Macroeconomic Dynamics, **5**, 2001, 204–224. Printed in the United States of America.

PERFORMANCE OF RATIONAL AND BOUNDEDLY RATIONAL AGENTS IN A MODEL WITH PERSISTENT EXCHANGE-RATE VOLATILITY

JASMINA ARIFOVIC
Simon Fraser University

The model is a two-country overlapping generations economy with boundedly rational agents who update their decision rules using a version of the stochastic replicator dynamic. The results show that stationary rational expectations equilibria of this model are unstable under this type of evolutionary adaptation. The paper also derives a two-period-ahead forecast of the values of average fractions of savings placed in each of the two currencies. This forecast is used in decisionmaking of a rational agent who has a full knowledge of the evolutionary economy. The performance of the rational agent is compared to the performance of boundedly rational agents, based on the average utility received over time. Results show that the difference between utilities earned by rational and boundedly rational agents is small. In addition, the average utility of the best-performing boundedly rational agents is higher than the average utility of the rational agents.

Keywords: Bounded Rationality, Exchange-Rate Fluctuations

1. INTRODUCTION

Persistent fluctuations of exchange rates have proven to be an empirical phenomenon that is difficult to explain theoretically. The studies by Meese and Rogoff (1983a, b), for instance, show that the structural models aimed at explaining the behavior of the exchange rate in terms of the fundamentals fail to improve on the random-walk out-of-sample forecasting accuracy. It seems that no model based on fundamentals such as money supplies, real income, interest rates, inflation rates, and current account balances has been able to explain or predict a high percentage of variation in the exchange rate, at least at short- or medium-term frequencies. Nonstructural models (univariate and vector autoregression) have not had much more success in terms of outperforming the random walk's forecasting accuracy. [See Frankel and Rose (1995) for a survey of models with fundamentals and of time-series models of the exchange-rate behavior.]

I would like to thank Albert Marcet for useful discussions regarding the main theme of this paper. I would also like to thank Peter Bossaerts, Jim Bullard, Tim Cogley, and participants at the seminars at the University of Amsterdam, Federal Reserve Bank of San Francisco, European Central Bank, and California Institute of Technology for helpful comments. Address correspondence to: Jasmina Arifovic, Department of Economics, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada; e-mail: arifovic@sfu.ca.

 1365-1005/01 $9.50

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

Explicit modeling of the changes in agents' expectations might be a possible way to capture and explain the exchange-rate volatility. Arifovic (1996) studies a two-country overlapping generations environment with boundedly rational agents. Agents make savings and portfolio decisions and update their decision rules using a genetic algorithm. The results show that stationary rational expectations equilibria of this model are unstable under genetic algorithm adaptation, resulting in persistent fluctuations of the nominal exchange rate. The observed instability and fluctuations of the exchange rate are due to the interaction between the evolutionary algorithm and the underlying structure of the model, which is characterized by the indeterminacy of rational expectations equilibria.[1]

This paper studies the same two-country overlapping generations environment. Agents are boundedly rational and use a version of the stochastic replicator dynamic to update their decision rules. The model displays the same type of instability of the stationary rational expectations equilibria and persistence in the behavior of the nominal exchange rate as the one presented by Arifovic (1996). This is not surprising because both the genetic algorithm and the stochastic replicator dynamic share a set of common principles that govern the evolution of decision rules. However, they do differ in the specific implementation of these principles, such as in the representation of beliefs and the types of operators that are used for updating. The paper shows that, despite these differences, the main qualitative features of the observed behavior are the same for the two algorithms. We are not aware of any other study that examines the robustness of results of evolutionary adaptation in economic environments under different rules' representation and different updating schemes. This result is important because it shows that the observed dynamics are robust to changes in specific algorithmic details or fine-tuning of the parameter values.[2]

Generally, applications of evolutionary algorithms in macroeconomic models are simulation-based because the environment is usually too complex to be characterized analytically. Making these algorithms more tractable represents a step toward a more formal characterization of their behavior and possibly toward the analysis of their convergence properties. The stochastic replicator dynamic used in this paper allows a degree of analytical tractability that could not be obtained with a genetic algorithm model. We derive one-period- and two-period-ahead forecasts of the model's endogenous variables and we add to a model a rational agent who uses this forecast to make portfolio decisions.[3] We compare the performance of generations of rational agents to the performance of generations of boundedly rational agents in terms of their average utilities.

