market is worth examining in the OLG environments in which a fraction of agents has their endowment pattern with $w^1 < w^2$. This would make it worthwhile for these agents to learn to borrow the utility maximizing amount in the first period of their lives. On the other hand, if $w^1 > w^2$ for all agents, the rules that instruct agents to borrow would result in the relatively low levels of utility and would, over time, disappear from the population through the workings of the reproduction operator.

Table 1
Crossover and mutation rates for genetic algorithm

Set	1	2	3	4
pcross	0.6	0.6	0.75	0.75
pmut	0.0033	0.033	0.033	0.0033
Set	5	6	7	8
pcross	0.9	0.9	0.3	0.3
pmut	0.0033	0.33	0.033	0.33

Table 2
Overlapping generations model with constant money supply

Model	1	2	3	4
w^1	150	120	100	7
w^2	10	20	90	1
h/N	1000	500	1000	100
p^*	14.286	10	200	33.3
$c^{1,*}$	80	70	95	4

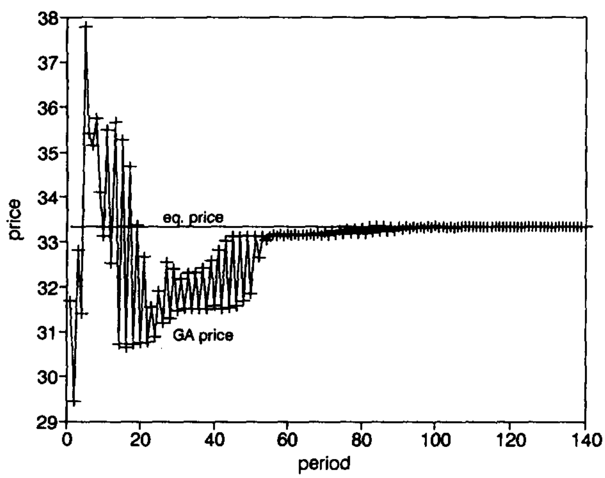

Fig. 1. GA OLG economy with constant money supply; GA price; set 4 of OLG parameter values; set 2 of genetic rates values.

Beliefs about the amount that should be consumed in the first period in both populations converge to the stationary value, $c^{1,*}$. Thus, once the algorithm converges, the members of both populations make decisions about how much to consume in the first period as if they had learned to maximize their utility functions and make the correct price prediction.

4.2. Constant deficit

GA simulations of constant deficit OLG model were conducted for thirteen different sets of OLG model parameter values (endowment patterns and values of deficit). Values from these simulations, together with the values of rational expectations stationary equilibria inflation rates (π_1^* and π_2^*) and the corresponding stationary values of first- and second-period consumption ($c_1^{1,*}$ and $c_1^{2,*}$ associated with π_1^* and c_2^{1*} and c_2^{2*} associated with π_2^*) are given in Table 3.

The GA OLG economies converged to the equilibrium with the low stationary inflation rate for all sets of OLG parameter values and all sets of genetic rate

Table 3
Overlapping generations model with constant real deficit

Set	1	2	3	4	
w^1	100	10	150	10	
w^2	90	2	30	9	
d	0.02	0.001	15	0.0007695	
π_1^*	1.0041	1.00025	1.382	1.00156	
π_2^*	1.1065	4.99987	3.618	1.10938	
$c_1^{1,*}$	95.1878	6.00025	95.7295	9.50702	
$c_2^{1,*}$	99.7922	9.99875	129.2705	9.9922	
Set	5	6	7	8	
w^1	2	2	2	2	
w^2	1.8	1.8	1.8	1.8	
d	0.0019	0.00234	0.0024	0.00263	
π_1^*	1.026	1.0357	1.0377	1.052	
π_2^*	1.0838	1.0728	1.0707	1.056	
$c_1^{1,*}$	1.9227	1.9321	1.9339	1.9469	
$c_2^{1,*}$	1.9754	1.9655	1.9637	1.9505	
Set	9	10	11	12	13
w^1	10	10	10	10	10
w^2	4	4	4	1	1
d	0.001	0.5	0.67544	0.001	1.5
π_1^*	1.00033	1.25	1.57922	1.00022	1.55051
π_2^*	2.49917	2.00	1.58306	9.99778	6.44949
$c_1^{1,*}$	7.00067	7.5	8.15843	5.50011	5.77526
$c_2^{1,*}$	9.9983	9.00	8.16613	9.99889	8.22474

