Macroeconomic Dynamics, **2**, 1998, 22–48. Printed in the United States of America.

LEARNING AND THE STABILITY OF CYCLES

JAMES BULLARD
Federal Reserve Bank of St. Louis

JOHN DUFFY
University of Pittsburgh

We investigate the extent to which agents can learn to coordinate on stationary perfect-foresight cycles in a general-equilibrium environment. Depending on the value of a preference parameter, the limiting backward (direction of time reversed) perfect-foresight dynamics are characterized by steady-state, periodic, or chaotic trajectories for real money balances. We relax the perfect-foresight assumption and examine how a population of artificial, heterogeneous adaptive agents might learn in such an environment. These artificial agents optimize given their forecasts of future prices, and they use forecast rules that are consistent with steady-state or periodic trajectories for prices. The agents' forecast rules are updated by a genetic algorithm. We find that the population of artificial adaptive agents is able eventually to coordinate on steady state and low-order cycles, but not on the higher-order periodic equilibria that exist under the perfect-foresight assumption.

Keywords: Learning, Multiple Equilibria, Coordination, Genetic Algorithm

1. INTRODUCTION

1.1. Overview

The possibility of multiple stationary perfect-foresight equilibria in certain classes of general equilibrium models is now well established.[1] Some of these stationary equilibria may be steady states, but others can be periodic or even chaotic. Standard general-equilibrium theory makes no prediction as to which of these stationary equilibria might be achieved in economies driven by actual human behavior. Many economists are trying to make sense of this situation, often by arguing that some equilibria are more likely to eventually be observed than others, and therefore that these equilibria are the more relevant ones for making predictions based on the model. One common approach involves replacing the perfect-foresight assumption with an adaptive learning scheme to determine which of the multiple stationary-equilibrium trajectories are stable under the learning dynamics.[2] However, the use of learning as an equilibrium selection device often has been limited to examples

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of St. Louis or the Federal Reserve System. We thank the referees for helpful comments and suggestions.
Address correspondence to: James Bullard, Research Department, Federal Reserve Bank of St. Louis, 411 Locust Street, St. Louis, MO 63102, USA; e-mail: bullard@stls.frb.org.

© 1998 Cambridge University Press 1365-1005/98 $9.50

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

in which the degree of multiplicity of stationary equilibria is not too severe. Moreover, in cases in which stationary periodic equilibria have been considered, the learning analyses to date typically have been strictly local; the question of which among many equilibria would be selected under a plausible global learning analysis remains largely open.[3]

We implement a certain type of global learning dynamic in a model that sometimes possesses coexistent steady-state, periodic, and chaotic perfect-foresight equilibria. Our technique involves analysis of the dynamics generated by a population of artificial adaptive agents. The environment is Grandmont's (1985) endowment overlapping-generations economy, in which there is a constant supply of fiat currency and consumption in the first and second periods of life is not necessarily a gross substitute. The model supports two steady-state equilibria: the Pareto inferior autarchic steady state, which is characterized by a zero demand for fiat currency; and the Pareto optimal monetary steady state, which is characterized by a positive demand for fiat currency. Under time-separable preferences, and provided that the value of the coefficient of relative risk aversion for the old agents is high enough and that of the young agents is low enough, it can be shown that stationary perfect-foresight equilibria also may exist in which the equilibrium dynamics are characterized either as periodic or chaotic trajectories for real money balances, and that these complicated stationary equilibria are also Pareto optimal. However, in such cases, the forward perfect-foresight dynamics are not well defined, so that it is not clear that there is a meaningful forward dynamic that could plausibly be assumed to converge to one of these more complicated stationary equilibria. Thus, although complicated perfect-foresight equilibria may exist in this model, whether or how such equilibria might be achieved is an open question. Grandmont (1985) provided a complete analysis of the model under the well-defined, backward (direction of time reversed) perfect-foresight dynamics, and demonstrated that the limit of these backward dynamics could be a complicated stationary equilibrium.

Our approach is to create a meaningful *forward dynamic* by relaxing the perfect-foresight assumption and examining how a heterogeneous population of artificial adaptive agents might learn to forecast future prices in such an environment. Agents are differentiated (in addition to birth dates) by the forecast rule that they employ, and each agent solves an optimization problem based on an individual-specific forecast for the future price. Agents' forecast rules are updated by a *genetic algorithm*, a population-based, stochastic, directed search algorithm that generates new rules while retaining and improving upon those rules that have performed well in the past.[4] We conduct computational experiments with economies defined in this way and report the results.

We interpret genetic-algorithm learning as a powerful representation of trial-and-error learning which has important advantages over many other models in the literature. Among these are that (1) beliefs are initially heterogeneous across agents, (2) the information requirements on agents are minimal, (3) the genetic algorithm offers a natural model for experimentation by agents with alternative forecast rules, (4) the heterogeneity of beliefs allows parallel processing to be an

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

important feature of the economy, (5) genetic-algorithm learning has been shown in other research to successfully mimic the behavior of human subjects in controlled laboratory settings, (6) the learning model can be applied even in complicated problems, and finally (7) the initial heterogeneity of the population allows us to initialize the system randomly, so that we are able to obtain some sense of the global properties of our system under learning as opposed to the local analysis that often is employed in the learning literature.[5] These features suggest to us that genetic-algorithm-based models of learning have interesting features from an economic standpoint, and that such a model of learning will allow us to address the question of what type of behavior reasonably might be expected to arise in an actual economy or economic situation with a pervasive multiplicity of stationary equilibria.

We report the results of a large number of computational experiments. These experiments are organized around a single preference parameter: the coefficient of relative risk aversion of the old agents (those in the second period of life) which we allow to vary. Following Grandmont (1985), the relative risk aversion of the young agents (those in the first period of life) remains unchanged in all of the experiments. The main finding is that the population of artificial adaptive agents is able to coordinate their initially heterogeneous beliefs so as to implement steady-state or low-order periodic equilibria for real money balances. For low values of the preference parameter, the artificial agents always coordinate on the monetary steady state; in this case, complicated stationary equilibrium trajectories do not exist. However, as the preference parameter is increased, and the multiplicity of stationary perfect-foresight equilibria becomes more pronounced, the population of artificial agents frequently fails to coordinate on the cyclic equilibria that are picked out by the limiting backward perfect-foresight dynamics. Instead, coordination occurs on steady-state and periodic equilibria that are of lower periodicity than the equilibria predicted by the limiting backward perfect-foresight dynamics. Thus, our long-run outcomes under learning tend to be simpler (lower order) than outcomes under backward perfect foresight. However, we also find that the time it takes agents to achieve coordination tends to increase with the relative risk aversion of the old agents over a large portion of the parameter space. In addition, we find that, when cycles exist, our systems can display qualitatively complicated dynamics for long periods of time before eventually converging to relatively simple, low-periodicity equilibria.

We reach two main conclusions based on these results. First, in the economies we study, the initially heterogeneous beliefs of the agents eventually give way, leading to a situation in which all agents coordinate on a perfect-foresight equilibrium. Thus, our results suggest that *agents can achieve coordination even in relatively complicated situations*. Second, the stationary equilibria on which agents coordinate are always relatively simple—either a steady state or a low-order cycle—and are always Pareto optimal. Our results suggest that *it is difficult for an economy comprised of optimizing agents with initially heterogeneous beliefs to coordinate on especially complicated stationary equilibria*, such as those characterized by high-order periodicities, even though such coordination is a distinct possibility in

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

this model a priori. On the other hand, we are unable to rule out coordination on periodic equilibria entirely, and we find that the transient dynamics of our systems can be persistent and qualitatively complicated if (and only if) stationary periodic equilibria exist.

We note that Grandmont's (1985) example economy and other related examples have been criticized for being empirically implausible.[6] Our focus, however, is on whether a population of initially heterogeneous adaptive agents is capable of achieving coordination in a general-equilibrium environment with a pervasive multiplicity of equilibria. For this purpose, we find Grandmont's (1985) example economy especially useful. The issue of whether this economy is empirically plausible is not one that we address. In this sense, our exercise is in the spirit of other studies of coordination problems, for example the Van Huyck et al. (1994) study discussed below.

1.2. Recent Related Literature

Several recent studies have focused on certain aspects of learning in general-equilibrium models closely related to those studied in this paper. Woodford (1990) analyzes the local stability of stationary sunspot equilibria in an overlapping-generations economy in which the sunspot variable follows a two-state Markov process. All agents in this economy use the same forecast rule, which is a version of the Robbins-Monro stochastic approximation algorithm. Woodford provides conditions under which a stationary sunspot equilibrium is (locally) an attractor under this learning algorithm. Our results amplify Woodford's conclusion in a closely related model, in that, first, a stationary perfect-foresight periodic equilibrium of order 2 can be an attractor on the basis of our more global learning algorithm, and second, these deterministic periodic equilibria can be viewed as limiting cases of the two-state stationary sunspot equilibria studied by Woodford. Whereas Woodford's stability conditions only generally apply to two-state, stationary sunspot equilibria, our results go further, suggesting that higher-order stationary periodic equilibria (with periodicity greater than 2) are unlikely to be attractors. That is, although equilibria characterized by high-order periodicities in principle may be locally stable in an analysis such as Woodford's, they evidently have negligible basins of attraction relative to the space of definition of the model in our analysis.