A number of recent papers that examine the statistical learning models have compared the performance of rational and boundedly rational agents in terms of their forecasting errors [Marcet and Nicollini (1999), Bullard and Duffy (2001), and Hommes and Sorger (1998)]. The objective is to make a tighter link between models of bounded rationality and rational expectations models, and to provide restrictions on the type of boundedly rational models that can be useful in economic modeling.

This paper follows the same line of research in an attempt to assess the performance of a model of bounded rationality. However, given that this is a model of bounded rationality with heterogeneous beliefs, in which agents do not make forecasts of the next period's prices, the same methodology that is based on the comparison of the forecasting errors cannot be used in this setup. Instead, the comparison is based on the utilities that the two types of agents receive. The results show that the best within-a-generation boundedly rational agent performs better than the rational agent. In addition, although the rational agent performs better than the average and the median boundedly rational agent, the differences in performance are relatively small.

The paper proceeds as follows: A description of the economic model is given in Section 2. Updating of agents' rules is described in Section 3. Dynamics of adaptation are examined in Section 4. Derivation of one-period- and two-period-ahead forecasts is provided in Section 5. That section also contains the comparison of the performance of the two types of agents. Finally, concluding remarks are presented in Section 6.

2. DESCRIPTION OF THE MODEL

The model is a version of the Karaken and Wallace (1981) two-country, overlapping generations economy. It is a pure endowment economy with fiat money. At each date t, $t \geq 1$, there are born N young people, in each country, said to be of generation t. They are young at period t and old at period $t+1$. Each agent of generation t is endowed with w^1 units of a single consumption good at time t, and w^2 of the good at time $t+1$ and consumes $c_t(t)$ of the consumption good when young and $c_t(t+1)$ of the good when old. Agents in both countries have common preferences given by

$$u_t[c_t(t), c_t(t+1)] = \ln c_t(t) + \ln c_t(t+1).$$

This is a free-trade, flexible-exchange-rate-regime environment in which agents in the two countries are permitted to freely borrow from and lend to each other and to hold each other's currencies. An agent of generation t solves the following maximization problem at time t:

$$\max \ln c_t(t) + \ln c_t(t+1)$$

so that

$$c_t(t) \leq w^1 - \frac{m_1(t)}{p_1(t)} - \frac{m_2(t)}{p_2(t)},$$

$$c_t(t+1) \leq w^2 + \frac{m_1(t)}{p_1(t+1)} + \frac{m_2(t)}{p_2(t+1)},$$

where $m_1(t)$ are the agent's nominal holdings of currency 1, $m_2(t)$ are the agent's nominal holdings of currency 2 acquired at time t, $p_1(t)$ is the nominal price of the good in terms of currency 1 at time t, and $p_2(t)$ is the nominal price of the good in terms of currency 2 at time t. The agent's savings, $s(t)$, in the first period of life, are equal to the sum of real holdings of currency 1, $m_1(t)/p_1(t)$, and real holdings of currency 2, $m_2(t)/p_2(t)$.[4]

The exchange rate $e(t)$ between the two currencies is given by $e(t) = p_1(t)/p_2(t)$. Because there is no uncertainty in the model, an equilibrium condition requires equal rates of return on all assets. Thus, the rates of return on currency 1 and currency 2, $R_1(t+1)$ and $R_2(t+1)$, respectively, have to be equal to

$$R(t+1) = \frac{p_1(t)}{p_1(t+1)} = \frac{p_2(t)}{p_2(t+1)}, \qquad t \geq 1, \tag{1}$$

where $R(t)$ is the gross real rate of return between t and $t+1$. Rearranging (1), we obtain

$$\frac{p_1(t+1)}{p_2(t+1)} = \frac{p_1(t)}{p_2(t)}, \qquad t \geq 1. \tag{2}$$

From equation (2), it follows that the exchange rate is constant over time:

$$e(t+1) = e(t) = e, \qquad t \geq 1. \tag{3}$$

An individual's savings $s(t)$ that are derived from the agent's maximization problem are given by

$$s(t) = \frac{m_1(t)}{p_1(t)} + \frac{m_2(t)}{p_2(t)} = \frac{1}{2}\left[w^1 - w^2\frac{1}{R(t)}\right]. \tag{4}$$

The equilibrium condition in the loan market requires that aggregate savings equal real-world money supply; that is,

$$S(t) = N\left[w^1 - w^2\frac{p_1(t+1)}{p_1(t)}\right] = \frac{H_1(t)}{p_1(t)} + \frac{H_2(t)e}{p_1(t)}, \tag{5}$$

where $H_1(t)$ is the nominal supply of currency 1 at time t, and $H_2(t)$ is the nominal supply of currency 2 at time t. The supply of each currency is kept constant and thus the amount of currency 1 is given by $H_1(t) = H_1(0) = H_1$ for all t, and the amount of currency 2 is given by $H_2(t) = H_2(0) = H_2$ for all t.