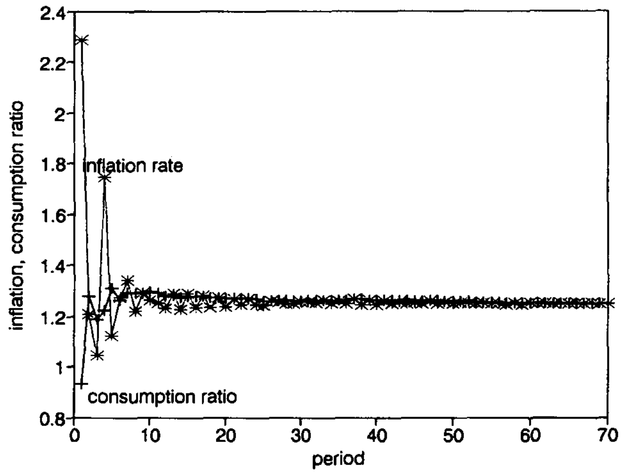

Fig. 2. GA OLG economy with constant deficit – GA inflation rate and consumption ratio; set 1 of OLG parameter values; set 2 of genetic rates values.

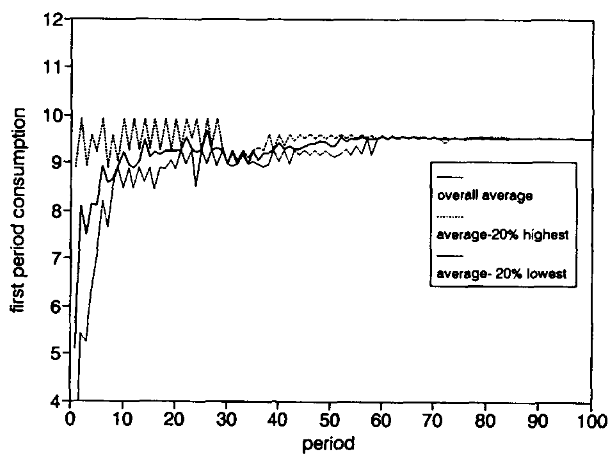

Fig. 3. GA OLG economy with constant real deficit – GA first period consumption; set 4 of OLG parameter values; set 2 of genetic rates values.

values for which the simulations were conducted. Again, set 2 of genetic rate values resulted in the fastest convergence of GA for all sets OLG model parameter values.

Fig. 2 exhibits the behavior of the average ratio between first- and second-period consumption and the inflation rate. Both of these variables fluctuate at the beginning of a simulation, while after period 50 fluctuations start to wear off as both variables approach and, at the end, converge to the common value equal to the low stationary inflation rate.

The evolution of the beliefs about first-period consumption is presented in Fig. 3. Beliefs in odd time periods are the beliefs of population 1, while those in even periods are the beliefs of population 2. Both populations show a great extent of diversity of decision rules at the beginning of the simulation. The diversity starts to decrease after period 40, as rules that have proven more successful in the environment up to that point begin to make up an increasing fraction of both populations via the application of the reproduction operator. Around simulation period 100, the beliefs of both populations converge to a single value, which is equal to the value of the first-period consumption in the rational expectations low-inflation stationary equilibrium. Beliefs of both populations, given as binary strings and as decoded real number values are given in the Appendix.

GA simulations in which the election operator is not included do not result in the algorithm's convergence. This simple GA (Goldberg, 1989) supports considerable population diversity until the end of simulations (3000 periods with no sign of convergence) due to the continuing effects of mutation. The election operator is used in this environment to offset these effects of mutation on population diversity.[7]

4.3. Comparison with least squares learning

In the computer simulations of their model, Marcet and Sargent obtained examples in which the least squares beliefs, $\{\beta_t\}$, failed to converge. These examples were generated using values of deficit very close to the maximum value attainable under perfect foresight dynamics and (or) using low values of initial belief, β_0. The objective of this section is to examine the behavior of the GA under the conditions which resulted in the divergent behavior of the least squares algorithm.