In a related study, Guesnerie and Woodford (1991) find conditions under which cycles will be stable under a learning dynamic in a one-step-ahead forward-looking model. In their analysis, all agents form expectations according to a simple adaptive rule in which the expected price is a convex combination of the actual and expected prices k periods in the past. Guesnerie and Woodford (1991) conclude that cycles of any order might be attractors under the learning dynamics in the systems they study, provided a condition on the map describing the equilibrium conditions is met. When this condition is not met, and such a possibility is shown always to exist, depending on the nature of the equilibrium map, their systems might converge to *learning equilibria* of the type studied by Bullard (1994), limiting learning dynamics that are

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

stationary but that are not perfect-foresight equilibria and which therefore involve forecast errors in perpetuity. The Guesnerie and Woodford (1991) study suggests that a priori, the possibilities for the limiting learning dynamics in our system are quite open. However, our learning model is considerably more complex than one in which there is a simple adaptive learning rule used by all agents.

In a similar vein, Grandmont and Laroque (1986, 1991) as well as Grandmont (1985) study general differentiable learning rules in one-step-ahead systems and find conditions under which these systems may or may not converge to cycles and steady states.[7] Grandmont and Laroque (1991) argue that it is easy to construct examples in which the steady state of an economic model is locally unstable under reasonably defined learning dynamics, whereas Grandmont and Laroque (1986) and Grandmont (1985) suggest that periodic equilibria, perhaps complicated, can be the ultimate outcome of similarly defined systems.

Evans and Honkapohja (1995a) also study one-step-ahead forward-looking systems, both stochastic and deterministic. They imagine that agents all use a certain recursive learning algorithm to compute the expected value of the variable of interest as a function of past observations on that variable. Evans and Honkapohja assume that agents prespecify a value for k, the periodicity of the local equilibrium, so that in each case, agents are assumed to know the periodicity of the equilibrium they are attempting to learn. The analysis makes use of the theory of recursive stochastic systems associated with Ljung (1977) and imported into economic problems of this type by Marcet and Sargent (1989). Evans and Honkapohja relate their necessary and sufficient conditions for the stability of equilibria under the recursive learning dynamic to the concept of *expectational stability* (E-stability) as formulated by Evans (1989).[8] The Evans–Honkapohja conditions suggest that local stability under learning will depend on the properties of the map describing the equilibrium conditions, and that in principle it is possible to have many equilibria that are (strongly E-) stable in the learning dynamics. Thus, although recursive learning might succeed in reducing the set of equilibria that may be regarded as plausible, in general it cannot suggest a unique equilibrium. However, the case studied in the present paper is somewhat special in this regard in that Grandmont's (1985) example involves a unimodal map with a negative Schwartzian derivative, which in turn implies that a single periodic equilibrium among many will be (strongly E-) stable in the learning dynamics according to the Evans–Honkapohja analysis. The monetary steady state also will be (weakly E-) stable in the learning dynamics in this case, so that Evans and Honkapohja conclude that the equilibrium actually observed would depend on the exact specification of the learning rule used by the agents, in particular whether the rule is consistent with a steady state or with a stationary periodic equilibrium of order k.[9] In our model, agents' learning rules are always consistent with steady state and stationary periodic equilibria of many different orders.

We know of no research that analyzes stationary periodic equilibria such as those we are interested in under genetic-algorithm learning. However, research on genetic-algorithm learning by Arifovic (1996) and Arifovic and Eaton (1995) does

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

focus on coordination problems that can be severe. Arifovic (1996), for example, analyzes genetic-algorithm learning in the Kareken and Wallace (1981) model of exchange rates, in which the exchange rate is indeterminate in the sense that any fixed exchange rate $e \in (0, \infty)$ is a rational-expectations equilibrium. Arifovic (1994, 1995, 1996) also shows that the dynamics from experiments with human subjects compare favorably with the dynamics generated by genetic-algorithm learning in a number of different economies.

Although experiments with human subjects are beyond the scope of this paper, we are enthusiastic about testing the predictions of our genetic-algorithm learning model against such a benchmark. Van Huyck et al. (1994) have performed an experiment with human subjects in a repeated coordination game that shares some of the properties of the model studied in this paper. In particular, their coordination game has two strict, pure-strategy Nash equilibria, one of which is an interior equilibrium. This interior equilibrium is shown to be stable under myopic best-response learning dynamics, provided that the game's payoff parameter falls within a certain range; otherwise, the myopic best-response dynamic leads to periodic or chaotic trajectories that cycle about the interior equilibrium. The authors consider an experimental treatment where the payoff parameter lies either within the range that ensures stability of the interior equilibrium under the myopic best-response dynamic or lies far enough outside this range so that the myopic best-response dynamic converges to a period-7 cycle. They find that, in contrast to the predictions of the myopic best-response dynamic, subjects always learn over time to coordinate on the interior equilibrium, even in the parameterization in which the best-response dynamics would imply that the interior equilibrium is unstable (and the period-7 cycle is stable). We view these laboratory results as complementary to the findings that we present in this paper, because in both papers agents fail to coordinate on relatively high-order cycles. However, we differ from Van Huyck et al. (1994) in that the environment we consider sometimes displays a dense set of periodic equilibria. We also note that in the overlapping-generations environment that we study, the complicated trajectories that agents are attempting to coordinate on are *equilibrium* trajectories, whereas the period-7 cycle in the Van Huyck et al. (1994) experiment arises from the use of the disequilibrium, best-response learning dynamic.

Experiments with human subjects that have been conducted in an overlapping-generations environment similar to the one studied in this paper include Marimon and Sunder (1993, 1994, 1995), Marimon et al. (1993), Lim et al. (1994), and Arifovic (1995), with the latter study being the most relevant to the focus of this paper. Marimon et al. (1993) study a two-period overlapping-generations economy with preferences and endowments such that a period-2 perfect-foresight equilibrium coexists with the monetary steady state.[10] Subjects make predictions of future prices, and, analogous to our setup, the optimal level of real money balances is computer generated for each subject, given their price prediction. The number of subjects in a generation ranged from two to five, whereas the number of new generations (iterations) ranged from 27 to 67. The design of their experiment is somewhat different from ours in that the authors systematically varied the size

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

of the incoming young generation as a means of inducing real variation in the equilibrium price level for the first phase of the experiment, and then subsequently kept the generation size constant to see if the subjects might learn to coordinate on the period-2 equilibrium once cyclicality in expectations had been established. In the experiments with changing population sizes (real shocks), they also employed a sunspot variable by alternating the color of some of the information presented on subjects' computer screens in a manner that was perfectly correlated with the real shock. The sunspot variable was left on at all times, even when the real shock was turned off, and agents were not informed of the presence of real shocks. The authors' main finding is that in treatments in which there were real shocks in the initial phase and sunspots throughout, prices continued to display marked periodicity of order 2 even after the real shock was turned off. However, in treatments without real shocks in the initial phase (but with sunspots), subjects always learned to coordinate on the monetary steady state. The authors conclude that it is possible for human subjects to learn to coordinate both on steady state and on period-2 equilibria. These results match up well with our own, because our genetic-algorithm-based learning model predicts that both steady state and period-2 equilibria will be observed as long-run outcomes in a similarly specified economy. Furthermore, their time series for the price level [Marimon et al. (1993, Fig. 3, p. 89)] are qualitatively similar to those generated in our computational experiments, even though our systems had neither real shocks nor a sunspot variable. Finally, these authors argue that equilibrium selection appears to be *path dependent*, because periodic price dynamics were observed only when there were real shocks in the initial phase. We note that our model also predicts that equilibrium selection will be path dependent. By conducting a large number of computational experiments, we are able to let conditions develop spontaneously that in certain instances led our system to converge to a period-2 equilibrium, whereas Marimon et al. (1993) needed inducements in the form of the real shocks to avoid conducting a prohibitively large number of experiments.[11] Indeed, the major advantage of our computational approach is that we are able to conduct many more experiments with many more iterations across a wider variety of overlapping-generations economies, thus enabling us to obtain a more complete picture of the long-run outcomes that might reasonably be expected in these economies.

Computational experiments with evolutionary algorithms (such as genetic algorithms) in economic settings have been suggested by Holland and Miller (1991) and Sargent (1993) and have been performed by Axelrod (1987), Marimon et al. (1990), Rust et al. (1994), Routledge (1995), Arthur et al. (1996), Ho (1996), and Miller (1996) among others. These applications generally have not been in competitive general-equilibrium environments like we have in mind. An important exception is the work of Arifovic (1994, 1995, 1996).