The exchange-rate indeterminacy proposition [Kareken and Wallace (1981)] states that if there exists a monetary equilibrium in which both currencies are valued at some exchange rate e, then there exists a monetary equilibrium at any exchange rate $e \in (0, \infty)$. Consider an exchange rate $\hat{e}$, $\hat{e} \neq e$, and the price sequences $\{\hat{p}_1(t)\}$ and $\{\hat{p}_2(t)\}$, $\hat{p}_1(t) \neq p_1(t)$, and $\hat{p}_2(t) \neq p_2(t)$ for $t \geq 1$ such that

$$\hat{p}_1(t) = \frac{(H_1 + \hat{e}H_2)\,p_1(t)}{H_1 + eH_2} \tag{6}$$

and

$$\hat{p}_2(t) = \hat{p}_1(t)/\hat{e}. \tag{7}$$

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

The price sequences defined in equations (5) and (6) result in the same sequence of real rates of return as the price sequences $\{p_1(t)\}$ and $\{p_2(t)\}$ and, in turn, in the same values of individual and aggregate savings. Solving (5) for $\hat{p}_1(t)$ and substituting into (4) gives the following equilibrium condition:

$$S(t) = \frac{H_1 + \hat{e}H_2}{\hat{p}_1(t)}. \quad \textbf{(8)}$$

Price levels $\hat{p}_1(t)$ and $\hat{p}_2(t)$ adjust enough to achieve identical values of savings in a monetary equilibrium with the exchange rate e and in a monetary equilibrium with the exchange rate $\hat{e}$. Except for the initially old, who experience different consumption allocations for different initial nominal price levels, all other generations face the same consumption allocations in the equilibrium with the exchange rate e as they do in the equilibrium with the exchange rate $\hat{e}$.

The indeterminacy of the exchange rate in this model results from the fact that there is only one equation for the real-world money demand [equation (5)]. The equations for the individual real demands for each currency are therefore not well defined. Note that if there were a restriction that residents of each country could use only their country's currencies, real money demands for currency 1 and currency 2 would be well defined and equal to the respective real money supplies of the two currencies.

For a given exchange rate e, $e \in (0, \infty)$, there is a stationary equilibrium with constant price levels, constant rates of return on two currencies, and Pareto-optimal, constant consumption allocations such that $c_t(t) = c^{1,*} = c_t(t+1) = c^{2,*}$.

3. BOUNDEDLY RATIONAL AGENTS

There are two classes of boundedly rational agents in the economy. One class makes decisions in every odd period, and the other class makes decisions in every even period. Each class of agents is represented by a population of decision rules. Thus, at each t, there are two populations of rules, one that represents young agents of generation t, and the other that repesents old agents of generation $t-1$. A decision rule of agent i of generation t is given by a string that consists of two real numbers. The first number is used to determine agent i's savings, $s_i(t) \in [0, w^1]$. The second number is used to determine agent i's portfolio fraction, $\lambda_i(t) \in [0, 1]$. The portfolio fraction defines the fraction of an individual's savings $s_i(t)$ that are placed in currency 1. Thus agent i of generation t places the amount of $\lambda_i(t)s_i(t)$ into currency 1, and the remaining part given by $(1 - \lambda_i(t))s_i(t)$ in currency 2.

Aggregate savings in terms of currency 1 and 2 are used to determine the nominal price levels, $p_1(t)$ and $p_2(t)$:

$$p_1(t) = H_1 \bigg/ \sum_i^N \lambda_i(t)s_i(t), \quad \textbf{(9)}$$

$$p_2(t) = H_2 \bigg/ \sum_i^N (1 - \lambda_i(t))s_i(t). \quad \textbf{(10)}$$

Given the market-clearing prices, $p_1(t)$ and $p_2(t)$, and the fraction $\lambda_i(t)$, the nominal holdings of currency 1 and currency 2, $m_{1,t}$ and $m_{2,t}$, of agent i, $i \in [1, N]$, of generation t are determined:

$$m_{i,1}(t) = \lambda_i(t) s_i(t) p_1(t) \tag{11}$$

$$m_{i,2}(t) = (1 - \lambda_i(t)) s_i(t) p_2(t). \tag{12}$$

In period $t+1$, agents of generation t use all of their money balances to purchase the consumption good at prices that clear the markets at $t+1$. Second-period consumption of agent i is then given by

$$c_{i,t}(t+1) = w^2 + \frac{m_{i,1}(t)}{p_1(t+1)} + \frac{m_{i,2}(t)}{p_2(t+1)}. \tag{13}$$