The failure of convergence of the least squares beliefs, $\{\beta_t\}$, to π_1^* is more likely to happen when π_2^* is sufficiently close to π_1^*. The stationary inflation rates, π_1^* and π_2^*, approach the common value $(w^1/w^2)^{1/2}$ as the deficit, d, approaches from below the maximal feasible value of d_{max}. The divergence of least squares

[7]It should be noted that election operator does not require any additional information that is not already used by GA agents. Application of this operator on population of generation t requires information about the inflation rate of period $t-1$, which is used by GA agents in computation of their utilities at the end of $t-1$, i.e., in computation of fitness values of their decision rules. The election operator also does not impede the ability for adjustment of the GA in environments in which parameters of the economic model change. For detailed discussion of the impact of election operator on GA performance see Arifovic (1991, 1994a).

may also occur if the initial belief β_0 is sufficiently smaller than π_1^* (when $\{\beta_t\}$ is approaching π_1^* from below).

The values of the first- and second-period endowment that Marcet and Sargent used in their computer simulations were $w^1 = 2$ and $w^2 = 1.8$. For these parameter values, the maximum value of deficit sustainable under rational expectations is $d_{max} = 0.00263$ (set 8 in Table 3). Three different values of deficit were tested, $d = 0.0019$ (set 5 in Table 3), $d = 0.00234$ (set 6 in Table 3), and $d = 0.0024$ (set 7 in Table 3), combined with different values of initial beliefs. The results of the least squares simulations were the following:

(a) For the relatively low value of deficit, $d = 0.0019$, $\{\beta_t\}$ exhibited fast convergence to π_1^*, even though the initial belief was set to 1.00, a value that is small compared to the value of π_1^* (1.026).

(b) When the value of deficit was increased to $d = 0.00234$, the least squares algorithm diverged away for the same value of initial belief, $\beta_0 = 1.00$. Decreasing the difference between β_0 and π_1^* (1.0357), by setting $\beta_0 = 1.02$, resulted in the convergence of the algorithm for the same value of deficit.

(c) Further increase in the value of deficit to $d = 0.0024$ resulted in the divergence of $\{\beta_t\}$ for the values of initial beliefs $\beta_0 = 1.00$, $\beta_0 = 1.02$, and $\beta_0 = 1.0376$. Even a value of β_0 extremely close to π_1^* (1.0377) could not prevent divergence of $\{\beta_t\}$.[8]

While least squares agents update their beliefs about next period's inflation rate, GA agents update their beliefs about the value of their first-period consumption. Moreover, at generation 0, beliefs of GA agents differ across the population, while beliefs of least squares agents are unanimous for all t. These differences exclude the possibility of simply setting the value of initial belief β_0 in GA simulations equal to the value of β_0 used in the least squares simulations. To start the GA simulation with the initial conditions that would correspond to Marcet and Sargent conditions, the GA was adjusted in the following way: Given β_0, the level of the first-period consumption of young agents at generation 0 was calculated from the first-order conditions of utility maximization. This value was taken as the average consumption of young GA agents at generation 0. The individual beliefs of the young at generation 0 about their first-period consumption were then diversified around this average value.

Thus adjusted, the GA population converged to the low-inflation stationary equilibrium for all three cases, (a), (b), and (c). In all of these simulations, the algorithm starts out with relatively large fluctuations around π_1^* which subside

[8]See Bullard (1994) for a detailed analysis of the Hopf bifurcation underlying these least squares results.