2. ENVIRONMENT

The example economy is taken from Grandmont (1985). Time is discrete with integer $t \in (-\infty, \infty)$. There is a single, perishable consumption good and a constant

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

supply of fiat money $M > 0$. There are $N/2$ agents per generation, where N is the total population alive at time t. Agents live for two periods, and can save between periods by holding fiat money. Agents receive strictly positive endowments of the consumption good in each period of life, $\{e_1, e_2\}$. Preferences are given by

$$U = \frac{c_t(t)^{1-\rho_1}}{1-\rho_1} + \frac{c_t(t+1)^{1-\rho_2}}{1-\rho_2}, \tag{1}$$

where $\rho_1, \rho_2 \in (0, \infty)$ denote the coefficient of relative risk aversion of the young and old agents, respectively. The notational convention is that subscripts denote birthdates and real time is recorded in parentheses, so that $c_i(j)$ denotes the time j consumption of the agent born at time i.[12] The agent born at time t maximizes utility subject to the pair of budget constraints

$$c_t(t) \leq e_1 - s_t(t), \tag{2}$$

$$c_t(t+1) \leq e_2 + \frac{P(t)}{P(t+1)} s_t(t), \tag{3}$$

where $s_t(t)$ denotes the amount the agent born at time t chooses to save at time t, and $P(t)$ denotes the time-t price of the consumption good in terms of fiat money. Combining the first-order conditions with the two budget constraints, one obtains

$$c_t(t) + c_t(t)^{(\rho_1/\rho_2)} \beta(t)^{[(\rho_2-1)/\rho_2]} = e_1 + e_2 \beta(t), \tag{4}$$

where $\beta(t) = [P(t+1)]/P(t)$ denotes the gross inflation factor between dates t and $t+1$. It follows from the compactness of the budget set and the strict concavity of the utility function that the young agent's consumption decision, $c_t(t)$, and therefore his savings decision, $s_t(t)$, are *uniquely* determined. Because a closed-form solution for $c_t(t)$ is unavailable, we instead rely on numerical methods to obtain $c_t(t)$. Once the optimal consumption and savings amounts are determined, it becomes possible to define a perfect-foresight equilibrium sequence for prices (equivalently, for real money balances) using the market clearing condition

$$S(t) = M/P(t), \tag{5}$$

where $S(t)$ is aggregate savings at time t. Using this condition, and the fact that the supply of fiat money M is constant, one can derive a first-order difference equation that characterizes all perfect-foresight equilibria in this economy:

$$P(t) = \frac{S(t+1)}{S(t)} P(t+1). \tag{6}$$

Following Grandmont (1985), let us write this difference equation more compactly as

$$P(t) = \Phi[P(t+1)]. \tag{7}$$

A *perfect-foresight equilibrium* is defined as any sequence of prices $\{P(t)\}$ that satisfies equation (7). A *steady-state* equilibrium is a price level $\bar{P}$ such that $\bar{P} = \Phi(\bar{P})$. It is well known [see, e.g., Gale (1973)] that in this model, there

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

can be, at most, two steady-state equilibria, one in which aggregate savings is zero and agents consume endowments, and possibly another in which aggregate savings is positive and the steady-state gross rate of return on savings and its reciprocal, the steady-state gross inflation factor, are both equal to unity. In the analysis that follows, we choose an endowment sequence that guarantees that this latter equilibrium exists, that is, that we have the *Samuelson case* in Gale's terminology.

As Grandmont (1985) demonstrates, these two steady states are not the only stationary equilibria that this economy may possess. If the offer curve is sufficiently backward bending, or equivalently, if the coefficient of relative risk aversion of the old agents, ρ_2, is large enough and the coefficient of relative risk aversion of the young agents, ρ_1, is less than or equal to unity, then in addition to steady states, the set of equilibria also may include periodic as well as chaotic equilibria.[13] Our focus is limited to the steady-state and periodic equilibria that may arise in this environment. A *periodic* equilibrium of order k consists of a sequence of k prices, $\{\bar{P}_1, \bar{P}_2, \ldots, \bar{P}_k\}$, such that $\bar{P}_j = \Phi^k(\bar{P}_j)$, $j = 1, 2, \ldots, k$, where Φ^k denotes the kth iterate of the map Φ.

Grandmont (1985) noted that the forward perfect-foresight dynamics, that is, iterates of the map (7), may not be defined uniquely, depending on the properties of the map $\Phi(\cdot)$. Grandmont (1985) studied this system by limiting attention to the *backward perfect-foresight dynamics*, that is, sequences of prices that solve the map (7) with the direction of time reversed. A periodic equilibrium of order k in the backward perfect-foresight dynamics is a sequence of k prices $\{\bar{P}_k, \bar{P}_{k-1}, \ldots, \bar{P}_1\}$ of the map Φ such that $\bar{P}_j = \Phi^{-k}(\bar{P}_j)$, $j = 1, 2, \ldots, k$. Grandmont's (1985) well-known main result was to show that periodic equilibria of any order and chaos could exist as long-run outcomes in the backward perfect-foresight dynamics without abandoning the classical assumptions of utility maximization and market clearing.[14] At least for periodic equilibria, this is enough to prove existence of complicated equilibrium dynamics in the more meaningful forward perfect-foresight system, because even though the forward dynamics are not uniquely defined, it is, at least in principle, possible that agents would choose appropriately in the forward system so as to replicate the trajectory uniquely defined in the backward perfect-foresight dynamics.

In Grandmont's (1985) model, when periodic equilibria exist, they coexist with other equilibria, and in particular, they coexist with steady states. Sarkovskii's (1964) theorem implies that if a periodic equilibrium of order 3 exists, then equilibria of every other order q, where q is an element of the set of positive integers, also exist. This raises the question of which among these many stationary equilibria might be achieved, and Grandmont (1985) addressed this question by considering the limiting backward perfect-foresight dynamics and arguing that plausible learning rules exist that would, if used by all agents, cause the system to be locally convergent to the periodic equilibria isolated under backward perfect foresight. In particular, Grandmont (1985, Proposition 3.2, p. 1012) interpreted a stationary equilibrium that was stable in the backward perfect-foresight dynamics as stable in the forward learning dynamics, and Grandmont gave conditions on a differentiable learning rule under which such statements would be true.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

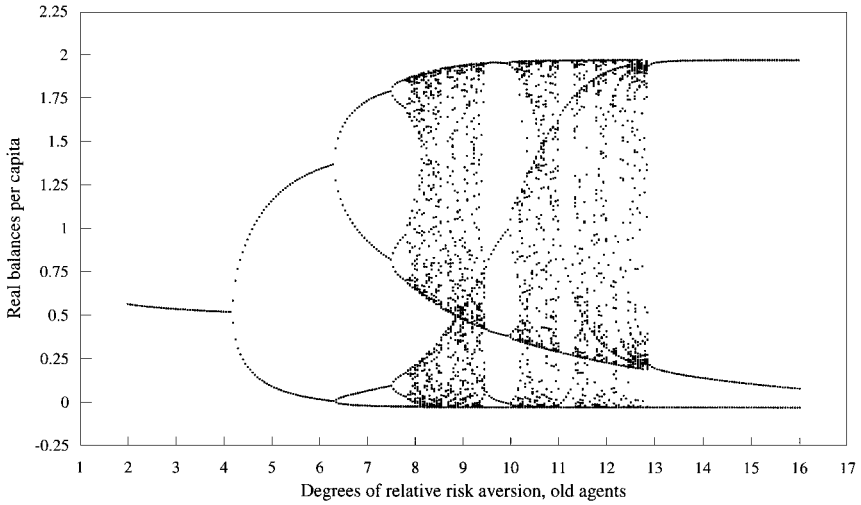


FIGURE 1. Bifurcation diagram, backward perfect-foresight dynamics (last 50 of 1,000 iterations).

In Figure 1, we replicate Grandmont's (1985, p. 1030) Figure 4, which is a bifurcation diagram in the backward perfect-foresight dynamics. In this example, which we use throughout the remainder of the paper, $\{e_1, e_2\} = \{2, 0.5\}$, and the relative risk aversion of the young is fixed at $\rho_1 = 0.5$. The figure shows a plot of real money balances per capita, which can range between zero and the first-period endowment of 2. For each value of the parameter representing the relative risk aversion of the old agents, ρ_2, the backward perfect-foresight map was iterated 1,000 times, using as an initial condition the peak of the unimodal map. The last 50 of the 1,000 iterations are plotted in the diagram. As the relative-risk-aversion parameter is increased, a standard period-doubling bifurcation pattern emerges, with the limiting backward perfect-foresight dynamics being the monetary steady state for low values of the risk-aversion parameter and a periodic equilibrium of order 3 for high values of the risk-aversion parameter. Between these extremes, the limiting dynamics involve cycles of mostly higher order, or perhaps of a very high order or even chaos, in which case the plot of the last 50 iterations is inadequate to characterize the equilibrium. Qualitatively, however, the figure is clear.

We now want to consider relaxing the perfect-foresight assumption and introducing genetic-algorithm learning into this system.