At the end of period $t+1$, agents' utilities are computed on the basis of their first- and second-period consumption values. These utilities are used to determine fitness values of rules that were used by the members of generation t. The *fitness*, $\mu_{i,t}$, of a rule i is given by the ex-post value of the utility function of agent i of generation t:

$$\mu_{i,t} = \ln c_{i,t}(t) + \ln c_{i,t}(t+1). \tag{14}$$

3.1. Updating

At the end of a cycle of two periods t and $t+1$, agents update their rules by imitating rules that have proven to be relatively successful and by occasionally experimenting with new rules. Prior to adoption of new rules, their performance is tested using data from periods t and $t+1$. If the experimentation takes place, an agent tests the performance of a new rule using the election operator [Arifovic (1994)]. Imitation, experimentation, and election are applied in the following way.

3.1.1. Imitation. Agents of generation $t+2$ inherit rules from members of generation t. Rules that were more successful at the end of period $t+1$ are more likely to be inherited. Each young agent of generation $t+2$ is assigned a copy of one of the rules of generation t. The probability that a rule i of generation t is assigned is equal to its relative fitness and is given by

$$\Pr_i(t+2) = \frac{\mu_{i,t}}{\sum_{i=1}^{N} \mu_{i,t}}. \tag{15}$$

Each rule occupies an interval of a measure equal to its relative fitness in the interval [0, 1]. A copy of a rule is assigned in the following way: For each new agent j of generation $t+2$, $j \in [1, N]$, a random number between 0 and 1, r_j, is

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

drawn from a uniform distribution. A rule i, $i \in [1, N]$, that occupies the range of values where r_j belongs is determined. Then, a copy of that rule is assigned to agent j. This copy is denoted $C_j(t+2)$.

These steps are repeated N times to generate a population of N copies of the rules of generation t. The process promotes rules with high fitness values that are, on average, imitated more frequently.

3.1.2. Experimentation. Next, every agent is given an opportunity to experiment with new rules. An agent experiments with only one of the two parts of the rule, either savings or portfolio decision. The part of the rule that undergoes experimentation is randomly determined. Both parts have equal probability of being selected for experimentation. Once one of the two parts is selected, experimentation takes place with probability π_{ex}. In case that experimentation takes place on a savings part of the rule, a new rule is determined by drawing a random number from the uniform distribution in the interval $[0, w^1]$, and in case that experimentation takes place on a portfolio fraction part, a new rule is determined by drawing a random number from the uniform distribution in the interval [0, 1]. Denote a resulting rule that belongs to agent j, $j \in [1, N]$, by $E_j(t+2)$.

3.1.3. Election operator. Prior to final determination of the population of rules of generation $t+2$, each new rule $E_j(t+2) \neq C_j(t+2)$ that was generated via experimentation is evaluated using the previous period's rates of return on the two currencies. Thus a *potential fitness* of a new rule is calculated. The value of potential fitness of a new rule $E_j(t+2)$ is compared to the fitness value of a copy $C_j(t+2)$ that was assigned to agent j. If it is higher than the fitness value of the copy $C_j(t+2)$, the new rule $E_j(t+2)$ replaces the copy of the old rule and is accepted into the population of rules of generation $t+2$. However, if its fitness value is lower than the fitness of the old rule $C_j(t+2)$, an agent j keeps the copy of the old rule and the new rule $E_j(t+2)$ is discarded.

At $t=1$, two populations of rules that will represent two classes of agents are randomly generated. These populations begin as populations of rules of agents of generation 0 (initially old) and generation 1 (initially young). The economy is simulated for $T_{\max}$ periods.

4. DYNAMICS OF ADAPTATION

The dynamics of exchange-rate behavior exhibit persistent volatility with no sign of settling to a constant value (Figure 1). At the same time, average values of the first-period consumption and savings remain close to the stationary equilibrium values (Figure 2). The exhibited dynamics are robust in regard to the changes in the parameter values. The observed persistence in the simulated data is due to the joint effects related to the indeterminacy of equilibria and the evolutionary dynamics. Arifovic (1996) demonstrates that a stationary equilibrium of this model in which