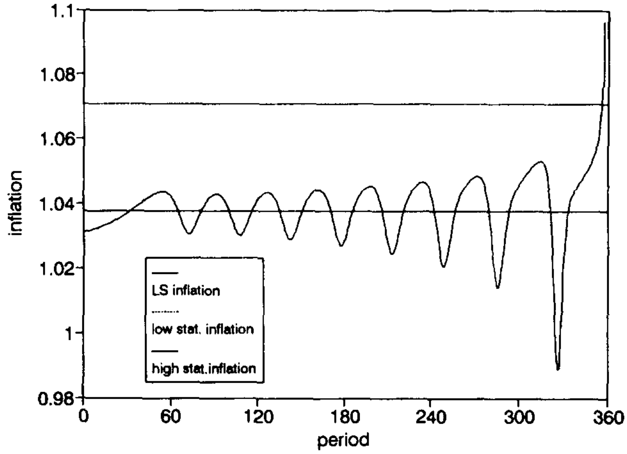

Fig. 4. OLG economy with constant deficit – least squares.

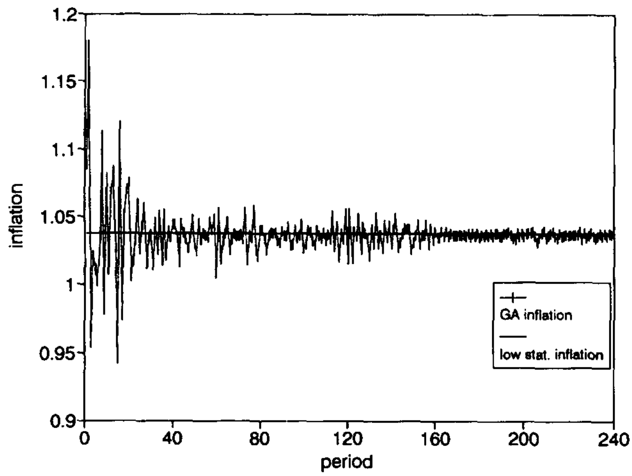

Fig. 5. OLG economy with constant real deficit – GA inflation rate; set 7 of OLG parameter values; set 2 of genetic rates values.

around period 200. The diverging path of the inflation rate generated by the least squares algorithm for $d = 0.0024$ is given in Fig. 4. The converging path of the GA inflation rate for the same value of deficit is presented in Fig. 5. These results indicate that the stability conditions for the convergence of GA economies and for the convergence of least squares learning in this OLG model are different.

5. Genetic algorithm and the experimental evidence

Lim, Prescott, and Sunder (1994) conducted experiments with human subjects of OLG economies with a constant money supply as described in Section 2. Their basic result is that trading patterns and prices converged to the Pareto-optimal rational expectations equilibrium of constant consumption (prices show an upward bias that may be explained as an outcome of imperfect competition). Exploding paths of prices that result in an autarkic solution where money has no value were not observed in any of these experimental economies.

In experiments with human subjects, Marimon and Sunder (1993) simulated OLG economies where government finances a fixed real deficit through seignorage. The structure of the economies and the level of deficits were common knowledge and participants also observed past prices. The results show that experimental inflationary paths lie close to the low-inflation stationary equilibrium. Comparing rational expectations nonstationary paths and least squares nonstationary paths with experimental inflation paths, Marimon and Sunder conclude that least squares nonstationary paths have much better explanatory power than rational expectations nonstationary paths. However, least squares paths are smoother than the observed experimental paths.

In addition, when experimental economies start with initial inflation rates and deficit values for which the least squares learning algorithm does not converge to the low-inflation stationary equilibrium, inflation patterns in these experimental economies do not follow diverging paths, but rather paths that converge to the neighborhood of the low-inflation stationary equilibrium.[9] The GA generates inflation patterns that fluctuate more than least squares patterns and thus it captures better this feature of the experimental data. Furthermore, the GA converges for deficit values and initial conditions for which least squares do not and thus it explains better the same behavior observed in the OLG experiments.