3. LEARNING

3.1. Representation of a Heterogeneous Population

At every date t, there is a finite population of N agents participating in the economy. The population is the same at every date and is divided equally between agents born at time t and those born at time $t-1$, so that the size of each generation is $N/2$.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

Each agent in the population is differentiated by a *forecast rule* for next period's price level. This forecast rule completely characterizes the agent's behavior in both periods of life, because the agent takes the optimal action (makes the optimal savings decision) based on the forecast of next period's price produced by his own forecast rule. This forecast rule is represented by an integer value $k \in [0, \bar{k}]$. The rule is to use the price level that was realized $k+1$ periods in the past as the forecast of next period's price level. For example, consider a young agent i, born at time t, who has chosen k_i. If we let $F_t^i[P(t+1)]$ denote agent i's time-t forecast of the price level expected to prevail at time $t+1$, we can write agent i's forecast rule as

$$F_t^i[P(t+1)] = P(t - k_i - 1). \tag{8}$$

This specification of the forecast rule is simple, but has the important implication that each young agent i can make optimal savings decisions in any type of stationary equilibrium of order q, so long as $q \in [0, \bar{k}+1]$. In fact, we chose this specification precisely because it allows agents to adopt behavior consistent with steady state or periodic trajectories for prices up to the limit $\bar{k}+1$.[15] We note that this class of forecast rules is more general than might first appear. Although one could allow agents to use past data in more complicated ways, perhaps even econometrically sophisticated ways, for our exercise to be interesting the forecast rule must ultimately allow the agent to behave in a way consistent with any of the equilibria of the model (up to some limit). In particular, if the economy is in an equilibrium with periodicity $q \leq \bar{k}+1$, the chosen forecast rule must deliver a forecast consistent with that equilibrium.[16] If it does not, then the agents cannot coordinate on that equilibrium. For this reason, we think that it is reasonable to abstract from the question of *how* agents manipulate the data and concentrate on specifying a class of rules such that agents in principle can coordinate on any of the equilibria of the model up to some limit.[17]

To apply the genetic algorithm, we must first encode each agent's forecast rule as a string of length ℓ, with elements of the string chosen from the binary $\{0, 1\}$ alphabet. For example, if $\ell = 8$, the string might be 00100101. This string value can be decoded to the base 10 integer 38, which is the k-value that this agent uses in forecasting future prices. Thus, this agent would use the *actual* price level that was realized $k+1 = 39$ periods ago [i.e., $P(t-39)$] as his forecast of next period's price level. Of course, if the economy had been in a steady state for each of the past 39 periods, then this forecast rule would give the same predicted price as the string 00000000. In fact, all strings representing integer values $k \in [0, \bar{k}]$ are consistent with a steady-state equilibrium, and many are also consistent with periodic equilibria of other orders. The nature of this situation is described in more detail in Table 1.

3.2. How Forecasts Are Used

Given the young agent i's time-t forecast of the price level expected to prevail at time $t+1$, $F_t^i[P(t+1)]$, we can denote this same agent's forecast of the *gross*

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

TABLE 1. Consistency of strings with stationary equilibria

Bit string	k	Consistent with
00000000	0	steady state
00000001	1	steady state, 2-cycle
00000010	2	steady state, 3-cycle
00000011	3	steady state, 2-cycle, 4-cycle
00000100	4	steady state, 5-cycle
00000101	5	steady state, 3-cycle, 6-cycle
...	...	...
11111111	255	steady state, 2-cycle, ..., 256-cycle

inflation factor between dates t and $t+1$ by

$$\beta_t^i(t) = \frac{F_t^i[P(t+1)]}{P(t)}. \tag{9}$$

Using this forecast for gross inflation, the computer algorithm then numerically solves equation (4) for agent i's consumption amount $c_t^i(t)$. The young agent i's savings decision is given by $s_t^i(t) = e_1 - c_t^i(t)$. Aggregate savings at time t is therefore given by

$$S(t) = \sum_{i=1}^{N/2} s_t^i(t). \tag{10}$$

The law of motion for the price level $P(t)$ in the *forward* dynamics of our learning model is given by

$$P(t) = \frac{S(t-1)}{S(t)} P(t-1), \tag{11}$$

and thus $P(t)$ is determined once $S(t)$ is known. However, the quality of the time-t young agents' forecasts cannot be evaluated until period $t+1$, when $S(t+1)$ and thus $P(t+1)$ is known. We now turn to a discussion of how we evaluate and update forecast rules.

3.3. Updating Forecast Rules Using a Genetic Algorithm

Although both the genetic algorithm and the overlapping-generations model have some rather obvious biological overtones, we do not interpret our model of learning as one of a strict passing of genetic information from one generation to the next. Instead, we view the genetic algorithm as a convenient device with which to model communication among individual agents alive at date t.[18] We imagine that agents who will enter the productive portion of their life at time $t+1$ (that is, agents who will be young next period) undergo a *social learning process* prior to time $t+1$ that involves having conversations with some of the older agents alive in the population at time t.[19] The soon-to-be-born agents devise their own forecast strategies based on

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

these conversations. The result of this process is a new generation of agents, which we refer to as the newborns, who become the next generation of young agents once the agents in the model are aged appropriately. Because we create the newborn generation at time t and allow these newborn agents to enter their productive lives at time $t + 1$, we can make use of the entire set of genetic information available at time t to create the new generation. That is, we can allow the newborns to converse with any of the agents participating in the economy, young or old, at time t. Accordingly, we subject the entire population of N strings available at time t to the four genetic operators that make up the genetic algorithm, and create the newborn generation from this information. We now describe our implementation of these four operators: reproduction, crossover, mutation, and election.

3.3.1. Reproduction. To create the newborn generation at time t, we begin by calculating the *fitness* of each string in the entire population of N agents alive at time t. First, each string is decoded to determine its forecast rule. Next, the forecast rule is applied using the history of prices available through time $t - 1$, and a price forecast is determined. The reason for restricting the forecast rule to the history of prices available through time $t - 1$ is that this restriction enables us to assess how well the forecast rule *would have performed* had it been in use in the previous period. Performance is measured in terms of the two-period, lifetime utility that the agent would have attained had the agent been young in the previous period and possessed the forecast rule in question. In assessing the lifetime utility value, we use the most recent realization of the gross inflation rate $\beta(t-1) = P(t)/P(t-1)$. Thus, the lifetime fitness value is a measure of the accuracy of a particular forecast rule; the higher the lifetime fitness value, the more accurate is the forecast rule that was used.

Once the fitness values have been determined, we conduct one *selection tournament* for each of the $N/2$ newborns that we wish to create. Each tournament consists of a random selection of two strings, with replacement, from the entire population of N strings alive at time t. The string with the higher fitness value wins the tournament and a copy of this string is included in the group of newborn strings.[20] The result of this selection process is a set of $N/2$ newborn strings with fitness values that are, on average, higher than those of the population alive at time t. Although we now have a set of newborn strings, we are not yet finished because we want to allow these newborn agents to experiment with some new possible forecast rules that may not be part of the genetic-information set available at time t. The next three operators accomplish the task of introducing new information into the system.

3.3.2. Crossover, mutation, and election. The main idea of the crossover and mutation operators is to create new forecast rules by mixing portions of existing strings and by changing individual bits with small probability. Accordingly, we randomly pair the newborn strings. Then for each pair we perform the crossover operation: With some probability $p^c > 0$, we divide the pair of strings at some

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

randomly chosen point, $s \in [1, \ell - 1]$, where ℓ is the length of each bit string, and swap the bits to the right of the crossover point. With probability $1 - p^c$, we do not apply the crossover operation to the pair of strings. As an example, suppose that we have a pair of strings of length $\ell = 8$ and that crossover is to be performed on these two strings. Suppose also that the randomly chosen point is $s = 6$. The pair of strings is divided at this point:

$$010100 \mid 11$$

$$111001 \mid 01.$$

The portion of the strings to the right of the dividing point then is swapped, creating two new strings:

$$01010001$$

$$11100111.$$

The two resulting strings (even in cases in which the two strings were not changed by crossover) then are subjected to some mutation: We consider each bit value $b = 0, 1$ of the two strings that result from the crossover operation, and with some probability $p^m > 0$, we replace the bit b with the bit $1 - b$. With probability $1 - p^m$, the bit value is not mutated.