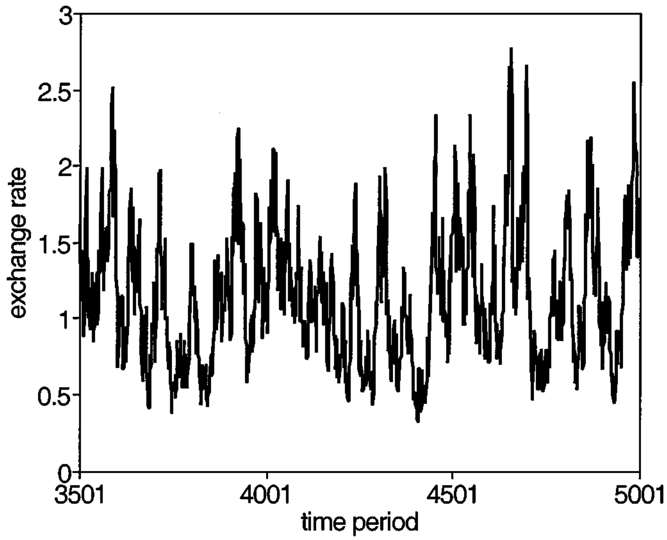

FIGURE 1. Exchange-rate behavior.

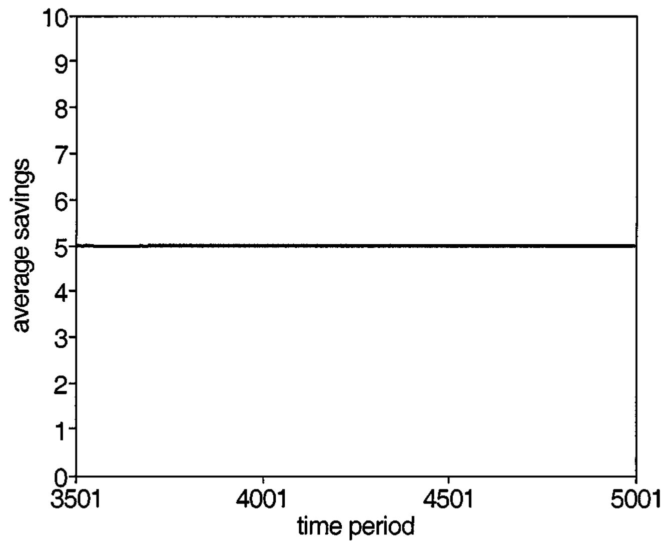

FIGURE 2. Average savings.

both currencies are valued is unstable under the genetic algorithm adaptation. The same argument can be used for the stochastic replicator dynamic to show the instability of a stationary equilibrium.

Suppose that the economy has been in a stationary equilibrium in $t-1$ and t. All agents make the same savings and portfolio fraction decisions. The nominal

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

price levels in terms of the two currencies are constant, and the rates of return on the two currencies are equal.

At the beginning of $t+1$, the updating of decision rules takes place. Imitation has no effect as all the rules are identical. Experimentation brings in new rules, different from the equilibrium ones, but whether they become members of the actual population of rules of $t+1$ depends on the election operator. This operator will admit all those rules whose portfolio fractions decode to numbers different from the stationary equilibrium ones, but that still have stationary equilibrium values of savings. They pass the election operator test because their fitness is evaluated at the previous-period rates of return on two currencies. Because the economy was in a stationary equilibrium in the previous two periods, the two rates of return are equal. This is the reason why the new rules with stationary values of the savings and portfolio fractions different from the stationary ones will have potential fitness value equal to the fitness value of the old equilibrium rules. From the standpoint of the election operator, actual fractions placed in each currency do not matter because the rates of return are equal.

Once the diversity is brought into the populations, the rates of return on the two currencies will no longer be equal. Further adaptation will favor those decision rules that place higher fractions of savings into the currency with a higher rate of return. Consequently, even if the economy reaches a stationary equilibrium by chance or if it is initialized at a stationary equilibrium, the evolutionary dynamics will take it away from that stationary equilibrium.

Thus, a stationary equilibrium of this model is evolutionarily stable with respect to invading new rules $E_j(t+1)$ whose savings decisions $s_j(t+1) \neq s^*$, where s^* is a stationary value of savings, and $\lambda_i(t+1) \neq \lambda^*$, where λ^* is a stationary value of the portfolio fraction. However, it is not evolutionarily stable with respect to invading rules $E_j(t+1)$ with $s_j(t+1) = s^*$ and $\lambda_j(t+1) \neq \lambda^*$.