Arifovic (1992) conducted the same type of OLG experiments with human subjects as Marimon and Sunder. Five different sets of endowment patterns and deficit values were used. These values are given in Table 3 (sets 9 through 13). Observations from these experiments confirm that the experimental inflation rate

[9]The fact that both adaptive algorithms and human subjects in the OLG environments considered in this paper select Pareto-superior stationary outcomes should not be interpreted to argue that this holds in general. There are results in the learning literature as well as in the experimental work that show that this need not be the case. For example, Woodford (1990) and Evans and Honkapohja (1994, 1995) derive the conditions for the convergence of learning dynamics to sunspot equilibria in OLG environments. Duffy (1994) provides an example of a learning process that converges to Pareto-inferior stationary equilibria of OLG economies. In the laboratory experiments of coordination games, Van Huyck, Battalio, and Beil (1990, 1991) show that experimental economies can select Pareto-inferior equilibria.

is close to the low stationary inflation rate and the experimental inflationary paths do not diverge away for the values that are critical for the least squares paths.

Observed experimental inflation rates for set 9 of the OLG parameter values are presented in Fig. 6. It is clear that the experimental inflation rate fluctuates around the low stationary inflation rate value. Least squares inflation rates for the same set of OLG parameter values are given in Fig. 7. Least squares learning exhibits smooth convergence to π_1^*, without the fluctuations that characterize

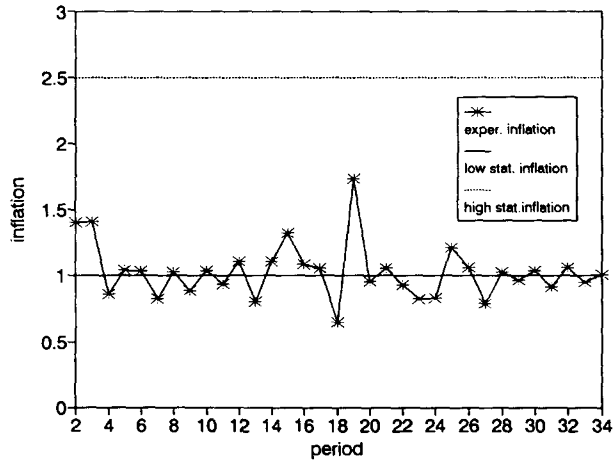

Fig. 6. OLG economy with constant real deficit – experimental inflation rate; set 9 of OLG parameter values.

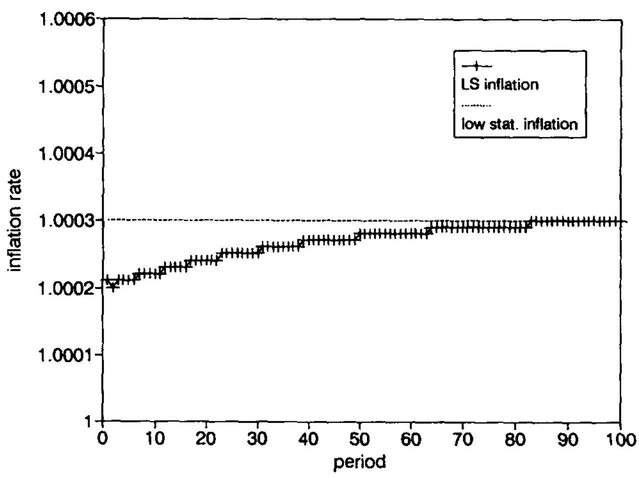

Fig. 7. OLG economy with constant deficit – least squares inflation rate; set 9 of OLG parameter values.

experiments. In contrast, the observed GA inflation rate fluctuates prior to its convergence to the low-inflation stationary equilibrium and thus captures this feature of the experimental data. Fig. 8 presents the behavior of GA inflation rates for set 9 of the OLG parameter values and set 6 of the genetic rates values.

