Macroeconomic Dynamics, **5**, 2001, 303–325. Printed in the United States of America.

GENETIC ALGORITHM LEARNING TO CHOOSE AND USE INFORMATION

BRYAN R. ROUTLEDGE
Carnegie Mellon University

A genetic algorithm (GA) is used to model learning in a financial model similar to the Grossman–Stiglitz model. Individuals need to learn how to use a signal, how to make an inference about a signal from a market-clearing price, and whether or not a signal is worth acquiring. We provide examples in which the GA does and does not converge to the rational expectations equilibrium. Similar to earlier results, the behavior depends heavily on the rate of experimentation or mutation in the GA and the size of the risky-asset supply noise in the economy.

Keywords: Learning, Genetic Algorithm, Rational Expectations

1. INTRODUCTION

Beginning, perhaps, with Keynes (1936) and "animal spirits," the effect of irrational behavior on stock prices has been a frequently debated subject. Irrationality or, more precisely, violations of expected-utility decision theory, has been documented recently in individual investor portfolio decisions.[1] Some research links these individual behavioral irregularities to stock-market return anomalies.[2] Others argue that noise caused by irrationality has important first-order consequences.[3] However, modeling markets and irrationality requires an understanding of how individual behavior manifests itself in equilibrium prices. More practically, it is necessary in order to address questions of security market design and regulation.[4] The difficulty, however, of modeling irrationality is that it is not static. For example, do financial markets act to drive out suboptimal behavior or reinforce it? It is important, therefore, to develop tools for tractable models that can incorporate less than fully rational but adaptive individuals. We attack these issues in two ways. First, we investigate genetic-algorithm (GA) learning in a specific asymmetric information asset-pricing context. Second, more generally, we demonstrate the extent and limit that a GA can be used to represent suboptimal, adapting individuals in a complex market.

Previous versions of this paper circulated with the title "Artificial Selection: Genetic Algorithms and Learning in a Rational Expectations Model." This research has benefited from the comments of Jasmina Arifovic, Jonathan Berk, James Bullard, Murray Carlson, Peter Danielson, John Duffy, Murray Frank, Alan Kraus, Michele Piccione, Norman Schuerhoff, and an anonymous referee. Address correspondence to: Bryan Routledge, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA; e-mail: rout@andrew.cmu.edu.

 1365-1005/01 $9.50

We build on Routledge (1999), which uses the stationary one-period risky asset model of Grossman and Stiglitz (1980) (GS model) and considers adaptive learning. Here, we model the imitation and random experimentation of adaptive learning specifically as a GA. In the classic GS model, traders in a one-period economy choose to acquire a costly signal of the risky asset's terminal dividend (informed traders). Those traders choosing not to acquire the signal (uninformed traders) make an inference about the signal from the market-clearing asset price. The GS model is a canonical model of the informational efficiency of financial markets. It addresses the important issue of how financial markets aggregate and convey information. In the economy that we consider, traders must learn the relevant exogenous relationship between the model's parameters as well as the parameters of the endogenous equation that determines equilibrium price. This setting highlights the difference between learning an exogenous parametric relationship and learning an endogenous equilibrium relationship. In the model, traders need to learn the physical connection between a signal and a dividend [described formally in equation (1)]. Learning about this connection is straightforward because it is independent of the trading decisions of other individuals. Since the signal affects informed-trader demands, the equilibrium price will be influenced by the signal. Uninformed traders trying to infer the signal from the price must learn the equilibrium connection between signal and price. This learning task is complicated because the equilibrium relationship depends on the behavior of other individuals, and therefore depends on the learning of the others.

Routledge (1999) considers the general properties of adaptive learning in a GS-based model. In this paper, we focus more specifically on a GA as an example of a stochastic adaptive learning process. A GA, developed primarily by Holland (1975), is a stochastic optimization technique. Here, we use a GA to simulate the adaptive learning process where successful strategies proliferate and unsuccessful strategies die out. The GA also generates novel strategies by combining and mutating existing strategies in an efficient search of the strategy space. GA's have intuitive appeal as a learning process because they employ both success-based imitation and random experimentation.

We report four representative simulations that demonstrate the stability and convergence properties of GA learning in the GS model. The main result is an illustration that the convergence of adaptive GA learning to the rational expectations equilibrium depends on the level of noise in the GA relative to the noise in the economy. When the GA is quite noisy (a high mutation rate, for example) relative to the economy (high variance in the asset supply), the learning task of the uninformed trader is too difficult and the GA converges to a population consisting only of informed traders and not to the GS equilibrium.

There are several papers that have used GA's to model learning. Axelrod (1987), for example, used GA's to choose an optimal strategy for the repeated prisoner's dilemma game in a fixed population of strategies. Arifovic and Eaton (1998) use a GA in a complex coordination problem; Marimon et al. (1990) consider adaptive agents in a monetary economy; Arifovic (1994) looks at GA's and cobweb cycles; and Arifovic (1996) and Bullard and Duffy (1997) implement GA's in overlapping generations models. Most closely related to the model we present here, Arifovic's (1989) work simulates learning in a rational expectations asset-pricing model [based on Bray (1982)] using a GA to perform least-squares learning.[5]

The model is presented in Section 2, which describes the economy, individuals, and the GA adaptive learning concept. Section 3 describes the results by focusing on four representative simulations. Section 4 concludes.

2. MODEL

To address learning, it is necessary to consider trading behavior through time. However, to maintain a comparison with the original static GS model, we restrict ourselves to the one-period repeated economy in which each realization of the exogenous variables in the economy is independently and identically distributed.[6] The model consists of three parts: the financial market, static individual behavior, and GA adaptive learning. The financial market is a repeated version of the Grossman and Stiglitz (1980) one-period (two-date) endowment economy in which traders can choose whether or not to purchase a costly signal of the terminal dividend (risky-asset payoff). Individual behavior in a given period is characterized by their information choice and the input parameters that determine their asset demand. Finally, the GA is used to simulate adaptive learning, which determines how behavior changes from one period to the next.

2.1. Financial Market

At the beginning of each date, t, a generation of N individuals trades a risk-free and a risky security in a frictionless, competitive market.[7] At the end of date t, the payoffs are realized and end-of-period wealth is fully consumed. The fact that wealth is not transferred between generations preserves the one-period nature of the GS model. Successive dates of the economy are linked by the GA learning and not by wealth accumulation. Because there is no intertemporal consumption decision, the payoff and the price of the risk-free asset are normalized to one. At the beginning of the period, the risky asset has price P_t. At the end of the period, the asset pays a dividend of d_t characterized by

$$d_t = B_0 + B_1 y_t + z_t. \tag{1}$$

B_0 and B_1 are constant parameters of the economy. Unlike in the classic GS model, we assume that individuals do not know the values of B_0 and B_1. The y_t and z_t are independent and identically distributed random variables. The per-capita supply of the risky asset is $\overline{e} + e_t$. As in the original GS model, it is stochastic with e_t distributed independently and identically. Jointly, y_t, z_t, and e_t are mean-zero, uncorrelated, and normally distributed, with respective variances of σ_y^2, σ_z^2, and σ_e^2. The z_t and e_t are not observed until the end of period t. However, agents

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

may choose to pay cost c to observe the signal y_t before trading. This choice endogenously separates the individuals into two groups: informed and uninformed. The proportion of informed traders plays an important role and is denoted λ_t.

2.2. Individual Behavior

At the beginning of each period, all traders are identically endowed with W_0 (recall wealth is not accumulated across periods). Individual n chooses to be informed or uninformed, $\iota \in \{i, u\}$, and the demand for the risky asset, x_t^n. End-of-period wealth, W_{t1}^n, is determined by these choices, the realized asset payoffs, and the budget constraint as

$$W_{t1}^n = W_0 + x_t^n(d_t - P_t) - c1_{\{\iota=i\}}. \quad \textbf{(2)}$$

The cost of information, c, is paid only if the individual is informed ($1_{\{\cdot\}}$ is an indicator function). Finally, individuals have identical constant absolute risk aversion (CARA) preferences for end-of-period consumption, $U(W_{t1}^n) = -\exp(-aW_{t1}^n)$, where a is the risk aversion coefficient. This utility provides the fitness for the GA adaptive learning that follows.

If the trader chooses to be informed, risky-asset demands may depend on the signal y_t as well as the asset price P_t. If the trader chooses to remain uninformed, asset demands may depend only on the price. For reasons that are discussed later, we consider linear demands. For individual n,

$$x_t^n = \begin{cases} \gamma^I\left(\beta_{0t}^{in} + \beta_{1t}^{in} y_t - P_t\right) & \text{if} \quad \iota^n = i \ (n \text{ is informed}) \\ \gamma^U\left(\beta_{0t}^{un} + \beta_{1t}^{un} P_t - P_t\right) & \text{if} \quad \iota^n = u \ (n \text{ is uninformed}), \end{cases} \quad \textbf{(3)}$$

where γ^I and γ^U are specified constants (across traders and time). Given the linearity, we can choose these constants without loss of generality, and so, we choose them to facilitate comparison with the original GS model.