We refer to the strings that result from crossover and mutation as *alternatives*, because we think of the newborns (strings selected through reproduction) as contemplating the adoption of some alternative forecast rules that may not be in use in the current population. In the standard, biological application of the genetic algorithm, these alternatives simply would become the newborn generation, but such an approach is less than satisfactory in economic applications in which agents are appropriately viewed as making intelligent choices. In particular, the standard application of the genetic algorithm would allow our alternative strings to enter the population even if they were not strictly better (in terms of fitness) than the newborn strings from which they were created. However, we want to think of our economic decisionmakers as being somewhat less naive. Therefore, we adopt Arifovic's (1994) *election* operator: Each pair of newborns is compared with the corresponding pair of alternatives. Of the four strings, only the two with the highest fitness values are allowed to remain in the generation of newborns, and the other two are discarded. Fitness of the alternatives is computed in the same way as it is in the reproduction operator. The crossover, mutation, and election operators combined can be interpreted as a method of allowing new forecast rules to enter the system without forcing agents to adopt rules that are unlikely to yield high utility.[21]

The reproduction, crossover, mutation, and election operators have a simple economic interpretation. Being "born" in this economy means leaving one's formative years and entering the productive portion of one's life. These newborn agents just leaving their formative years initially have no plans for the future—they are blank

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

slates. They acquire the forecast rule that they will need by communicating with a few other members of society, those either one or two generations ahead of them. This communication is modeled via the reproduction operator. In our implementation, each newborn agent communicates with two randomly selected members of society.[22] The newborns evaluate the forecast rules that belong to these two older agents by calculating how much utility the rules would have delivered had they been in use one period in the past. Each newborn then copies the forecast rule of the two that would have delivered the most utility. This completes the first step in attaching a forecast rule to each of the incoming members of the society, but the newborns communicate further when they talk with each other and contemplate alternative forecast rules that might not be in use in the society at that time—that is, the newborns conduct a mental experiment with other possible forecast rules. This additional communication is modeled via the crossover and mutation operators. In our implementation, the newborns are paired and each pair creates two alternative forecast rules by combining parts of their existing rules into two new possibilities, and also by randomly changing small parts of the recombined forecast rules. The two alternative forecast rules are evaluated to see how much utility they would have delivered had they been in use one period in the past. Alternatives that improve upon the newborns' existing forecast rules are adopted, whereas alternatives that do not are discarded. This last step is modeled using the election operator. Thus the incoming generation learns from the experience of the agents older than themselves and also can be innovative in introducing new forecast rules into the society.

Once the election operator has been applied, and a generation of newborn agents has been chosen, time advances to the next period. The time-t population of newborn agents, created via reproduction, crossover, mutation, and election, becomes the young generation alive at time $t + 1$. The generation of young agents alive at time t becomes the old generation alive at time $t + 1$. The generation of old agents alive at time t ceases to exist. The forecasts of the time $t + 1$ young generation are determined, their savings decisions are calculated, and the aggregate savings amount $S(t + 1)$ then determines the price level $P(t + 1)$ according to the law of motion for prices,

$$P(t + 1) = \frac{S(t)}{S(t + 1)} P(t).$$

The genetic algorithm then is begun anew. The algorithm ends when either a convergence criterion (described below) is satisfied or the algorithm has reached the maximum number of iterations (generations) allowed.

3.4. Some Advantages of Genetic-Algorithm Learning

We interpret genetic-algorithm learning as a useful model of trial-and-error learning in a heterogeneous-agent economy. This learning model has many important advantages relative to other models in the literature. First, beliefs are initially heterogeneous across agents, a feature not often modeled in the learning literature to date.[23] Second, the information requirements on agents are minimal, because they

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

only need to know their own utility and their own forecast rule in order to make a decision. Third, the genetic algorithm offers a natural model for experimentation by agents with alternative forecast rules, an important characteristic of learning also rarely modeled in competitive general-equilibrium environments in the literature to date. Fourth, the heterogeneity of beliefs allows parallel processing to be an important feature of the economy. That is, some agents are trying one forecast rule whereas other agents are trying other forecast rules, with the better forecast rules propagating and the poorer ones dying out. We think that this is closely akin to what goes on in actual economies, where communication among agents encourages successful strategies to be quickly copied and unsuccessful ones to be discarded. Fifth, the approach to learning that we study can be applied even in complicated problems such as the one studied here. Finally, the initial heterogeneity of the population allows us to initialize the system randomly, so that we are able to obtain some sense of the global properties of our system under learning as opposed to the local analysis that often is employed in the learning literature. We think that these features provide good reasons to investigate the properties of genetic-algorithm-based models of learning in systems in which learning may have an important role to play. In the present model the potentially important role for learning is equilibrium selection.

We chose to use a genetic algorithm, rather than some other learning/search algorithm, because the genetic algorithm has a number of attractive features that are particularly well suited to the problem that we examine. First, unlike other search/learning algorithms, the genetic algorithm is *population based*, and involves parallel processing of a finite set of initially heterogeneous strings; each of these strings can be viewed as representing different candidate solutions to a particular optimization problem, or, as in our interpretation, each string can be viewed as representing the belief of a different agent in the population. The population-based genetic algorithm can be regarded as a global search algorithm whereas representative–agent-type search algorithms are necessarily local. The implicit parallelism of the genetic algorithm also works to reduce computation time as compared with enumerative search strategies and other search algorithms. As we have stressed, we think parallelism is an important feature of coordination in actual economies or economic situations. Second, the genetic algorithm is applied readily to difficult optimization problems sometimes heuristically described as rugged surfaces. Unlike gradient–based learning algorithms [e.g., the Robbins-Monro or the least-squares learning algorithms studied by Woodford (1990) and Marcet and Sargent (1989) or other hill-climbing algorithms], the genetic algorithm does not require the taking of derivatives. This feature makes the genetic algorithm an attractive model of how populations of economic agents might update their forecast rules over time in highly nonlinear environments such as the one we consider here. Third, genetic algorithms are known to behave as excellent function optimizers. Holland (1975) has shown that genetic algorithms optimize on the trade-off between searching for new rules (*exploration*) and utilizing information discovered in the past (*exploitation*).[24] Optimization of this trade-off is

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

important: An algorithm that engages in too much exploration discards, early in the search, feedback information that may prove useful, whereas an algorithm that engages in too much exploitation is prone to local hill climbing and may be overly sensitive to noise. Finally, modeling learning behavior using a genetic algorithm has the advantage that the exact specification of the problem that agents are attempting to solve is known at the outset. Indeed, the representation of the agent's problem is a necessary prerequisite to the application of the genetic algorithm. Thus, interpretation and evaluation of the results from applying the genetic algorithm become relatively straightforward. By contrast, other artificial-intelligence techniques (e.g., artificial neural networks or genetic programming) evolve structures and programs that often are difficult to interpret. For all of these reasons, we chose the genetic algorithm over several alternative methods as a way of modeling learning behavior. This is not to say that other methods are not interesting, but only to note that the genetic-algorithm approach has a number of desirable features.

3.5. Design of Computational Experiments

We conducted a large number (1,410) of computational experiments using the genetic algorithm described above in Grandmont's example economy. We used the same set of parameter values that were used to generate Figure 1; we set $\{e_1, e_2\} = \{2, 0.5\}$, and we set $\rho_1 = 0.5$ in all of our experiments. The experiments were organized around the value of the single preference parameter ρ_2. We initially set ρ_2 equal to 2 and we then increased the value of this parameter up to 16 by increments of 0.1. Thus we considered 141 different values for ρ_2. We set the bit string length $\ell = 8$ which allows the agents to take actions consistent with a periodic equilibrium of an order as high as 256.[25] We set $N = 100$ (so that each generation consisted of $N/2 = 50$ agents), $p^c = 1$, and $p^m = 1/\ell = 0.125$; we chose these values in part on the basis of the optimal values recommended by Grefenstette (1986) and Bäck (1993) and in part because the election operator that we use ensures that strings that emerge from crossover and mutation with particularly low fitness values will not enter the population, and so, little is lost by allowing agents to experiment extensively with alternative strings.

We began each of our 1,410 experiments by randomly generating an initial population of 50 old and 50 young agents (i.e., each bit in each of the 100 strings was set equal to 1 with probability 0.5 and to 0 with probability 0.5). We initialized a set of prices by randomly choosing 256 real prices from a uniform distribution on the unit interval. We then repeatedly applied the genetic algorithm and computed market clearing prices until either convergence was obtained or 2,000 generations had failed to coordinate on an equilibrium (i.e., after 2,000 iterations). Following each iteration of each experiment, we recorded several pieces of information, including the price level, the gross inflation rate, and the mean value of real money balances held by agents in the young generation.

Because, as noted in Table 1, all strings are consistent with a steady-state equilibrium and nearly all strings are consistent with multiple stationary periodic

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

equilibria, we examined the path followed by the actual price level when ascertaining whether our system had achieved convergence on some type of stationary outcome. We checked for convergence following each iteration of the algorithm. In particular, we checked for the presence of cycles of every possible periodicity, and we declared that convergence had been obtained only if the most recent 10 prices matched the 10 corresponding prices that occurred $k + 1$ periods in the past [$k \in (0, 255)$] up to a tolerance of 1×10^{-10}. This criterion[26] was met before the 2,000-iteration limit in most (93%) of our experiments. Our convergence criterion is relatively easy on high-order equilibria, because we check only 10 prices on those cycles, but nevertheless we did not observe high-order equilibria.

Once convergence was obtained or the maximum number of iterations was reached, the same experiment was repeated again nine more times, so that we have 10 experiments for each of the 141 values of ρ_2 for a total of 1,410 experiments. We stress that our convergence criterion is price based, so that the fact that strings are identical does not imply that the system has converged. Indeed, we sometimes found that strings were identical before prices met our convergence criterion.

4. EXPERIMENTAL RESULTS

4.1. Overview

The main finding from our set of computational experiments is that in the economies that we study, the artificial agents eventually learn to coordinate on stationary perfect-foresight equilibria of order $k + 1$, where $k = 0$, 1, or 3, and where $k = 0$ refers to a degenerate cycle, the monetary steady state. None of the economies that we studied displayed coordination on stationary equilibria of any other periodicities. In particular, we never observed coordination on the *non*monetary steady state. We conclude from these results that simplicity may be a virtue for equilibria, because simpler equilibria are more likely to be achieved in our systems.[27] As mentioned in the introduction, our finding that artificial adaptive agents tend to coordinate on relatively simple equilibria is consistent with experimental findings with human subjects in similar environments studied by Van Huyck et al. (1994) and Marimon et al. (1993).