Fig. 9 shows the behavior of the inflation rate in the experimental economy in which the OLG parameters of set 11 were used. For these parameter values,

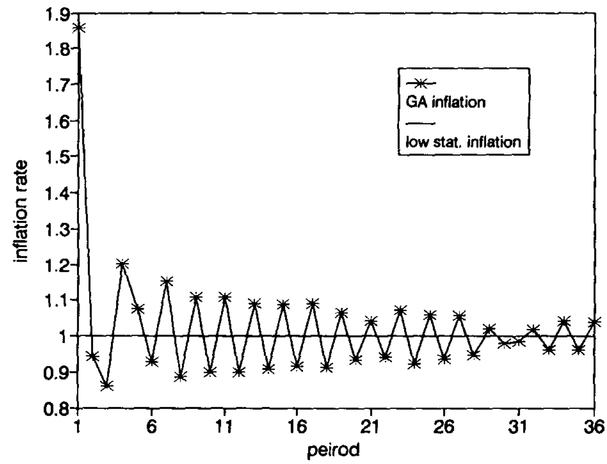

Fig. 8. OLG economy with constant real deficit – GA inflation rate; set 9 of OLG parameter values; set 6 of genetic rates values.

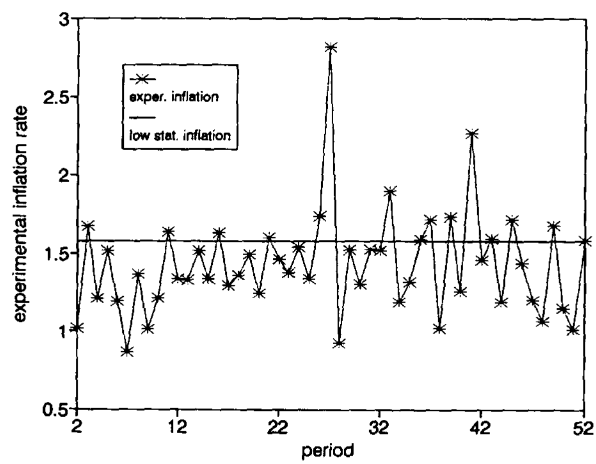

Fig. 9. OLG economy with constant real deficit – experimental inflation rate; set 11 of OLG parameter values.

least squares learning leads to divergent behavior of the inflation rate (Fig. 10). Although the extent to which the inflation rate fluctuates in this experimental economy is greater than in the others that were conducted for lower values of deficit, it does not exhibit a diverging pattern. The genetic algorithm does not diverge away for this set of OLG values either. Its fluctuating, but not diverging, inflation pattern for set 11 of OLG parameter values and set 7 of genetic rates values is given in Fig. 11.

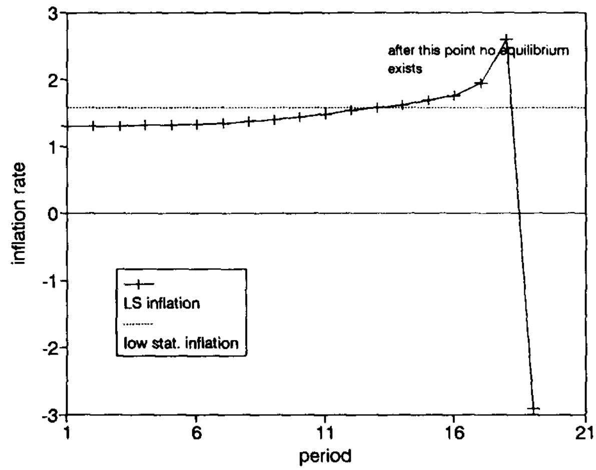

Fig. 10. OLG economy with constant real deficit – least squares inflation rate; set 11 of OLG parameter values.

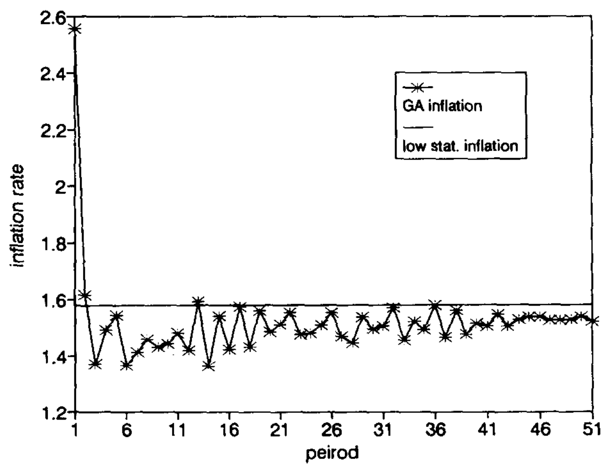

Fig. 11. OLG economy with constant real deficit – GA inflation rate; set 11 of OLG parameter values; set 7 of genetic rates.