An individual's behavior at date t is completely characterized by the vector $\boldsymbol{\ell}_t^n = \{\iota_t^n, \beta_{0t}^{in}, \beta_{1t}^{in}, \beta_{0t}^{un}, \beta_{1t}^{un}\}$, where ι captures the agent's choice to be informed or not and the $\beta_{jt}^{\iota n}$ determine the portfolio. The behavior of the economy is determined by population $L_t = \{\boldsymbol{\ell}_t^n\}_{n=1}^N$. In particular, L_t determines the market-clearing price (as a function of the realizations of y_t and e_t) for the risky asset. Not surprisingly, the linearity of the asset demands implies a market-clearing price that is linear in the signal y_t and the noisy asset supply as

$$P_t = \alpha_0(L_t) + \alpha_1(L_t)y_t + \alpha_2(L_t)e_t, \quad \textbf{(4)}$$

where the parameters α_j depend on the behavior in the population L_t [see Routledge (1999) for the explicit algebra]. L_t is also sufficient to calculate the expected utility of individual n, given that she is using $\boldsymbol{\ell}_t^n$ to choose information and form her portfolio. Define $f(\boldsymbol{\ell}_t^n, L_t) = E[U(W_{t1}^n)|\ell_t^n, L_t]$. The GA in Section 2.3 uses this measure of fitness of the behavior ℓ_t^n.

The rational expectations equilibrium can be described in terms of a population. This allows us to compare a GA-evolved population to the traditional GS equilibrium one. The equilibrium is characterized by two conditions: First, the GS assumption of individuals' rational expectations implies that their asset demands are optimal. For informed traders, this means that the informed traders know the values of B_0 and B_1 in equation (1). Uninformed traders not only know the exogenous parameters in equation (1), they also know the endogenous α_j in equation (4), allowing them to make a rational inference about the signal. Second, individuals are indifferent about becoming informed. Informed traders have a more precise estimate of the terminal dividend and uninformed traders do not pay the cost of the information. In the GS equilibrium, all agents are identical except for their information choice. Formally, the GS equilibrium, denoted L^*, is defined for all individuals n, $\ell^n = \ell^{i^*} = \{i, \beta_0^{i^*}, \beta_1^{i^*}, \beta_0^{u^*}, \beta_1^{u^*}\}$ or $\ell^n = \ell^{u^*} = \{u, \beta_0^{i^*}, \beta_1^{i^*}, \beta_0^{u^*}, \beta_1^{u^*}\}$, where $\beta_j^{i^*}$ and $\beta_j^{u^*}$ maximize expected utility. In addition, $f(\ell^{i^*}, L^*) = f(\ell^{u^*}, L^*)$. The proportion of informed traders in the GS equilibrium is denoted λ^*. In the simulations that follow, we choose parameters of the economy so that $\lambda^* = 1/2$.

In the classic GS model with maximizing individuals and rational expectations, it is not necessary to specify the functional form for the asset demands because rationality implies that they follow from maximization.[8] In contrast, since individuals are not fully rational, we specify traders as adaptive learners (captured below with the GA). Since this form of learning is non-Bayesian, we need to specify a form for the "rules-of-thumb" that the individuals will use and that the GA will evolve. Practically, this means specifying the functional form for the asset demands. In this paper, we have chosen this class to be linear as in equation (3). Focusing on linear demands is a binding restriction because, in a model in which B_0 and B_1 in equation (1) are not common knowledge, demands of a rational Bayesian would not be linear. It is feasible in the GA simulations to abandon the linearity assumption in equation (3) and use some other ad-hoc specification. For example, we could use a polynomial representation or an artificial neural network. However, it is not clear that the extra numerical complexity would yield additional insight. Besides keeping the numerical calculations simple [the linearity of the price in equation (4)], the specification is convenient for comparing the learning results to the GS equilibrium. In addition, Routledge (1999) shows that fitness-maximizing demands are linear, and so, equation (3) is not *a priori* inconsistent with rationality.

In the GS model, the informed utility-maximizing agent will have asset demands of $\gamma^I(E[d|y] - P)$. By choosing the normalizing constant $\gamma^I = (a\sigma_z^2)^{-1}$, we can interpret the adaptive-learning individual as trying to set β_{jt}^{in} in equation (3) equal to the corresponding B_j from equation (1) (for $j = 0, 1$). Note the linearity in the functional form of the asset demands makes the choice of a normalizing constant, γ^I, without loss of generality. Similarly, for the uninformed, the rational expectations equilibrium demands are $\gamma^U(E[d|P] - P)$. The linearity of (4) and normality imply that in the GS rational expectations equilibrium, $E[d|P]$ is linear in the price. Again, by choosing the γ^U appropriately,[9] we can interpret the

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

individual behavior represented in β_{t0}^{nu} and β_{t1}^{nu} as the parameters used to extract the relevant information about the expected dividend from the price.

2.3. Genetic Algorithm

So far, we have described one period of the economy, given a population of individuals L_t. A GA, representing adaptive learning, will determine how population changes to L_{t+1}. The GA is based on three operators: selection, crossover, and mutation. The operators are applied to a binary string representation of ℓ_t^n (a chromosome). One bit (gene) determines the agent's information choice, ι_t^n. The remaining bits are the binary representation of the real-valued $\beta_{jt}^{\iota n}$ parameters in ℓ_t^n.[10] From the point of view of individual n, the GA determines ℓ_{t+1}^n as follows: First, with probability proportional to fitness, two strings from L_t are selected. These strings are combined using crossover, which takes the first part of the first string and the second part of the second string. Finally, mutation perturbs some of the elements of the combined string state, yielding ℓ_{t+1}^n.[11] The selection operator captures the imitation or adaptation component of a learning process while the crossover and mutation operators drive the experimentation. The crossover step helps to guide experimentation to the most promising strategies, and mutation ensures that the learning process always has some degree of experimentation. Note that, when the population consists of similar individuals, the selection and crossover steps have little effect. It is the mutation rate that determines the limiting size of experimentation.

For the GA to create successive populations based on the success of trading strategies, we need to determine how fitness is defined and calculated. We define fitness of strategy ℓ_t^n in population L_t as the expected utility; that is, $f(\ell_t^n, L_t) = E[U(W_{t1}^n)|\ell_t^n, L_t]$. To calculate the expected utility fitness measure, the model is repeated with a fixed population. For each generation, the one-period model is simulated R times, using the current population, L_t. The repetitions of the economy are independent of one another and, for $t_r = t_1, \ldots, t_R$, consist of a random draw of the economic variables y_{t_r}, z_{t_r} and e_{t_r}, calculating asset demands using the agents' parameters ℓ_t^n according to (3), solving for the equilibrium price and determining the individual's terminal wealth $W_{t_r1}^n$ using equation (2). By repeating the single-period economy many times, we can approximate the expected utility of the behavior ℓ_t^n, given the population L_t. Given the normality and CARA preferences, we can work with the mean-variance specification of

$$\bar{f}^n\left(\ell_t^n, L_t\right) = m\left(W_{t_r1}^n\right) - \frac{a}{2}v\left(W_{t_r1}^n\right). \tag{5}$$

The $m(\cdot)$ and $v(\cdot)$ are the sample mean and variance operators over the R repetitions.

The choice of a fitness definition is not innocuous. Alternative definitions of fitness are available. For example, Blume and Easley (1992) assume fitness is based on average return or wealth. By defining success purely on mean wealth, the utility function is typically unimportant for the long-run behavior. In Blume and Easley (1992), only investors with log utility functions survive. Since we are investigating the specific GS model, it is important that the adaptive algorithm preserves the utility function.

The calculation of fitness is also important. The GA used here selects strings for imitation according to fitness. Specifically, the probability that strategy ℓ^n is selected is proportional to $\overline{f}^n/(\Sigma_m \overline{f}^m)$. Thus, even though wealth does not affect asset demands, the parameter for initial wealth, W_0 [which appears in equation (2)], is not innocuous. An increase in initial wealth, while not affecting any of the equilibrium relationships, makes the probability of any one agent's learning state being selected (imitated) closer to uniform. This makes the learning algorithm "noisier" by weakening the connection between success and representation in subsequent generations. The choice of how many times the economy with a fixed population is repeated (R) also influences the level of noise in the GA selection operator. In the simulations presented here, we choose $R = 1{,}000$ so that $\overline{f}^n$ gives a relatively precise measure of f^n. Alternatively, one could measure fitness with fewer repetitions including using just one realization of the economy ($R = 1$). This provides a noisier measure of fitness and adds more noise to the selection portion of the GA.[12] The specific results presented here are sensitive to these choices in the measurement of fitness. However, the general conclusion of the simulations presented below is robust. Convergence to the GS equilibrium depends on the level of noise in the GA relative to that in the economy. The measurement of fitness is one of the items that affects the level of GA noise.