A second, related finding is that our systems under learning rarely chose a single equilibrium from among the many possible equilibria as the only limit point. Instead, for most values of ρ_2, our systems sometimes converged to one stationary equilibrium, the monetary steady state, and sometimes converged to another stationary equilibrium, a period-2 cycle.[28] In one case, the learning system had nearly converged to a period-4 cycle, and in two other cases the system had not converged to any path resembling an equilibrium after 2,000 iterations. We conclude that although the introduction of learning into our economic systems can sharply limit the set of equilibria that might be considered reasonable in the sense that they might be achievable even in systems in which the decisions are made by humans, the introduction of learning does not imply a *unique* stationary

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

equilibrium, and there is no guarantee that convergence will even occur in finite time.

A third broad finding is that, although our systems almost always converged eventually, they sometimes displayed qualitatively complicated dynamics for very long periods of time. By complicated dynamics, we mean that either the price dynamics do not appear to be in a small neighborhood of any discernible equilibrium path, or that they are in such a neighborhood, but the path involves higher-order periodic motion. These complicated transient dynamics sometimes occurred in a region of the parameter space where cycles exist in the backward perfect-foresight dynamics, but never occurred in the region where the monetary steady state was the lone efficient equilibrium. Thus, when the set of equilibria is relatively complicated, the coordination problem is more pronounced and the economy may remain out of equilibrium for long periods of time. Some illustrations of these complicated transient dynamics are provided by Bullard and Duffy (1995).

4.2. Specific Results

In most cases, our systems did not converge to the same stationary equilibrium chosen by the limiting backward perfect-foresight dynamics studied by Grandmont (1985). For low values of ρ_2, in particular those below 4.2, we observed convergence to the monetary steady state in every experiment, and this is the same prediction made by the limiting backward perfect-foresight dynamics (see Figure 1). As ρ_2 is increased further, however, the limiting backward perfect-foresight dynamics display a bifurcation, with the monetary steady state losing stability and never regaining it for values of $\rho_2 \geq 4.2$. In contrast, in our systems with learning, the monetary steady state was always a limit point in at least 1 of the 10 experiments conducted for each different value of ρ_2. In addition, for higher values of ρ_2, and in particular for values past the first bifurcation point in the backward perfect-foresight dynamics (i.e., for $\rho_2 \geq 4.2$), our systems often converged, in at least one experiment, to a period-2 stationary equilibrium, even in cases in which that equilibrium, too, had lost its stability in the backward perfect-foresight dynamics. Other relatively simple periodic equilibria, such as the period-3 cycle that is stable in the backward perfect-foresight dynamics for values of $\rho_2 \geq 13$, appear not to be attractive in our learning dynamics, because they were never observed in our computational experiments.[29]

Of the 1,410 computational experiments that we performed, we encountered 102 cases of *nonconvergence*, in the sense that our convergence criterion had not been met within the allotted 2,000 iterations. These nonconvergent cases occurred only for values of $\rho_2 \geq 4.2$ and became increasingly common as the value of ρ_2 was increased. Nevertheless, we regard almost all of these nonconvergent cases as *near-convergent* situations because in every instance save two the system was very close to converging to a period-2 stationary equilibrium. If these experiments had been continued beyond the upper limit of 2,000 iterations, we have little doubt that, eventually, our convergence criterion would have been met and the

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

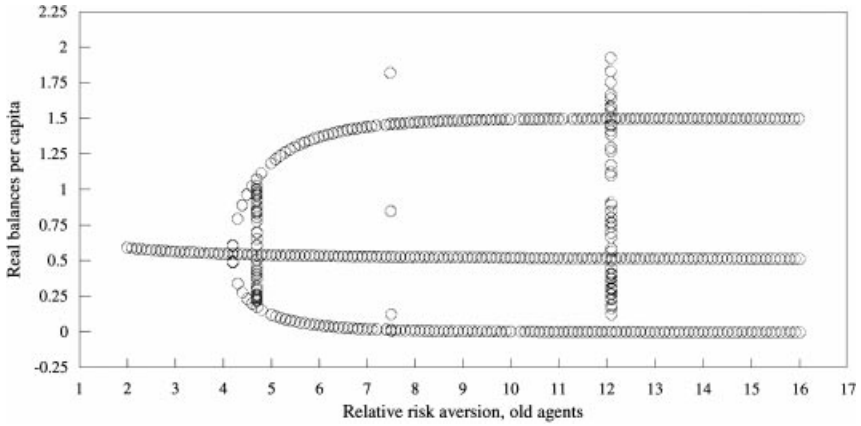


FIGURE 2. Limiting learning dynamics: ten replications at each old-agent relative risk aversion; convergence values or last 50 iterations of each replication plotted. ○ Observed after convergence or 2,000 iterations.

system in fact would have converged to a period-2 equilibrium. Generally, we found that when our systems converged to the monetary steady state, they spent relatively few iterations in the immediate vicinity of this steady state value, by which we mean that the last 10 prices were all equal up to five significant digits. In contrast, when our systems converged to a cycle of period 2, they often spent several hundred iterations in a similar immediate neighborhood of the equilibrium before finally meeting our convergence criterion. Because this process of locking on to the period-2 equilibrium often took a long time, it is not surprising that our 2,000-iteration limit often was encountered in these cases.[30] This situation might have been helped somewhat by a less stringent tolerance.

The specific convergence results from our experiments can be seen by comparing Figure 2, in which we plot the limiting dynamics of our adaptive learning system for each value of ρ_2, with Figure 1, which plots the limiting backward perfect-foresight dynamics for each value of ρ_2. There are slight differences in the methods used to construct these two figures. Figure 1 replicates Grandmont's (1985, p. 1030) Figure 4, and plots the last 50 of 1,000 iterations of the backward perfect-foresight system for each value of ρ_2, taking as the initial condition the peak of the unimodal map. Figure 2 plots the fixed points of any of the stationary equilibria to which the system under our adaptive learning algorithm converged in any of the 10 computational experiments completed at a given value of ρ_2. If the learning system failed to meet the convergence criterion within the allotted 2,000 iterations, then the last 50 of the 2,000 iterations were plotted.

As an example, consider Figure 2 for the value $\rho_2 = 5$. Some of the 10 experiments conducted for this value of ρ_2 resulted in convergence to the monetary steady state, whereas others resulted in convergence to the period-2 cycle that is predicted by the limiting backward perfect-foresight dynamics (see Figure 1).

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

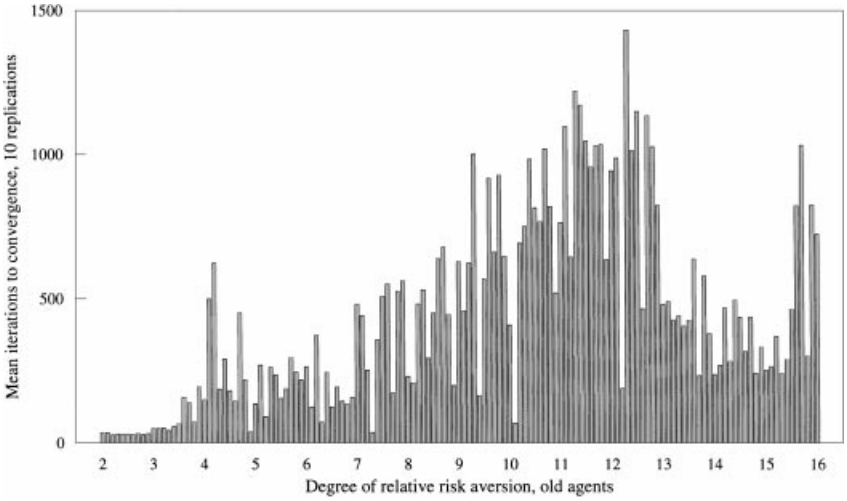


FIGURE 3. Mean number of iterations to convergence (cases of nonconvergence counted as 2,000 iterations).