6. Conclusion

The goal of this paper was to examine the behavior of GA OLG economies with fiat money in computer simulations and to compare the results to those generated by other learning algorithms in the same environments as well as to the observations from experiments with human subjects. Two types of policy rules were considered, one with a constant money supply and the other with a constant real deficit financed through seignorage.

In the constant money supply environment, OLG GA simulations resulted in convergence to the unique stationary monetary equilibrium with a constant price level. This is consistent with results obtained in the study of the adaptive scheme which uses the sample average of past price levels for price forecasting, as well as results from experiments with human subjects.

In the case of a constant deficit financed through seignorage, the GA converged to the low-inflation stationary equilibrium, the equilibrium which is unstable under perfect foresight dynamics, but locally stable under least squares learning for low values of the deficit. The inflation rates observed in the experiments with human subjects converged to the low-inflation stationary equilibrium as well. Computer simulations suggest that the GA is not sensitive to the initial conditions and values of deficit which cause the divergent behavior of the least squares algorithm. Further, the GA performs better in capturing fluctuations of the inflation rate recorded in the experimental economies.

Appendix

Table 4
GA population report, period 500

	String	First-period consumption
Population of old agents rules		
1)	111100110110000101001110011101	9.51
2)	111100110110000101001110011101	9.51
3)	111100110110000101001110011101	9.51
4)	111100110110000101001110011101	9.51
5)	111100110110000101001110011101	9.51
6)	111100110110000101001110011101	9.51
7)	111100110110000101001110011101	9.51
8)	111100110110000101001110011101	9.51
9)	111100110110000101001110011101	9.51
10)	111100110110000101001110011101	9.51
11)	111100110110000101001110011101	9.51
12)	111100110110000101001110011101	9.51
13)	111100110110000101001110011101	9.51

Table 4 (continued)

String	First-period consumption
14) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
15) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
16) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
17) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
18) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
19) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
20) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
21) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
22) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
23) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
24) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
25) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
26) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
27) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
28) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
29) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
30) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
Population of young agents rule	
1) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
2) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
3) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
4) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
5) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
6) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
7) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
8) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
9) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
10) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
11) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
12) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
13) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
14) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
15) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
16) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
17) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
18) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
19) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
20) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
21) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
22) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
23) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
24) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
25) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
26) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
27) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
28) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
29) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51
30) 1 1 1 1 0 0 1 1 0 1 1 0 0 0 0 1 0 1 0 0 1 1 1 0 0 1 1 1 0 1	9.51

References

Arifovic, J., 1991, Learning by genetic algorithms in economic environments, Doctoral dissertation (University of Chicago, Chicago, IL).

Arifovic, J., 1992, Genetic algorithm: Experimental evidence, Manuscript (McGill University, Montreal).

Arifovic, J., 1994a, Genetic algorithm learning and the cobweb model, Journal of Economic Dynamics and Control 18, 3–28.

Arifovic, J., 1994b, The behavior of the exchange rate in the genetic algorithm and experimental economies, Working paper no. 7 (Simon Fraser University, Burnaby).

Arthur, B., 1991, Designing economic agents that act like human agents: A behavioral approach to bounded rationality, American Economic Review: Papers and Proceedings of the 103rd Annual Meeting of the American Economic Association, 353–359.

Bullard, J., 1994, Learning equilibria, Journal of Economic Theory 64, 468–485.

Bullard, J. and J. Duffy, 1994, A model of learning and emulation with artificial adaptive agents, Manuscript (Federal Reserve Bank, St. Louis, MO; University of Pittsburgh, Pittsburgh, PA).

Cagan, P., 1956, The monetary dynamics of hyperinflation, in: M. Friedman, ed., The quantity theory of money (University of Chicago Press, Chicago, IL) 25–117.