Simulations begin with the creation of an initial string for each agent in the population. Four representative simulations are presented. In Simulations 1 and 2, the initial population contains the GS rational expectations equilibrium values (i.e., $L_0 = L^*$). In Simulations 3 and 4, initial learning states for each agent are selected randomly. However, the same randomly selected population, $\tilde{L}_0$, is used in both Simulations 3 and 4. The simulation algorithm is summarized in Figure 1. The parameters of the GA learning process are in Table 1. Table 2 contains the parameters of the economy that are common to all the simulations. Table 3 presents the parameters specific to each of the four simulations and summarizes the results.

3. RESULTS

Four simulations are presented in detail. They are representative of a large number of such simulations that have been conducted.[13] The characteristics that distinguish the simulations are the initial populations and whether or not they yield the GS equilibrium ($L_t \to L^*$). The description of the simulations focuses on the proportion of informed agents (λ_t) and the average demand parameters. Define

$$\beta_{0t}^{I} = \frac{1}{\lambda_t N} \sum_{\iota^n = i} \beta_{0t}^{in}$$

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

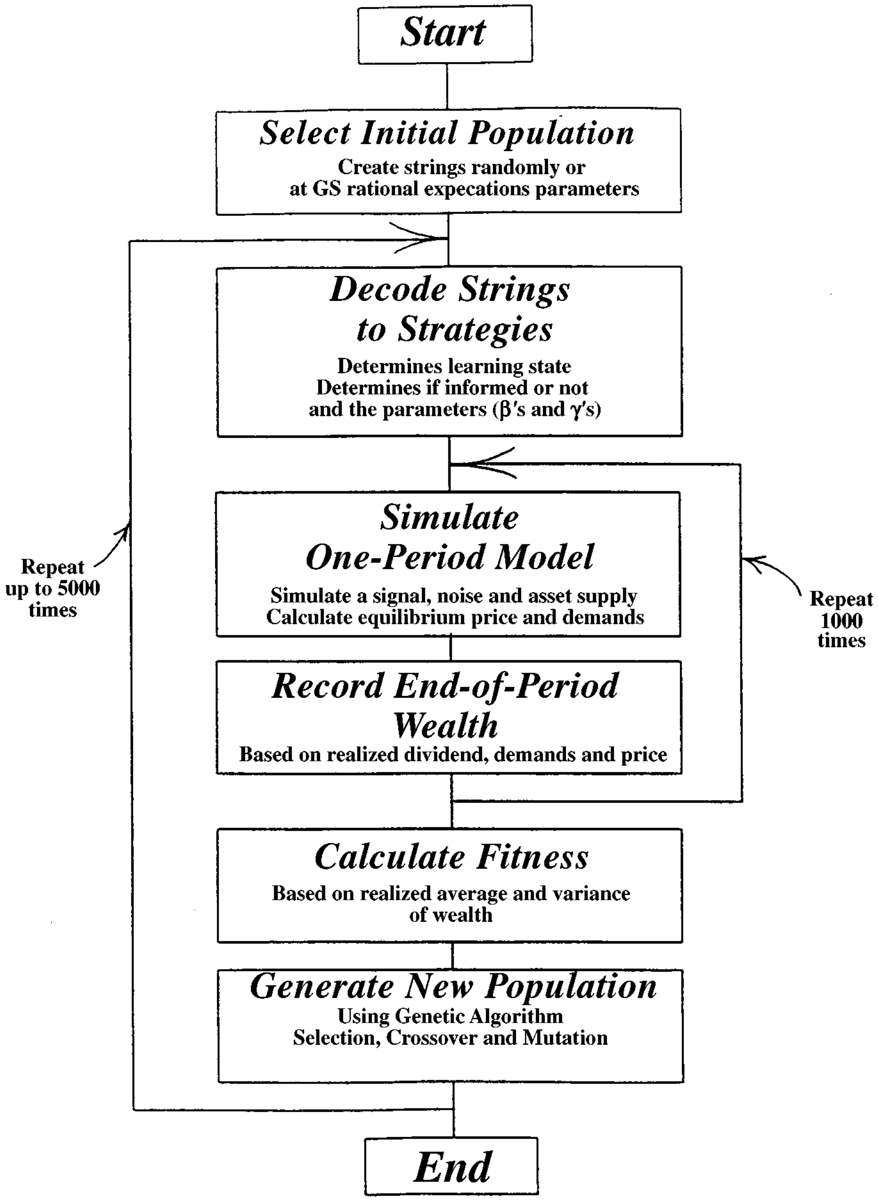

Figure 1. Schematic for GA simulation of GS.

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

Table 1. Parameters of the Genetic Algorithm[a]

Parameter	Symbol	Value
Population size	N	1,000
Repetitions of the single period per generation	R	1,000
Length of simulation		5,000
Length of string used to represent individual	ℓ_t^n	81
Crossover probability (if crossover does not occur one of the strings is selected)		0.7
Mutation probability (probability that a "0" in the string is changed to a "1" or vice versa.)		0.0001
Expected number of mutations per generation		8.3
Range length for asset demand parameters[b]		0.2

[a] Parameters are common to all simulations. They define the behavior of the GA adaptive learning.
[b] The minimum and maximum values for β_{jt}^{in} are chosen so that the GS rational expectations parameter values lie in the middle of the range.

Table 2. Parameters of the economy[a]

Parameter	Symbol	Value
Coefficient of absolute risk aversion	a	2.0
Initial wealth per repetition of economy	W_0^n	0.1
Average number of shares in economy	$\bar{e}$	1,000
Average dividend [see equation (1)]	B_0	0.1
Sensitivity of dividend to signal	B_1	1.0
Standard deviation of signal	σ_y^2	0.0004
Standard deviation of dividend	$\sigma_y^2 + \sigma_z^2$	0.0008
Informed-trader asset demand constant [equation (3)]	$\gamma^I = (a\sigma_z^2)^{-1}$	1,250
GS rational expectations equilibrium proportion of informed traders	λ^*	0.5
GS rational expectations equilibrium demand parameters of informed traders	B_0^{i*}	0.1
	B_1^{i*}	1.0

[a] Parameters for the financial market that are common to all simulations.

and

$$\beta_{0t}^U = \frac{1}{(1-\lambda_t)N} \sum_{\iota^n = u} \beta_{0t}^{un}.$$

Note that β_{0t}^I is an average across only those traders who are informed and, similarly, for those who are uninformed. To reduce the number of figures presented, we focus discussion on the average intercept terms. Figures for similarly defined β_{1t}^I and β_{1t}^U are qualitatively similar, and so, they are not presented. Finally, the optimal asset demand parameters given the current population (utility-maximizing, given knowledge of L_t) and the rational expectations parameters are shown for comparison.

TABLE 3. Simulation-specific parameters[a] and results

Parameter	Simulation number 1	2	3	4
Initial population	GS REE (L^*)	GS REE (L^*)	Random ($\tilde{L}_0$)	Random ($\tilde{L}_0$)
Number of generations that population is constrained to $\lambda_t = \lambda^*$	1	1	**100**	**50**
Asset supply noise (σ_e^2)	10.0	**1.0**	10.0	10.0
Cost of information (for $\lambda^* = 0.5$)	0.0146027	**0.001585**	0.0146027	0.0146027
Uninformed-trader asset demand constant	1,179.1	1,242.1	1,179.1	1,179.1
GS rational expectations demand parameters for uninformed traders				
$\beta_0^{u^*}$	0.000986	0.001119	0.000986	0.000986
$\beta_1^{u^*}$	0.968116	0.996800	0.968116	0.968116
Informativeness of price at GS equilibrium $V[y\|P, L^*]/\sigma_y^2$ (%)	6.015	0.636	6.015	6.015
Result of simulation presented	L^* stable (Figure 2)	L^* not stable (Figure 3)	Converge to L^* (Figure 4)	Converge $\lambda = 1.0$ (not L^*) (Figure 5)
Frequency (%) of convergence of GS equilibrium (based on repeating each simulation 100 times with a different random seed and, in Simulations 3 and 4, a different randomly selected initial population. Convergence to the GS equilibrium is defined at the terminal generation $t = 5{,}000$ as $\lambda^* - 0.05 \leq \lambda_t \leq \lambda^* + 0.05$)	100	0	94	0

[a] Important parameters are shown in bold.