In general, the out-of-equilibrium heterogeneity of the portfolio fraction values results in the inequality of the rates of return on two currencies. Agents seek to exploit this arbitrage opportunity by placing larger fractions of their savings into the currency that had a higher rate of return in the previous period. If the aggregate change of the portfolio fraction is large enough, the direction of the inequality is preserved and the value of the currency with the higher rate of return increases. On the other hand, if the aggregate change is not large enough, the reversal of the inequality of the rates of return occurs. The reversal will prompt the agents to place more savings into the currency whose value was decreasing prior to the reversal. As a result, the exchange rate changes the direction of movement. These dynamics bring about the fluctuations in the portfolio fraction and the exchange rate that persist over time. Intermediate values of the average portfolio fraction are more likely to be observed than the values closer to 1 or 0. Table 1 reports the frequency distributions of the average portfolio fraction for three values of the rate of experimentation. For each π_{ex}, five simulations initialized with different random seed numbers were conducted. Each simulation lasted for 10,000 periods. The table shows that the values tend to be more concentrated around the intermediate

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

Table 1. Frequency distribution of portfolio fraction values

	π_{ex}		
λ values	0.0033	0.033	0.33
0.1	0.0284	0.00	0.00
0.2	0.1067	0.00	0.00
0.3	0.1442	0.03	0.0001
0.4	0.2066	0.11	0.46
0.5	0.1566	0.31	0.4954
0.6	0.2772	0.37	0.0215
0.7	0.0695	0.15	0.00
0.8	0.0078	0.03	0.00
0.9	0.003	0.00	0.00
1.0	0.00	0.00	0.00

values for higher rates of mutation and more dispersed for lower rates of mutation. Nevertheless, regardless of the rate of experimentation, the most frequent are the intermediate values in the range between 0.4 and 0.6.

It is worthwhile to point out that, in experiments with human subjects [Arifovic (1996)] in which the same model was simulated, most of the portfolio fraction values were concentrated in the interval [0.4, 0.6): 47% in the interval [0.4, 0.5) and 43% in the interval [0.5, 0.6). In general, the evolutionary dynamics capture the features of the behavior observed in the experiments with human subjects. This behavior was characterized by the fluctuations of the exchange rates that did not settle down over time and by the levels of savings that converged and stayed at the values close to the stationary equilibrium values.

In general, the qualitative features of the dynamics of the evolutionary algorithm presented in this paper and of the genetic algorithm are very similar. In the genetic algorithm application, rules are represented by binary strings that are decoded and normalized to give real-number values of savings and portfolio fractions. They are updated using reproduction, crossover, mutation, and election. Reproduction plays the same role as imitation, while mutation has the same role as experimentation. However, crossover, the operator that performs recombination of the parts of existing binary strings has no counterpart in the current setup. Crossover, which plays an important role in the genetic algorithm adaptation, could be implemented in this framework as well.[5] Thus, the results show that the main dynamics are preserved under different types of rules' representation and application of different operators.

5. MODEL WITH A RATIONAL AGENT

Consider now a model in which, in addition to boundedly rational agents, a *rational* agent is born at each t. Each rational agent lives for two periods. When young, a rational agent makes savings and portfolio decisions, and at the end of the old

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

age, her utility is evaluated. Thus, in every t, there are two rational agents, one of generation $t-1$ and the other of generation t. Rational agents' decisions do not affect the outcomes of the economy. Their utilities are calculated using their own decisions and the relevant rates of return realized in the evolutionary economy as a result of decisions made by boundedly rational agents.

The performance of a rational agent is compared to the performance of the boundedly rational agents of the same generation. Utilities of three *classes* of boundedly rational agents are used for comparison. One is the *average* agent whose utility is computed as the average of all utilities of agents of a given generation. The second is the *median* agent whose utility is equal to the median in a given generation. Finally, the third is the *best within-generation* agent that received the highest utility in a given generation.

This analysis is done to define benchmarks that can be used to evaluate the performance of the evolutionary model. If an average boundedly rational agent does not do much worse (in terms of the received utility) than the rational agent, then it can be argued that the model can be used as a reasonable model of agents' adaptation. In addition, in a model with heterogeneous decision rules, there can be agents who perform better than rational agents.

Consider a rational agent born at the beginning of period $t+1$. She has to make optimal savings and portfolio fraction decisions, $s^r(t+1)$ and $\lambda^r(t+1)$. The agent consumes $c^r_{t+1}(t+1)$ and $c^r_{t+1}(t+2)$, and receives utility u^r_{t+1} at the end of period $t+2$.[6] To simplify the analysis, we assume the endowment pattern $w^1 > 0$ and $w^2 = 0$. For this endowment pattern (and previously defined preferences), the optimal savings decision, $\hat{s}(t)$, is always equal to $w^1/2$ regardless of the rate of return. With this simplification, we only need to derive the rational agent's optimal portfolio decision. This decision is based on one-period-ahead forecast of the rates of return on two currencies and two-period-ahead forecast of the average portfolio fraction.