Crawford, V.P., 1991, An 'evolutionary' explanation of Van Huyck, Battalio and Beil's experimental results on coordination, Games and Economic Behavior 3, 25–59.

Duffy, J., 1994, On learning and the nonuniqueness of equilibrium in an overlapping generations model with fiat money, Journal of Economic Theory 64, 541–553.

Evans, G. and S. Honkapohja, 1994, On the local stability of sunspot equilibria under adaptive learning rules, Journal of Economic Theory 64, 142–161.

Evans, G. and S. Honkapohja, 1995, Local convergence of recursive learning to steady states and cycles in stochastic nonlinear models, Econometrica 63, 195–206.

Goldberg, D.E., 1989, Genetic algorithms in search, optimization and machine learning (Addison-Wesley, New York, NY).

Grandmont, J.M. and G. Laroque, 1991, Economic dynamics with learning: Some instability examples, in: W.A. Barnett et al., eds., Equilibrium theory and applications (Cambridge University Press, Cambridge) 247–273.

Holland, J.H., 1975, Adaptation in natural and artificial systems (University of Michigan Press, Ann Arbor, MI).

Holland, J.H. and J.H. Miller, 1990, Artificial adaptive agents in economic theory, American Economic Review: Papers and Proceedings of the 103rd Annual Meeting of the American Economic Association, 365–370.

Lim, S., E. Prescott, and S. Sunder, 1994, Stationary solution to the overlapping generations model of fiat money: Experimental evidence, in: J.D. Hay, ed., Empirical economic theory, forthcoming.

Lucas, R.E., Jr., 1986, Adaptive behavior and economic theory, Journal of Business 59, 401–426.

Marcet, A. and T.J. Sargent, 1989, Least squares learning and the dynamics of hyperinflation, in: W.A. Barnett, J. Geweke, and K. Shell, eds., Economic complexity: Chaos, sunspots, bubbles and nonlinearity (Cambridge University Press, Cambridge).

Marimon, R. and S. Sunder, 1993, Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence, Econometrica 61, 1073–1108.

Marimon, R., E. McGrattan, and T.J. Sargent, 1989, Money as a medium of exchange in an economy with artificially intelligent agents, Journal of Economic Dynamics and Control 14, 329–373.

Marimon, R., S. Spear, and S. Sunder, 1993, Expectationally driven market volatility: An experimental study, Journal of Economic Theory 61, 74–103.

Miller, J.H., 1989, The coevolution of automata in the repeated prisoner's dilemma, Working paper no. 89-003 (Santa Fe Institute, Santa Fe, NM).

Miller, J.H. and J. Andreoni, 1990a, Auctions with adaptive artificially intelligent agents, Working paper no. 90-01-004 (Santa Fe Institute, Santa Fe, NM).

Miller, J.H. and J. Andreoni, 1990b, A coevolutionary model of free riding behavior: Replicator dynamics as an explanation of the experimental results, Working paper (Santa Fe Institute, Santa Fe, NM; University of Wisconsin, Madison, WI).

Samuelson, P.A., 1958, An exact consumption-loan model of interest with or without the social contrivance of money, Journal of Political Economy 66, 467–482.

Sargent, T.J., 1987, Dynamic macroeconomic theory (Harvard University Press, Cambridge, MA).

Sargent, T.J., 1993, Bounded rationality in macroeconomics (Oxford University Press, Oxford).

Van Huyck, J., R. Battalio, and R. Beil, 1990, Tacit coordination games, strategic uncertainty, and coordination failure, American Economic Review 80, 234–248.

Van Huyck, J., R. Battalio, and R. Beil, 1991, Strategic uncertainty, equilibrium selection principles, and coordination failure in average opinion games, Quarterly Journal of Economics 106, 885–910.

Wallace, N., 1980, The overlapping generations model of fiat money, in: J.H. Karaken and N. Wallace, eds., Models of monetary economics (Federal Reserve Bank of Minneapolis, Minneapolis, MN) 49–82.

Woodford, M., 1990, Learning to believe in sunspots, Econometrica 58, 272–307.