3.1. Stability Examples

Simulations 1 and 2 both begin from the GS rational expectations equilibrium, $L_0 = L^*$. They differ only in the noisiness of the risky-asset supply, σ_e^2, and corresponding cost of information, c. The cost of information is chosen so that, at the GS equilibrium, half the population is informed ($\lambda^* = 1/2$). The reason that σ_e^2 plays such an important role is discussed later. In Simulation 1, where supply noise is higher, the proportion of informed agents remains close to its GS equilibrium level of 0.5 (see Figure 2A) and the GS equilibrium is stable. In contrast, in Simulation 2, where the supply noise is smaller, λ_t drifts steadily toward 1 (see Figure 3A). In this simulation, by generation 2,000, with the exception of random-agent experimentation, all agents are informed.

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

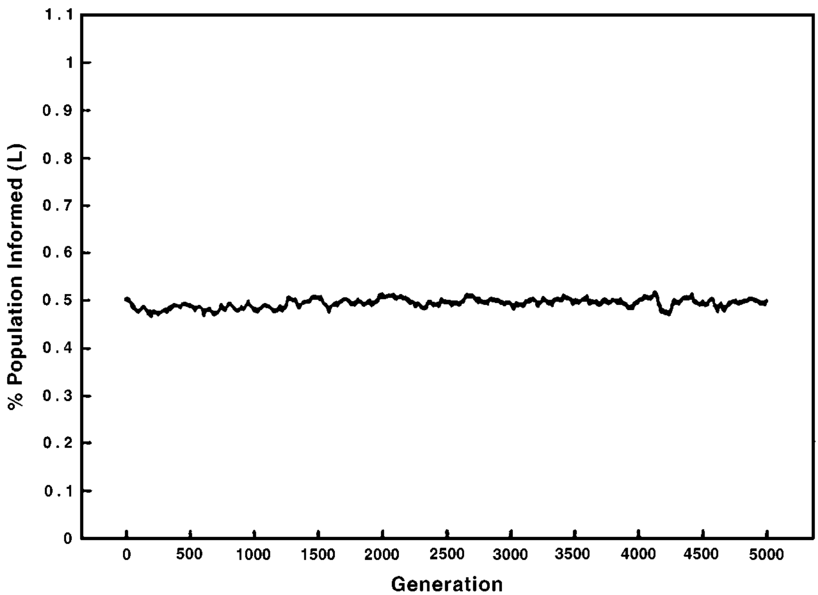

FIGURE 2A. Stable GA simulation—proportion of informed traders. This simulation starts from the GS rational expectations equilibrium (L^*) and remains in a close neighborhood. Parameters are listed in Tables 1–3.

The average demand parameters of the informed agents, β_{0t}^{I}, differ only slightly in the two simulations (compare Figures 2B and 3B). This reflects the fact that the learning about the exogenous signal–dividend relation is relatively unaffected by behavior of other traders or the noise in the risky-asset supply. The small variations from the optimal demand parameters [which are the B_j from (1)] are similar to an estimation error that reflects that the parameters are chosen on the basis of a sample path of dividend–signal pairs.

The behavior of the uninformed agents' asset demand parameters is quite different in the two simulations. In Simulation 1, where λ_t remains close to its GS equilibrium level (see Figure 2A), the average β_{0t}^{U} also remains close to its rational expectations level (see Figure 2C).[14] The distance between the evolving parameters and optimal ones is small in terms of fitness (or utility). In Simulation 2, most agents eventually become informed (see Figure 3A). The evolving average parameter of the uninformed, β_{0t}^{U}, diverges from its optimal level (see Figure 3C). After the first 250 generations, L_t differs substantially from the initial state ($L_0 = L^*$). The difference is large enough that asset demands are sufficiently far from optimal, the fitness of the uninformed is below that of the informed, and adaptation leads to more informed agents. The fitness difference persists even as λ_t approaches 1. By generation 1,600, almost all agents are informed. Figure 3C shows the average demand parameter, which is an average across only the uninformed traders. With

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

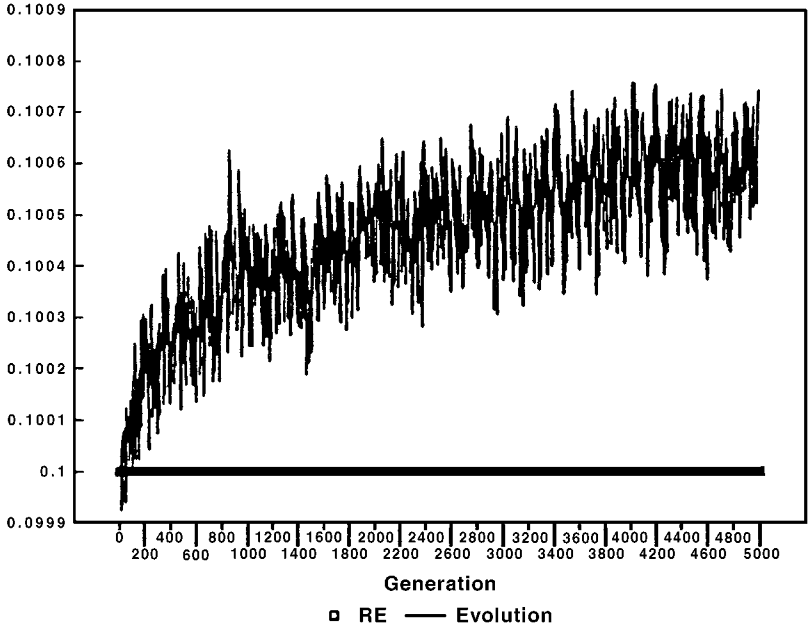

FIGURE 2B. Stable GA simulation—informed average demand parameter. $\beta^I_{0t} = 1/(\lambda_t N)\Sigma_{\iota^n=i}\beta^{in}_{0t}$, where β^{in}_{0t} is described in equation (3).

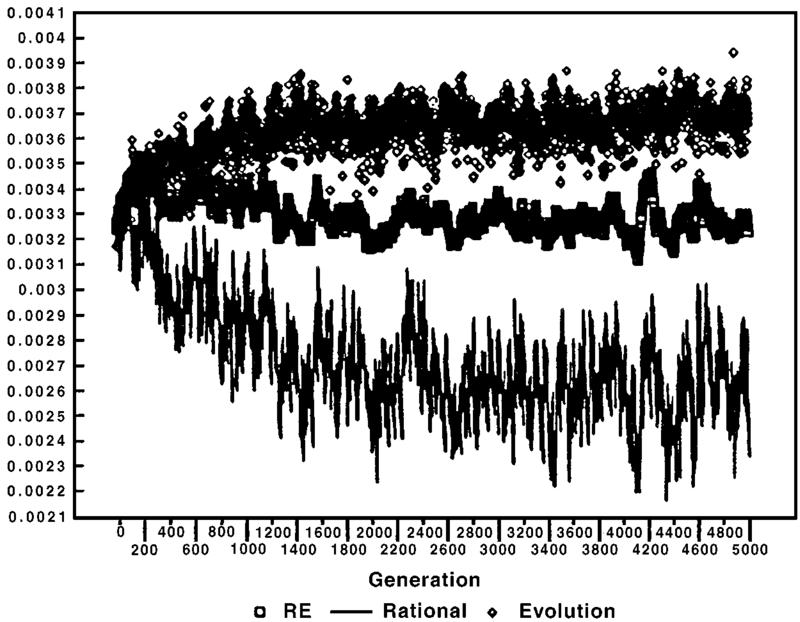

FIGURE 2C. Stable GA simulation—uninformed average demand parameter, where $\beta^U_{0t} = 1/[(1-\lambda_t)N]\Sigma_{\iota^n=u}\beta^{un}_{0t}$, where β^{un}_{0t} is described in equation (3).

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

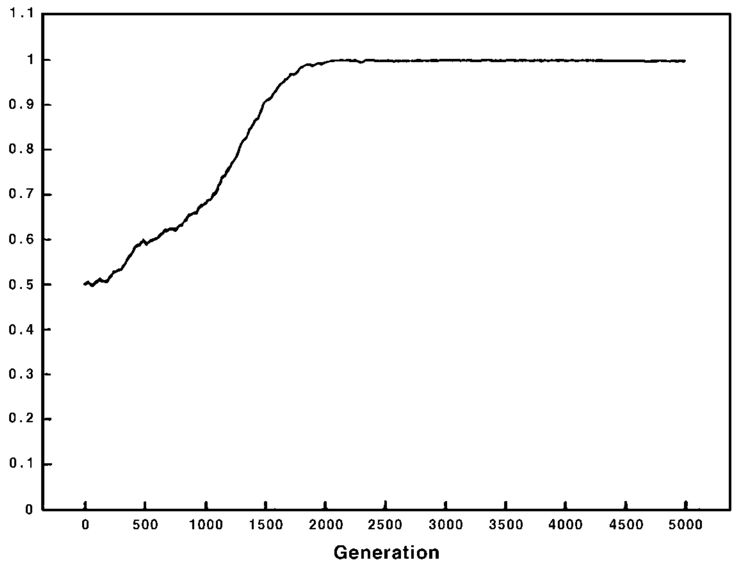

FIGURE 3A. Unstable GA simulation—proportion of informed traders. This simulation starts from the GS rational expectations equilibrium (L^*) and diverges toward $\lambda = 1$. Parameters are listed in Tables 1–3.

fewer uninformed, the average is more volatile. More importantly, with so few uninformed agents, uninformed agents' learning by imitation does not work. As a potential uninformed trader, there is no one to copy. For most traders (since $\lambda_t \approx 1$), the part of the genetic string that maps into the parameter β_{0t}^{un} has no bearing on their fitness. It is recessive. Therefore, the only learning about how to be an uninformed trader can come from experimentation (driven by mutation in the GA). Since this learning is not that efficient, the utility of the uninformed traders always lies below that of the informed. The result is that all traders become informed.