5.1. Derivation of Forecasts

5.1.1. Forecast of average savings at $t+1$. We first derive an expression for $s^e(t+1)$, the expected value of savings at $t+1$. Savings at $t+1$ are determined by applying imitation, experimentation, and election on the population of rules of generation $t-1$. Consider first the impact of imitation. Denote by $s^{m,e}(t+1)$ the expected value of the outcome of imitation at $t+1$. It is given as a weighted average of the individual values of savings at $t-1$, where weights are given as relative fitness values:

$$s^{m,e}(t+1) = \sum_{i=1}^{N} \Pr_i(t+1) s_i(t-1) \quad \textbf{(16)}$$

where $\Pr_i(t+1)$ is relative fitness of a rule i given by $\mu_{i,t-1}/\sum_{i=1}^{N}\mu_{i,t-1}$. As noted in the description of the algorithm, the savings part of a decision rule is selected for experimentation with probability 0.5. Thus, on average, for half of the population of savings decision rules, the updating is completed with the application of the imitation operator. The expected value of this fraction of savings decisions is thus equal to $s^{m,e}(t+1)$.

The other half of savings decision rules are selected for further modification using experimentation and election.[7] Denote by $s^{p,e}(t+1)$ the expected value of savings decisions that result from the application of these two operators. This expected value consists of three parts.

The first part, $\prod_1^s(t+1)$,[8] refers to those savings decisions that do not undergo experimentation. Their expected value is equal to $s^{m,e}(t+1)$. And their contribution to the value of $s^{p,e}(t+1)$ is weighed by $(1-\pi_{ex})$, the probability that experimentation does not take place. Thus, $\prod_1^s(t+1)$ is given by

$$\prod\nolimits_1^s(t+1) = (1-\pi_{ex})s^{m,e}(t+1). \quad \textbf{(17)}$$

Experimentation takes place with probability π_{ex} and is implemented by drawing a random number in the interval $[0, w^1]$ from the uniform distribution. However, not all newly generated values pass the election operator test. The second part, $\prod_2^s(t+1)$, is related to the expected value of those newly generated savings decisions that pass the election operator test.

Let us define the *range of admissible values of savings*, that is, those values that pass the election operator test. This range is defined by the value of $s^{m,e}(t+1)$ and optimal savings, $\hat{s}(t)$. Let $\Delta\hat{s}^m(t+1)$ denote the absolute value of a difference between $\hat{s}$ and $s^{m,e}(t+1)$:

$$\Delta\hat{s}^m(t+1) = |\hat{s} - s^{m,e}(t+1)|.$$

Then, the range of admissible values is given by $[\hat{s}-\Delta\hat{s}^m(t+1), \hat{s}+\Delta\hat{s}^m(t+1)]$. Any value of savings in this range will result in a higher fitness value than the savings decisions with the expected value of $s^{m,e}(t+1)$. The expected value of this admissible range is equal to

$$\frac{\hat{s}+\Delta\hat{s}^m(t+1)+\hat{s}-\Delta\hat{s}^m(t+1)}{2} = \hat{s}. \quad \textbf{(18)}$$

Since these values are drawn from the uniform distribution, in the interval $[0, w^1]$, the probability that a value is in this admissible range is equal to

$$\frac{2\Delta\hat{s}^m(t+1)}{w^1}. \quad \textbf{(19)}$$

Overall, the expected value of $\prod_2^s(t+1)$ is equal to

$$\prod\nolimits_2^s = \pi_{ex}\frac{2\Delta\hat{s}^m(t+1)}{w^1}\hat{s} = \pi_{ex}\Delta\hat{s}^m(t+1). \quad \textbf{(20)}$$

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

Finally, $\prod_3^s(t+1)$ defines the contribution of unsuccessful experimentations. Savings decisions that will not pass the election operator are those that fall outside the range of admissible values. Thus, these are the values in the range $[w^1 - (\hat{s} + \Delta\hat{s}^m(t+1)]$, and in the range $[0 + (\hat{s} - \Delta\hat{s}^m(t+1))]$. The probability that these values with lower fitness values are drawn via experimentation is

$$\frac{w^1 - (\hat{s} + \Delta\hat{s}^m(t+1))}{w^1} + \frac{0 + (\hat{s} - \Delta\hat{s}^m(t+1))}{w^1} = 1 - \frac{2\Delta\hat{s}^m(t+1)}{w^1}. \quad \textbf{(21)}$$