Therefore, there are three fixed points plotted in Figure 2 for $\rho_2 = 5$. The middle point corresponds to the steady-state value for real-money balances whereas the outer two points correspond to the stationary two-cycle values for real-money balances. As one can see from Figure 2, the outcome of the 10 experiments performed for $\rho_2 = 5$ is the same outcome observed for virtually all other values of $\rho_2 \geq 4.2$; in some experiments the learning system converged to the monetary steady state, whereas in others, the system converged to a stationary, monetary period-2 equilibrium. A comparison between Figures 2 and 1 conveys the relative simplicity of the set of long-run equilibria under learning as opposed to that set under backward perfect foresight. The exceptional cases in Figure 2 occur for a single computational experiment each when $\rho_2 = 4.7$, $\rho_2 = 7.5$, and $\rho_2 = 12.1$. In the first and third of these cases, the system did not approach anything resembling an equilibrium in 2,000 iterations, and the corresponding plots of the last 50 iterations show no clear pattern in Figure 2. In one experiment in which $\rho_2 = 7.5$, the system finished 2,000 iterations very close to a period-4 stationary equilibrium, and the last 50 iterations in this case also are plotted in Figure 2.[31]

The fact that there was wide variation in the number of iterations to convergence can be observed in Figure 3, which depicts the mean number of iterations to convergence for each set of the 10 computational experiments conducted for each different value of ρ_2. In cases in which the system failed to converge within 2,000 periods, the number of iterations to convergence included in the sum used to calculate the mean is 2,000. Thus, the reported mean number of iterations to convergence underestimates the actual mean number of iterations it would have taken the system to achieve convergence, had the system been allowed to continue in those experiments in which noncovergence results were obtained. In the figure, there is a noticeable spike in the mean number of iterations to convergence near the

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

first backward perfect-foresight bifurcation point at around $\rho_2 = 4.2$ (cf. Figure 1) as well as near (less obviously) the second and third bifurcation points. Thereafter, as ρ_2 is increased further, the bifurcations occur more and more frequently and the mean number of iterations to convergence shows a marked tendency to increase up to the point where $\rho_2 = 13$. At this point, the limiting backward perfect-foresight dynamics of Figure 1 indicate the presence of a period-3 cycle, and thus, via Sarkovskii's (1964) theorem, periodic equilibria of every order q exist, where q is an element of the set of positive integers. In this complicated region, there is a marked decrease in the mean number of iterations to convergence. This decrease is due to the increased tendency of the learning system in this region ($\rho_2 \geq 13$) to coordinate on the monetary steady-state outcome rather than on a period-2 or higher-order cycle. When the system converged to the monetary steady state, the number of iterations to convergence was generally less than when the system converged to a period-2 equilibrium.

5. CONCLUSION

We have studied learning in a model in which there can be an extensive multiplicity of stationary equilibria, many of which are periodic. It is often thought that introducing learning into such a model might help with the pervasive multiplicity problem because some of the equilibria might not be attractors under learning and hence might not be viewed as likely long-run outcomes in an actual economy or economic situation driven by human behavior. We have provided some evidence in favor of this notion: Introducing the criterion that equilibria should be learnable does lead to a sharp reduction in the set of stationary equilibria that one would view as plausible in our systems. However, we are unable to rule out convergence to periodic equilibria entirely, and our systems sometimes display qualitatively complicated transitory dynamics for long periods of time.

Experiments with human subjects might provide some further clarification of which equilibria are more likely to emerge in environments like this one. However, as noted in the introduction, genetic learning like that employed here already has been shown to be quite successful in mimicking the behavior of human subjects in controlled laboratory settings, and our results are consistent with the experimental work to date on this topic. Because experiments with human subjects presently are not feasible on the scale contemplated in this paper, we think our computational approach provides a reliable and practical alternative.

NOTES

1. For a survey, see Azariadis (1993).

2. For a summary of this literature, see the surveys by Sargent (1993), Evans and Honkapohja (1995b), and Marimon and McGrattan (1995).

3. Of course, if any type of stationary equilibrium is locally unstable under a particular learning dynamic, it will remain globally unstable under the same learning dynamic, and thus can be ruled out as a candidate for long-run equilibrium.

4. Genetic algorithms were invented by John Holland (1975). For an introduction to genetic algorithms, see Goldberg (1989) or Michalewicz (1996).

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

5. The genetic algorithm can be viewed as a global search algorithm in that it involves parallel processing of a large population of candidate solutions (beliefs) that are initially randomly generated. Dorsey and Mayer (1995) provide support for the notion of the genetic algorithm as a kind of global search algorithm. They compare the performance of the genetic algorithm with several other methods (including simulated annealing) on a test suite of difficult optimization problems (with many local optima) and conclude that the genetic algorithm comes closest to finding the global optimum in all cases. While we refer to our genetic learning algorithm as a global search process, we recognize that our analysis is not truly global, even computationally speaking, because we did not complete multiple experiments based on every possible initialization for a given parameterization. Such an approach is beyond the scope of this paper.

6. See, e.g., the discussion by Boldrin and Woodford (1990).

7. See Grandmont (1994) for a survey of this literature.

8. For a survey of this topic, see Evans and Honkapohja (1995b).

9. Evans et al. (1993) make a very different argument in the context of the same overlapping-generations model analyzed in this paper. They suggest that if a significant fraction of the population uses misspecified models because they view themselves as participating in an economy driven by random processes, and these agents attempt to optimize against forecasts generated by using the unconditional distribution of past prices, then the possibility of stationary periodic equilibria will be eliminated. In the present paper, agents have heterogeneous beliefs, but the existence of periodic equilibria is not threatened.

10. No higher-order perfect-foresight periodic equilibria exist.

11. In subsequent research, Marimon and Sunder (1995) observed two experimental overlapping-generations economies in which inflation realizations were persistently cyclical in the neighborhood of a monetary steady state of the model. These two economies were not subject to any real shocks, and there were no stationary periodic equilibria in the neighborhood of the steady state. The observed cyclic behavior for inflation was attributed to the stability of the adaptive learning system that subjects were thought to have used; the linearized learning system, evaluated at the monetary steady state, had eigenvalues that were complex, with modulus close to, or equal to, one.

12. Although in principle there are many agents born at time t, we do not distinguish among them in this section.

13. Similarly, Benhabib and Day (1982) demonstrate the possibility of periodic and chaotic equilibria in the classical version of this same model in which the time patterns of endowments and preferences are reversed.

14. The periodic equilibria in Grandmont's model can be viewed as deterministic versions of sunspot equilibria first studied by Shell (1977) and Azariadis (1981).

15. For the same reason other researchers who have studied the stability of cycles under learning, [e.g., Guesnerie and Woodford (1991) and Evans and Honkapohja (1995a)] also have assumed that agents form expectations based on prices that prevailed k periods in the past.

16. This is the requirement that the forecast rule be able to *detect the cycle* as discussed by Grandmont (1985) and Grandmont and Laroque (1986). This requirement would rule out price forecasts that were simple averages of past prices, although it would not rule out forecasts based on certain weighted averages of past prices. We chose not to allow such weighted-average specifications so as to keep the analysis of our learning system as straightforward as possible. We suspect that allowing for weighted past average specifications would increase the incidence of convergence to the steady state because the vast majority of these types of rules would not be consistent with any kind of periodic equilibrium.

17. For a discussion of the issue of *how* agents manipulate the data to derive forecasts, see, e.g, Bray and Savin (1986) or Duffy (1994).

18. See Arifovic (1995) and Bullard and Duffy (1998) for similar interpretations regarding the role of the genetic algorithm.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

19. Social learning refers to any process in which individuals learn from the behavior of others. See Vives (1996) for an introduction to this topic.

20. The tournament method of implementing the reproduction operator of the genetic algorithm has some advantages over the biased roulette wheel methodology found in Goldberg (1989) and many other early genetic-algorithm implementations. See Fogel (1994) or Michalewicz (1996) for a discussion. For a theoretical analysis of tournament selection, see Blickle and Thiele (1995).

21. The election operator is properly viewed as a modification of the selection/reproduction operator of the genetic algorithm. In the genetic-algorithm literature, a distinction frequently is made between *pure* selection procedures and *elitist* selection procedures. See, for instance, Grefenstette (1986) or Rudolph (1994). In a pure selection procedure, strings are chosen for reproduction solely on the basis of principles of natural selection. Elitist selection procedures involve two steps. First, pure selection is performed. Second, the elitist procedure stipulates that the best discovered string, either before or (in our case) after reproduction, always remains intact into the next generation. Thus, our reproduction operator, together with our election operator, comprise an elitist selection procedure. Rudolph (1994) has shown that an elitist selection procedure is necessary, though not sufficient for the genetic algorithm to converge asymptotically to one of the perfect-foresight equilibria of the model.

22. One could depart from the pairwise matching protocol that is standard in the genetic-algorithm literature and experiment with the possibility of allowing newborn agents to communicate with more than two other agents. However, one would not want to allow newborn agents to communicate with *all* of the agents alive in the population, because in that case, all newborn agents would adopt the same forecast rule, i.e., the one that yielded the highest level of fitness. Restricting communication to a finite subset of the population ensures that, over some period of time, there is sufficient heterogeneity in the population to allow for a good global search. The restriction that newborn agents communicate with just two other agents can be formalized by assuming that the cost of communicating with more than two agents is infinitely high.

23. For an alternative approach to systems with heterogeneous learning rules, see Evans et al. (1995).

24. Holland's schema theorem is proved by analogy with the two-armed-bandit problem and is found in Holland (1975, pp. 75–88). See Vose (1991) for a more general treatment.

25. With this choice for ℓ, we are explicitly ruling out the possibility of agents coordinating on periodic equilibria of order greater than 256. Experiments with higher values of ℓ, not reported here, did not appear to produce qualitatively different behavior in our systems.