The economies represented in Simulations 1 and 2 are very similar. However, they lead to different behavior. To understand the reason for the differences it is helpful to look at the components of fitness. The proportion of informed traders, λ_t, tends to increase when the informed traders have higher fitness. So, consider the fitness ratio f^i/f^u for an informed trader, i, and uninformed traders, u, in the population, L_t.[15]

$$\frac{f\left(\ell_t^i, L_t\right)}{f\left(\ell_t^u, L_t\right)} = \frac{E\left[U\left(W_{t1}^i\right)\middle|\ell_t^i, L_t\right]}{E\left[U\left(W_{t1}^u\right)\middle|\ell_t^i, L_t\right]}$$

$$= \left[\frac{f^{i^*}}{f^{u^*}}\right]\left[\frac{\xi^i}{\xi^u(L_t)}\right]$$

$$= \left\{\frac{\exp(-ac)}{\left[\sigma_z^2/\left(\sigma_z^2 + V[y|P, L_t]\right)\right]^{0.5}}\right\}\left[\frac{\xi^i}{\xi^u(L_t)}\right] \qquad \textbf{(6)}$$

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

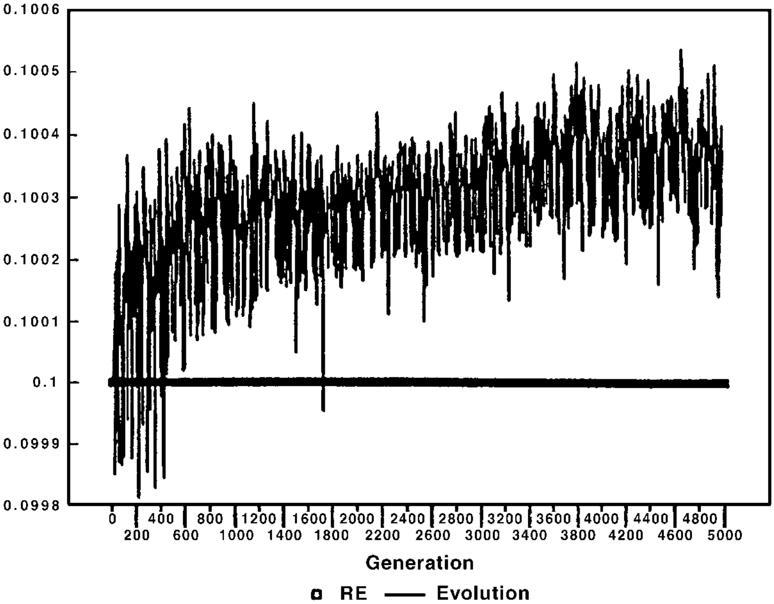

Figure 3B. Unstable GA simulation—informed average demand parameter, where $\beta_{0t}^{I} = 1/(\lambda_t N)\Sigma_{\iota^n=i}\beta_{0t}^{in}$ where β_{0t}^{in} is described in equation (3).

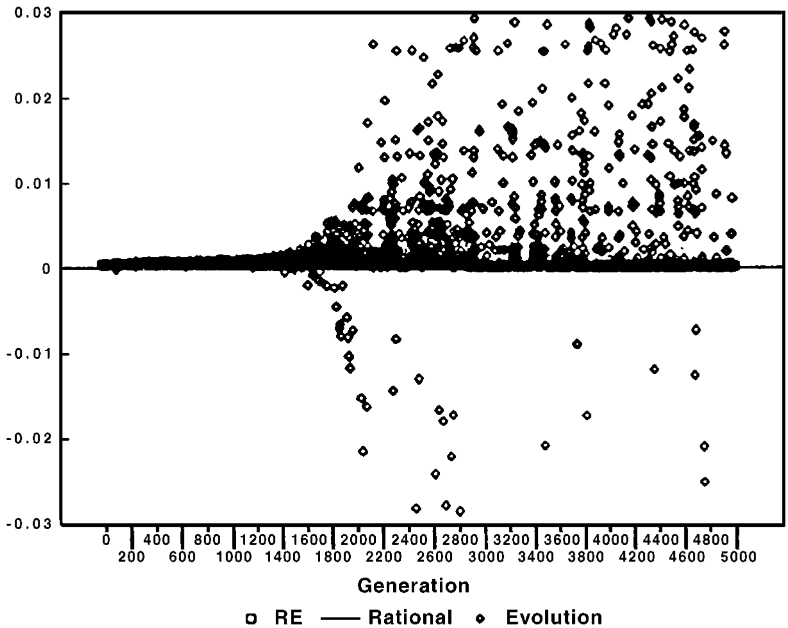

Figure 3C. Unstable GA simulation—uninformed average demand parameter, where $\beta_{0t}^{U} = 1/[(1-\lambda_t)N]\Sigma_{\iota^n=u}\beta_{0t}^{un}$, where β_{0t}^{un} is described in equation (3).

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

The explicit derivation is omitted since it parallels a similar one in Routledge (1999). The result is intuitive. The first line in equation (6) is the definition of fitness. The second line states the fitness ratio in terms of maximal fitness (i.e., expected utility at the GS rational expectations parameters) and error terms ξ^i and ξ^u. These error terms have the property that $\xi^i, \xi^u \in (0, 1]$ and reach 1 only when the traders have maximal fitness.[16] The final line highlights the two key elements that contribute to the relative success of the informed and uninformed. In the first square bracket is the cost of information borne only by the informed and the additional variance faced by the uninformed from inferring the signal from the equilibrium price. This trade-off between information cost and informativeness of the asset price is present in the rational expectations GS model. The second set of square brackets contains the relative errors in the adaptive strategies of the traders.

Change in the population has a different effect on informed and uninformed traders. Informed-trader fitness is unaffected by the behavior of the rest of the population since an informed trader is learning about the exogenous dividend–signal relation in equation (1). This is why ξ^i in equation (6) does not depend on L_t. In contrast, uninformed fitness is affected by the population in two ways. First, the quality of the inference from the equilibrium price about the signal increases when there are more informed traders ($V[y|P, L_t]$ is decreasing in λ_t). Second, since an uninformed trader is trying to make an inference based on the equilibrium condition in (4), the error that an uninformed trader makes, $\xi^u(L_t)$, depends on the behavior of all other traders. For example, consider an uninformed trader, u, who currently has demand parameters (in ℓ_t^u) that are appropriate given the current population L_t and so has high fitness [$\xi^u(L_t) \approx 1$]. Informally, the uninformed trader "knows" the endogenous price–signal relation in equation (4) and correctly forms asset demands. If there is a change in the population (through adaptation or experimentation by another individual), there are two effects on individual u's fitness. First, since the endogenous market-clearing price relationship in equation (4) has changed, the asset demand parameters in ℓ_t^u are less well adapted to the new population L_{t+1}. This reduces u's fitness [$\xi^u(L_{t+1}) < \xi^u(L_t)$]. Second, if the new population has a higher proportion of informed traders ($\lambda_{t+1} > \lambda_t$), then the price is more informative about the signal ($V[y|P, L_{t+1}] < V[y|P, L_t]$), which increases u's fitness. The magnitude of this latter effect is determined by the economy parameter σ_e^2. When there is a large variance in the risky-asset supply, the informed traders' demands are hard to infer from the market-clearing price. An increase in the number of informed traders substantially increases the aggregate informed trade relative to the asset supply noise. In contrast, when σ_e^2 is low, an extra informed trader makes little difference to the informativeness of the price since the price is already very informative.

Returning to Simulations 1 and 2, note that they differ only in the value of the risky-asset supply variance, σ_e^2, and associated cost of buying information, c (see Table 3). Simulation 2 begins at the GS equilibrium population $L_0 = L^*$. At this point, all traders have maximum fitness [$\xi^i = \xi^u(L^*) = 1$]. As the GA progresses, random experimentation (mutation) moves the population away from the

GS equilibrium. Informed fitness falls slightly because of individual experimentation, and uninformed fitness falls dramatically because of the experimentation by all of the other traders. Uninformed fitness is sensitive to the whole population [$\xi^u(L_t)$ depends on L_t]. The degree to which the fitness falls depends on the magnitude of the random experimentation in the population.[17] The low relative fitness of the uninformed leads the selection operator of the GA to choose relatively more informed traders and λ_t increases. Since, in Simulation 2, σ_e^2 is small, the increase in the proportion of informed traders has a small effect on uninformed fitness. The loss in fitness due to the changing population swamps the small benefit of more informed traders [the change in $\xi^u(L_t)$ has a greater impact than the change in $V[y|P, L_t]$]. The net result is that all traders eventually become informed. In Simulation 1, random experimentation (mutation) moves the population away from L^*, reducing the fitness of the uninformed. This leads to an increase in the proportion of informed traders. In this simulation with a larger degree of asset supply noise, σ_e^2, the higher proportion of informed traders increases the fitness of the uninformed enough to overcome the loss due to the changing population [the change in $V[y|P, L_t]$ is larger]. As a result, imitation–selection in the GA drives the fraction of informed traders down and λ_t remains near the GS equilibrium value of $1/2$ (see Figure 2A). This stability in the fraction of informed traders lets the uninformed traders "re-learn" the optimal demand parameters, and the population as a whole remains close to the GS equilibrium. Hence, the stability of the GS equilibrium in this GA learning setting is determined by the size of the variance in the asset supply, σ_e^2, relative to the level of experimentation (mutation and other noise) in the GA.

3.2. Convergence Examples

The preceding two simulations began at the GS rational expectations equilibrium. Simulations 3 and 4 begin from (the same) randomly selected population learning state ($\tilde{L}_0 \neq L^*$). In these simulations, significant learning is required because agents are not endowed with optimal initial conditions [ξ^i and $\xi^u(L_0)$ are significantly less than 1]. The parameters for this simulation are identical to those for Simulation 1. Under these parameters, the GS equilibrium is stable. Simulation 3 is an example of a simulation that converges to the GS equilibrium ($L_t \to L^*$). Simulation 4 does not. The difference in the two simulations is a constraint on the movement of λ_t in the initial generations of the simulations. In Simulation 3, the proportion of informed agents is constrained to the rational expectations equilibrium level for the first 100 generations: that is, $\lambda_t = 1/2$ for $t \leq 100$. In Simulation 4, $\lambda_t = 1/2$ for $t \leq 50$. This constraint is implemented by simply forcing the first bit in the genetic string representing ℓ_t^n during the initial generations. The constraint is seemingly rather minor (the simulations last 5,000 generations) yet influences the long-run properties of the simulations.

Simulation 3 converges to the GS equilibrium. The proportion of informed agents initially increases to $\lambda_t = 0.7$ by generation 500 (see Figure 4A). After this point, λ_t slowly declines back toward the GS equilibrium level and finally reaches

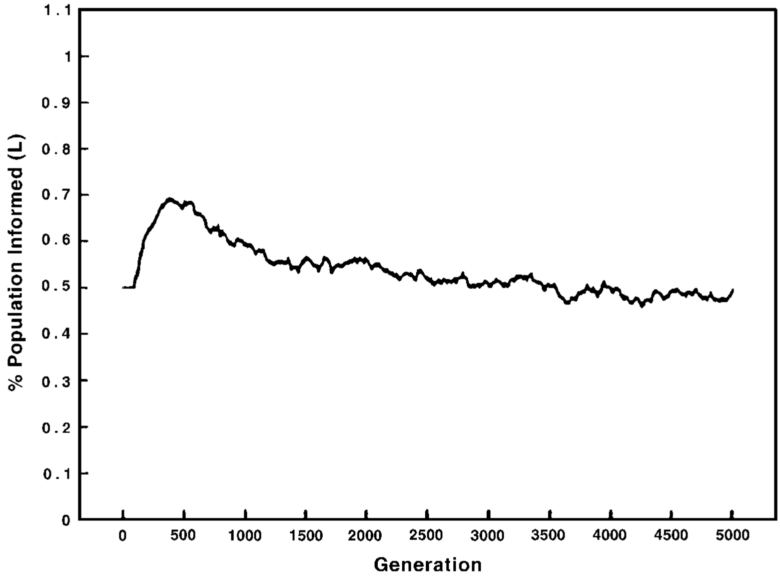

FIGURE 4A. GA simulation converges to rational expectations equilibrium—proportion of informed traders. This simulation starts from a randomly selected population ($\tilde{L}_0$) and converges (close) to the GS equilibrium (L^*). Parameters are listed in Tables 1–3.

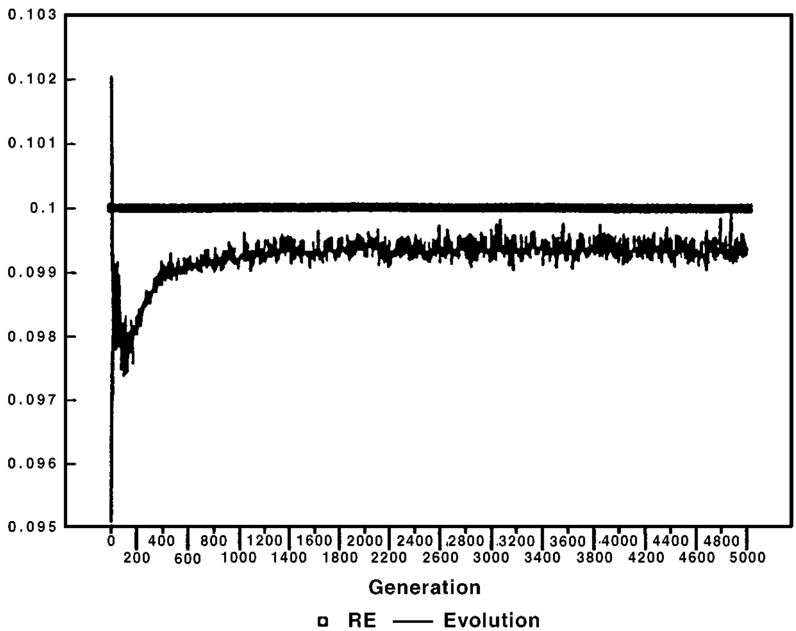

FIGURE 4B. GA simulation converges to rational expectations equilibrium—informed average demand parameter. $\beta_{0t}^{I} = 1/(\lambda_t N)\Sigma_{\iota^n=i}\beta_{0t}^{in}$, where β_{0t}^{in} is described in equation (3).

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

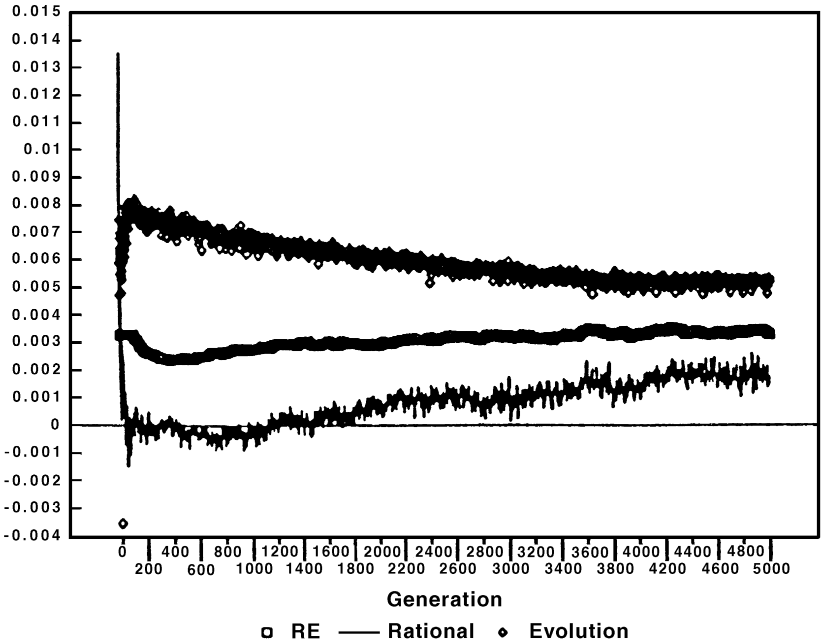

FIGURE 4C. GA simulation converges to rational expectations equilibrium—uninformed average demand parameter. $\beta^U_{0t} = 1/[(1-\lambda_t)N]\Sigma_{\iota^n=u}\beta^{un}_{0t}$, where β^{un}_{0t} is described in equation (3).