26. Our convergence criterion is described by the following pseudocode: convergence = yes; tolerance $= 1 \times 10^{-10}$; for $s = 1$, $s \leq$ maxcycle $= 256$; for $j = 0$, $j \leq 9$; if $[P(t-j) - P(t-j-s)] \geq$ tolerance, convergence = no $j = 10$ endif; next j; if convergence = yes, $s =$ maxcycle $+ 1$ endif; next s.

27. One could view this as a "strong stability result" in Evan's (1989) sense because it is robust to overparameterization, meaning in our case that agents consider higher-order cycles even though none exist for low values of ρ_2.

28. The exact number of times that our system achieved convergence on the monetary steady state or a period-2 cycle, or failed to satisfy our convergence criterion within 2,000 iterations for each different value of ρ_2, is reported by Bullard and Duffy (1995).

29. We note that the relatively simple equilibria that we observed all had a periodicity that is one of the first few entries in Sarkovskii's ordering. Because three is the last entry in that ordering, one might view it as a complicated equilibrium.

30. The 2,000-iteration limit was chosen to conserve on the amount of computer time it took to complete our computational experiments.

31. As in other cases of near convergence in which the system was close to converging to a period-2 cycle after 2,000 iterations, we are quite confident that this single run that is very close to a period-4 cycle at 2,000 iterations eventually would have met our convergence criterion.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

REFERENCES

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

——— (1995) Genetic algorithms and inflationary economies. *Journal of Monetary Economics* 36, 219–243.

——— (1996) The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy* 104, 510–541.

Arifovic, J. & C. Eaton (1995) Coordination via genetic learning. *Computational Economics* 8, 181–203.

Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer & P. Taylor (1996) Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. Working paper 96-12-093, Santa Fe Institute.

Axelrod, R. (1987) The evolution of strategies in the iterated prisoner's dilemma. In Lawrence Davis (ed.), *Genetic Algorithms and Simulated Annealing*, pp. 32–41. Los Altos: Morgan Kaufmann.

Azariadis, C. (1981) Self–fulfilling prophecies. *Journal of Economic Theory* 25, 380–396.

——— (1993) The Problem of Multiple Equilibria. Manuscript, University of California, Los Angeles.

Bäck, T. (1993) Optimal mutation rates in genetic search. In S. Forrest (ed.), *Proceedings of the Fifth Annual Conference on Genetic Algorithms*. San Mateo: Morgan Kaufmann.

Benhabib, J. & R.H. Day (1982) A characterization of erratic dynamics in the overlapping generations model. *Journal of Economic Dynamics and Control* 4, 37–55.

Blickle, T. & L. Thiele (1995) A mathematical analysis of tournament selection. In L.J. Eshelman (ed.), *Proceedings of the Sixth Annual Conference on Genetic Algorithms*. San Mateo: Morgan Kaufmann.

Boldrin, M. & M. Woodford (1990) Equilibrium models displaying endogenous fluctuations and chaos: A survey. *Journal of Monetary Economics* 25, 189–222.

Bray, M. & N.E. Savin (1986) Rational expectations equilibria, learning, and model specification. *Econometrica* 54, 1129–1160.

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

Bullard, J. & J. Duffy (1995) On Learning and the Stability of Cycles. Working paper 95-006B, Federal Reserve Bank of St. Louis.

——— (1998) A model of learning and emulation with artificial adaptive agents. *Journal of Economic Dynamics and Control* (in press).

Dorsey, R.E. & W.J. Mayer (1995) Genetic algorithms for Estimation problems with multiple optima, nondifferentiability, and other irregular features. *Journal of Business and Economic Statistics* 13, 53–66.

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.W. (1989) The fragility of sunspots and bubbles. *Journal of Monetary Economics* 23, 297–317.

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

——— (1995b) Adaptive learning and expectational stability: An introduction. In A. Kirman & M. Salmon (eds.), *Learning and Rationality in Economics*, pp. 102–126. Oxford: Basil Blackwell.

Evans, G.W., S. Honkapohja & T.J. Sargent (1993) On the preservation of deterministic cycles when some agents perceive them to be random fluctuations. *Journal of Economic Dynamics and Control* 17, 705–721.

Evans, G.W., S. Honkapohja & R. Marimon (1995) Global Convergence in Monetary Models with Heterogeneous Learning Rules. Working paper, University of Oregon, University of Helsinki, and Universitat Pompeu Fabra.

Fogel, D.B. (1994) An introduction to simulated evolutionary optimization. *IEEE Transactions on Neural Networks* 5, 3–14.

Gale, D. (1973) Pure exchange equilibrium of dynamic economic models. *Journal of Economic Theory* 6, 12–36.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*. Menlo Park: Addison–Wesley.
Grandmont, J.-M. (1985) On endogenous competitive business cycles. *Econometrica* 53, 995–1045.
Grandmont, J.-M. (1994) Expectations formation and stability of large socioeconomic systems. CEPREMAP working paper 9424.
Grandmont, J.-M., & G. Laroque (1986) Stability of cycles and expectations. *Journal of Economic Theory* 40, 138–151.
Grandmont, J.-M. & G. Laroque (1991) Economic dynamics with learning: Some instability examples. In W.A. Barnett, B. Cornet, C. D'Aspremont, J.J. Gabszewicz, and A. Mas-Colell (eds.), *Equilibrium Theory and Applications*. Cambridge: Cambridge University Press.
Grefenstette, J.J. (1986) Optimization of control parameters for genetic algorithms. *IEEE Transactions on Systems, Man, and Cybernetics* 16, 122–28.
Guesnerie, R. & M. Woodford (1991) Stability of cycles with adaptive learning rules. In W.A. Barnett, et al. (eds.), *Equilibrium Theory and Applications*. Cambridge: Cambridge University Press.
Ho, T.-H. (1996) Finite automata play repeated prisoner's dilemma with information processing costs. *Journal of Economic Dynamics and Control* 20, 173–207.
Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems.* Ann Arbor: University of Michigan Press.
Holland, J.H. & J.H. Miller (1991) Artificial adaptive agents in economic theory. *American Economic Review Papers and Proceedings* 81, 365–370.
Kareken, J.H. & N. Wallace (1981) On the indeterminacy of equilibrium exchange rates. *Quarterly Journal of Economics* 96, 207–222.
Lim, S.S., E.C. Prescott & S. Sunder (1994) Stationary solution to the overlapping generations model of fiat money: Experimental evidence. *Empirical Economics* 19, 255–277.
Ljung, L. (1977) Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control* AC-22, 551–575.
Marcet, A. & T.J. Sargent (1989) Convergence of least squares learning mechanisms in self–referential stochastic models. *Journal of Economic Theory* 48, 337–368.
Marimon, R. & E. McGrattan (1995) On adaptive learning in strategic games. In A. Kirman & M. Salmon (eds.), *Learning and Rationality in Economics*, pp. 63–101. Oxford: Basil Blackwell.
Marimon, R. & S. Sunder (1993) Indeterminacy of equilibria in a hyperinflationary world: Experimental evidence. *Econometrica* 61, 1073–1107.
——— (1994) Expectations and learning under alternative monetary regimes: An experimental approach. *Economic Theory* 4, 131–162.
——— (1995) Does a constant money growth rule help stabilize inflation?: Experimental evidence. *Carnegie–Rochester Conference Series on Public Policy* 43, 111–156.
Marimon, R., E. McGrattan & T.J. Sargent (1990) 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.E. Spear & S. Sunder (1993) Expectationally driven market volatility: An experimental study. *Journal of Economic Theory* 61, 74–103.
Michalewicz, Z. (1996) *Genetic Algorithms + Data Structures = Evolution Programs*, 3rd ed. New York: Springer-Verlag.
Miller, J.H. (1996) The co-evolution of automata in the repeated prisoner's dilemma. *Journal of Economic Behavior and Organization* 29, 87–112.
Routledge, B. (1995) Adaptive Learning in Financial Markets. Working paper, Carnegie–Mellon University.
Rudolph, G. (1994) Convergence analysis of canonical genetic algorithms. *IEEE Transactions on Neural Networks* 5, 96–101.
Rust, J., J.H. Miller & R. Palmer (1994) Characterizing effective trading strategies: Insights from a computerized double auction model. *Journal of Economic Dynamics and Control* 18, 61–96.
Sargent, T.J. (1993) *Bounded Rationality in Macroeconomics*. Oxford: Oxford University Press.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press

Sarkovskii, A.N. (1964) Coexistence of cycles of a continuous map of a line into itself. *Ukrainian Mathematics Journal* 16, 61–71.
Shell, K. (1977) Monnaie et Allocation Intertemporelle. Working paper, CNRS, Paris (text in English).
Van Huyck, J.B., J.P. Cook & R.C. Battalio (1994) Selection dynamics, asymptotic stability, and adaptive behavior. *Journal of Political Economy* 102, 975–1005.
Vives, X. (1996) Social learning and rational expectations. *European Economic Review* 40, 589–601.
Vose, M.D. (1991) Generalizing the notion of schema in genetic algorithms. *Artificial Intelligence* 50, 385–396.
Woodford, M. (1990) Learning to believe in sunspots. *Econometrica* 58, 277–307.

https://doi.org/10.1017/S1365100598006026 Published online by Cambridge University Press