$\lambda_t = 0.5$ around generation 2,500 and remains in that vicinity for the remainder of the simulation. In contrast, Simulation 4 does not converge to the GS equilibrium (see Figure 4A). After generation 50, λ_t increases quickly to $\lambda_t = 0.8$ by generation 200, and then continues to drift toward 1, reaching $\lambda_t = 1.0$ by generation 1,000, and remains near there for the duration (see Figure 5A).

As with the previous two simulations, the behavior of the informed demand parameters in Simulations 3 and 4 is similar. Note that, in both Figures 4B and 5B, there is a great deal of variation in β^I_{0t} in the initial few generations, reflecting the rapid initial learning. However, in both simulations the informed traders quickly hone in on the rational expectations (optimal) parameters. As an aside, the slightly higher variation apparent in Figure 4B relative to Figure 5B reflects that the β^I_{0t} is an average over more traders in Simulation 4 since λ_t is close to 1.

Figures 4C and 5C show the average demand parameters for the uninformed traders. As with the informed, there is a great deal of variation in the initial generations as the agents struggle to improve from the randomly assigned initial demand parameters. The large degree of variation in the initial populations, L_t, means that the uninformed traders have lower fitness than the informed. Since they are learning an endogenous relation, they are chasing a moving target and their learning is slower [ξ^i goes to 1 faster than $\xi^u(L_t)$]. For this reason, imitation–selection leads to more informed traders in both simulations (see Figures 4A and 5A). Similar to the

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press

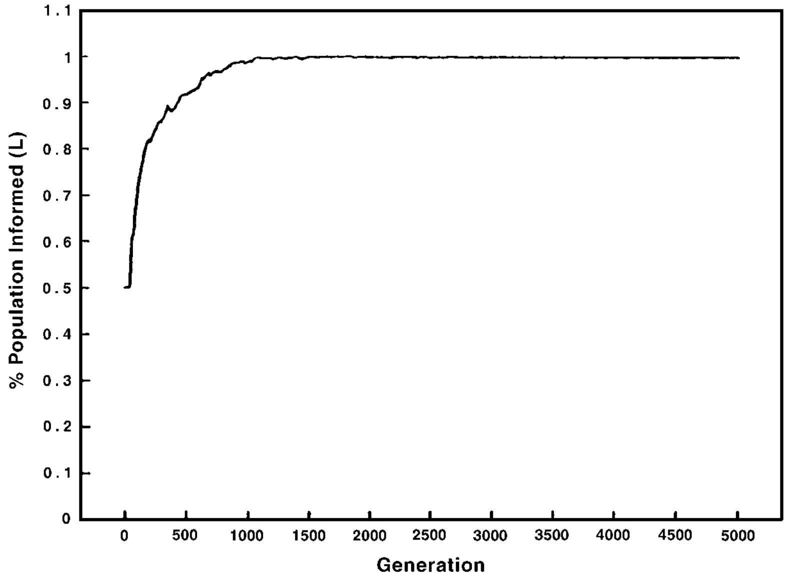

FIGURE 5A. GA simulation does not converge to rational expectations equilibrium—proportion of informed traders. This simulation starts from a randomly selected population (same $\tilde{L}_0$ as in Simulation 3) and does not converge to a neighborhood close to the GS equilibrium (L^*). Parameters are listed in Tables 1–3.

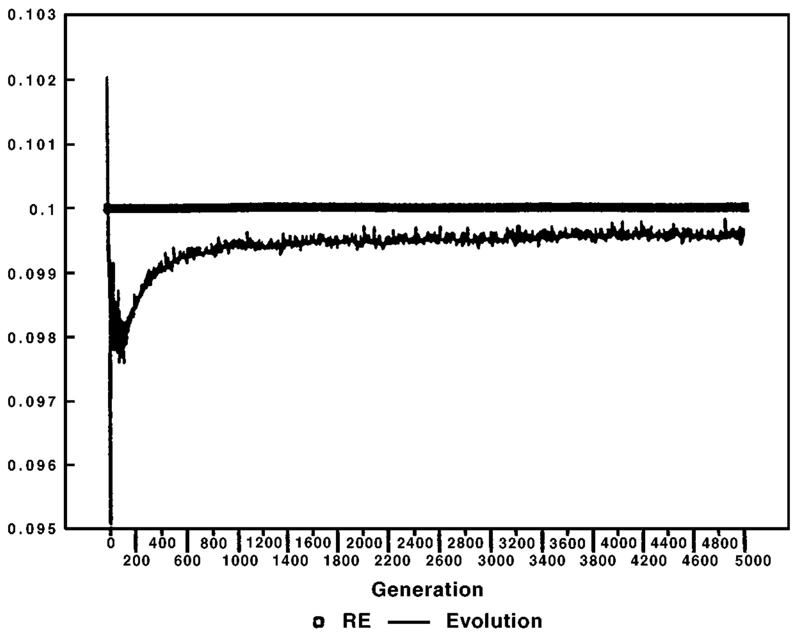

FIGURE 5B. GA simulation does not converge to rational expectations equilibrium—informed average demand parameter. $\beta^{I}_{0t} = 1/(\lambda_t N)\Sigma_{\iota^n=i}\beta^{in}_{0t}$, where β^{in}_{0t} is described in equation (3).

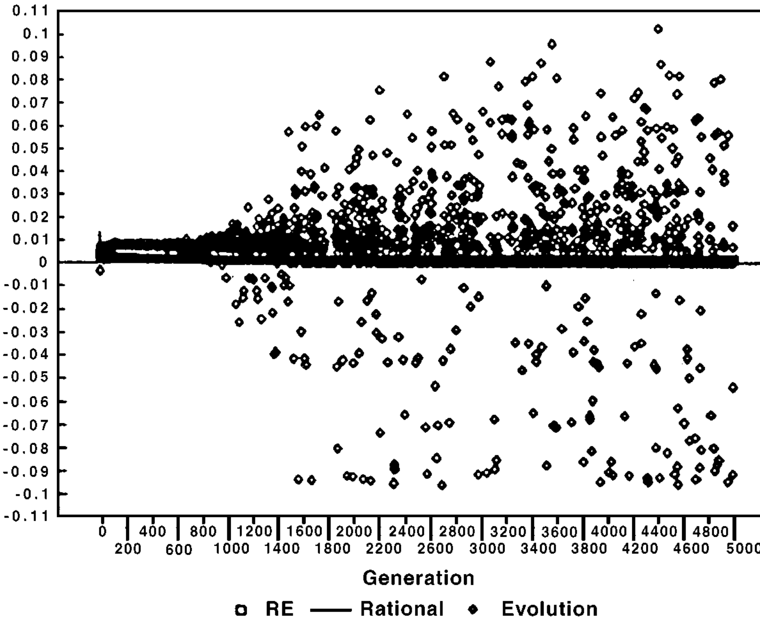

FIGURE 5C. GA simulation does not converge to rational expectations equilibrium—uninformed average demand parameter. $\beta_{0t}^{U} = 1/[(1-\lambda_t)N]\Sigma_{\iota^n=u}\beta_{0t}^{un}$, where β_{0t}^{un} is described in equation (3).

situation in Simulations 1 and 2, this improves the informativeness of the price and improves the fitness of the uninformed traders [$V[y|P, L_t]$ decreases]. However, the benefit of the larger number of informed traders does not immediately offset the fitness disadvantage due to slower learning. This is why λ_t continues to rise. In Simulation 3, the uninformed traders eventually learn enough about the endogenous relationship to benefit from the increased informativeness of the price. Their fitness increases above that of the informed traders and λ_t declines back to the GS-equilibrium value of 0.5. In contrast, in Simulation 4, where the uninformed traders have less time to learn before the constraint forcing them to remain uninformed is removed (recall $\lambda_t = 1/2$ for $t \leq 50$ in Simulation 4), uninformed traders struggle. They are never able to take advantage of the more informative price because they never learn the price–signal relation [$\xi^u(L_t)$ never gets close to 1]. As with Simulation 2, the learning ability of the uninformed decreases rapidly as traders become informed. As λ_t moves to 1, an uninformed trader has fewer people to imitate and uninformed learning is driven only by random experimentation.

4. CONCLUSION

There are many possible extensions to the GA model presented. For example, one could look to improve the learning rate of the uninformed by incorporating a more sophisticated GA learning algorithm like "election" as used by Arifovic (1994). The election operator is designed to make the population more stable by testing a strategy (created by selection, crossover, and mutation) against recent experience and adding it to the population only if it would have done well. Such a modification would slow the rate of change in the population because many strategies would not be altered. This increases the ability of the uninformed traders to learn the endogenous price–signal relationship. The less volatile "target" for learning is easier to learn. However, it turns out that incorporating election makes little difference to the qualitative results. First, one can always find mutation and asset supply noise combinations such that the GS equilibrium is stable or unstable.[18] In addition, it turns out that the election operator does not make the population more stable. Since the fitness of traders is highly correlated within the informed and uninformed groups, election affects one group of agents disproportionately, leading to greater variability in the population.

A GA is a simple tool that captures the spirit of learning through imitation and experimentation. It is encouraging to note that the GA is able to replicate a complex rational expectations equilibrium. It also illustrates the reasons and conditions for why learning can fail to produce the standard equilibrium. In the simulations presented here, convergence depends on the level of the asset supply noise relative to the level of experimentation or noise in the GA. Unfortunately, the convergence of the simulations hinges on a relatively deep parameter of the GA. Whereas calibrating the selection or imitation component of the GA to experimental data seems feasible, the crucial mutation parameter that controls experimentation is much harder to observe. The importance of the hard-to-calibrate mutation parameter is a limitation of the GA tool.

NOTES

1. For a survey of empirical decision theory, see Schoemaker (1982). For the link to portfolio decisions, see, for example, Odean (1998) or Barber and Odean (2000).
2. Shiller (1999) surveys this large area of research. Examples of using individual behavior to explain market phenomena include using prospect theory to understand the equity premium puzzle [Benartzi and Thaler (1995)] and overweighting recent evidence to explain the apparent success of contrarian investment strategies [Lakonishok et al. (1994)]. Daniel et al. (1998) take a slightly different approach and construct an equilibrium model from the premise that traders are overconfident and suffer from biased self-attribution of success.
3. Examples include Black (1986) and De Long et al. (1990).
4. For example, the SEC is considering its regulatory response to the Internet and the increase in "inexperienced" trade and recommendations from on-line chat [Unger (1999)].
5. The GA adaptive learning that we consider is a non-Bayesian approach to learning. In contrast, Bayesian models of learning do not relax the unbounded rationality assumption and model learning through beliefs that are updated with Bayes' rule. In a GS-type model, Bray (1982) studies non-Bayesian least-squares or econometric learning and Bray and Kreps (1987) consider Bayesian learning.
6. This differs from the approach taken by LeBaron et al. (1999). Restricting to a single period also avoids the complex issue of implementing an adaptive learning algorithm in a dynamic setting [see Uhlig and Lettau (1999)].
7. In the simulation, the competitiveness is approximated by choosing a large population size, N.

8. The linearity of the informed traders' demands follows direcly from CARA utility and normality [assuming that they know B_0 and B_1 in equation (1)]. For the uninformed, the linearity requires a conjecture that the market-clearing price function [equation (4)] is linear. In the rational expectations equilibrium, the conjecture is true.

9. We set $\gamma^U = \{a[\sigma_z^2 + V(y|P, L^*)]\}^{-1}$, where $V[y|P, L^*]$ is calculated at the GS rational expectations equilibrium, L^*. Again, this is simply a normalizing constant and is chosen only to facilitate interpretation.

10. To convert the finite binary string to a real number, the minimum and maximum values of the parameter $\beta_{jt}^{\iota n}$ need to be specified. These minimum and maximum values are chosen so that the GS rational expectations parameter values lie in the middle of the range. See Table 1 for the parameters used.

11. For crossover, the breakpoint in the string is determined randomly. As is standard in the GA, with some probability, crossover does not occur and the first string is passed unchanged to ℓ_{t+1}^n. Mutation involves randomly swapping a "0" for "1" or vice versa. The probability that any one bit is mutated is reported in Table 1 along with the crossover parameter and other GA parameters. The GA code used here is based on that of Goldberg (1989).

12. If f^n is the expected utility and $\bar{f}^n$ is the realized average utility over R independent repetitions of the economy (with a fixed population), then we can write $f^n = \bar{f}^n + \tilde{\epsilon}^n(R)$. The variance of $\tilde{\epsilon}^n(R)$ is decreasing in R. Practically, for smaller choices of R, one needs to work with the CARA utility function, $-\exp(-aW_{t1}^n)$, directly rather than the mean variance form in equation (5) [note that $R = 1$ implies $v(W_{t_r 1}^n) = 0$]. To use the CARA utility directly, we need to add a constant to fitness since fitness-proportional selection requires a positive fitness number. The alternative in this case is to use a rank-based selection mechanism.

13. Table 3 summarizes the results of repeating each simulation 100 times. The table reports the frequency that the simulation converges to the GS rational expectations value for the proportion of informed traders. There is little qualitative difference in the behavior of the other parameters across simulations that converge or across simulations that do not.

14. Along with the GS equilibrium parameters, the "rational" parameters are shown. These are the parameters an individual would use if she knew the details of the population in L_t. This is both more information and more rationality than the adaptively learning traders possess in our model.

15. The ratio is normalized so that we can interpret $f^i/f^u > 1$ as implying a likely increase in the proportion of informed traders.

16. For trader i, $\xi^i = 1$ if and only if $\beta_{0t}^{ii} = \beta_0^{i^*}$ and $\beta_{1t}^{ii} = \beta_1^{i^*}$. For trader u, $\xi^u(L_t) = 1$ if and only if $\beta_{0t}^{uu} = \beta_0^{u^*}$ and $\beta_{1t}^{uu} = \beta_1^{u^*}$. Recall from equation (3) that the fitness of a given strategy ℓ^n will not depend on all the elements in the vector since ι^n determines which asset-demand parameters are used.

17. The degree of experimentation and randomness in the GA is controlled primarily by the mutation probability. However, other sources of randomness, such as that from the selection step, also contribute to the magnitude of experimentation. This is why the measurement of fitness in equation (5) will affect the specific results (see note 12).

18. This is consistent with the analysis of deterministic learning studied by Routledge (1999).

REFERENCES

Arifovic, J. (1989) Learning by Generic Algorithms in Economic Environments. Working paper 90-001, Sante Fe Institute.

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

Arifovic, J. (1996) The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy* 104, 510–542.

Arifovic, J. & B.C. Eaton (1998) The evolution of type communication in a sender/receiver game of common interest with cheap talk. *Journal of Economic Dynamics and Control* 22, 1187–1208.

Axelrod, R.A. (1987) The evolution of strategies in the iterated prisoner's dilemma. In L. Davis (ed.), *Genetic Algorithms and Simulated Annealing*. London: Pitman.

Barber, B.M. & T. Odean (2000) Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55, 773–806.

Benartzi, S. & R.H. Thaler (1995) Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110, 73–92.

Black, F. (1986) Noise. *Journal of Finance* 41, 529–543.

Blume, L. & D. Easley (1992) Evolution and market behavior. *Journal of Economic Theory* 58, 9–40.

Bray, M. (1982) Learning, estimation and the stability of rational expectations. *Journal of Economic Theory* 26, 318–339.

Bray, M.M. & D. Kreps (1987) Rational learning and rational expectations. In G.R. Feiwel (ed.), *Arrow and the Ascent of Modern Economic Theory*. Houndmills, NY: New York University Press.

Bullard, J. & J. Duffy (1997) A model of learning and emulation with artificial adaptive agents. *Journal of Economic Dynamics and Control* 22, 179–207.

Daniel, K.D., D. Hirshleifer & A. Subrahmanyam (1998) Investor psychology and security market under- and overreactions. *Journal of Finance* 53, 1839–1885.

De Long, J.B., A. Shleifer, L.H. Summers & R.J. Waldmann (1990) Noise trader risk in financial markets. *Journal of Political Economy* 98, 703–737.

Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*. Reading, MA: Addison-Wesley.

Grossman, S.J. & J.E. Stiglitz (1980) On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.

Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.

Keynes, J.M. (1936) *The General Theory of Employment Interest and Money*. London: Macmillan.

Lakonishok, J., A. Shleifer & R.W. Vishny (1994) Contrarian investment, extrapolation and risk. *Journal of Finance* 49, 1541–1578.

LeBaron, B., W.B. Arthur & R. Palmer (1999) Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487–1516.

Marimon, R., E. McGrattan & T.J. Sargent (1990) Money as a medium of exchange with artificially intelligent agents. *Journal of Economic Dynamics and Control* 14, 329–373.

Odean, T. (1998) Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.

Routledge, B.R. (1999) Adaptive learning in financial markets. *Review of Financial Studies* 12, 1165–1202.

Schoemaker, P.J.H. (1982) The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature* 20, 529–563.

Shiller, R.J. (1999) Human behavior and the efficiency of the financial system. In J.B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, Vol. 1, pp. 1305–1340. Amsterdam: North-Holland.

Uhlig, H. & M. Lettau (1999) Rules of thumb versus dynamic programming. *American Economic Review* 89, 148–174.

Unger, L.S. (1999) Online Brokerage: Keeping Apace of Cyberspace. Discussion paper, Securities and Exchange Commission.

https://doi.org/10.1017/S1365100501019083 Published online by Cambridge University Press