Thus, the value of $\prod_3^s(t+1)$ is given by multiplying this probability by the probability of experimentation, π_{ex}, and the value $s^{m,e}(t+1)$ since the old value of savings is kept in case of an unsuccessful experimentation:

$$\prod\nolimits_3^s(t+1) = \pi_{ex}\left[1 - \frac{2\Delta\hat{s}^m(t+1)}{w^1}\right]s^{m,e}(t+1). \quad \textbf{(22)}$$

The value of $s^{p,e}(t+1)$ is the sum of $\prod_1^s(t+1)$, $\prod_2^s(t+1)$, and $\prod_3^s(t+1)$ and is equal to

$$s^{p,e}(t+1) = (1-\pi_{ex})s^{m,e}(t+1) + \pi_{ex}\frac{2\Delta\hat{s}^m(t+1)}{w^1}\hat{s}$$
$$+\pi_{ex}\left[1 - \frac{2}{w^1}\Delta\hat{s}^m(t+1)\right]s^{m,e}(t+1). \quad \textbf{(23)}$$

Finally, since a savings decision is selected for experimentation (and election) with probability $1/2$, the expected value of savings $s^e(t+1)$ is the weighted average of $s^{m,e}(t+1)$ and $s^{p,e}(t+1)$:

$$s^e(t+1) = \frac{1}{2}s^{m,e}(t+1) + \frac{1}{2}s^{p,e}(t+1). \quad \textbf{(24)}$$

5.1.2. Forecast of average portfolio fraction at t+1. Next, we derive the expected value of the portfolio fraction at $t+1$, $\lambda^e(t+1)$. Steps used in this derivation are similar to those used for deriving $s^e(t+1)$. We have to consider the impact of imitation, experimentation, and election on the portfolio fraction decisions of the rules of generation $(t-1)$.

First, we determine the expected value of a portfolio fraction that is the outcome of the process of imitation, $\lambda^{m,e}(t+1)$. It is given by

$$\lambda^{m,e}(t+1) = \sum_{i=1}^{N} \mathrm{Pr}_i(t+1)\lambda_i(t-1). \quad \textbf{(25)}$$

The weights, $\mathrm{Pr}_i(t+1)$'s, are again equal to relative fitness values. Portfolio decision rules are selected for further modification with probability 0.5. Thus, on average, half of the rules are selected for the application of experimenation and election. For the other half, updating is completed with imitation, and their expected value is given by $\lambda^{m,e}(t+1)$.

Again, the expected value of decision rules that are selected for experimentation and election, $\lambda^{p,e}(t+1)$, consists of three parts. Since experimentation takes place with probability π_{ex}, a fraction $(1-\pi_{ex})$ of agents keeps the portfolio decisions that resulted from imitation. The expected value of these fractions is $\lambda^{m,e}(t+1)$. The first part, $\prod_1^{\lambda}(t+1)$,[9] represents a contribution by these values of portfolio decisions, and is given by

$$\prod\nolimits_1^{\lambda}(t+1) = (1-\pi_{ex})\lambda^{m,e}(t+1). \tag{26}$$

The second and the third parts capture the impact of implementation of experimentation and election. Experimentation takes place with probability π_{ex}, and if it takes place, a random number is drawn from the uniform distribution, in the interval [0, 1]. Then, election is implemented and its impact depends on the direction of inequality between the rates of return on two currencies.

(i) If $R_1(t) > R_2(t)$, the values of λ greater than $\lambda^{m,e}(t+1)$ will pass the election operator test because they will result in higher fitness values (more savings is placed in the currency with higher rate of return). Thus, the admissible range of values is $[\lambda^{m,e}(t+1), 1]$. The probability that the values greater than $\lambda^{m,e}(t+1)$ are drawn is given by $(1-\lambda^{m,e}(t+1))$, and the expected value of the admissible range is equal to $(1+\lambda^{m,e}(t+1))/2$. The value of the second part, $\prod_2^{\lambda,1}(t+1)$,[10] is given by

$$\prod\nolimits_2^{\lambda,1}(t+1) = \pi_{ex}(1-\overline{\lambda}^{m,e}(t+1))\frac{(1+\overline{\lambda}^{m,e}(t+1))}{2}. \tag{27}$$

The third part captures the contribution of decision rules that keep the old values of λ's because experimentation results in values that do not pass the election operator test. The probability of drawing numbers less than $\lambda^{m,e}(t+1)$ from the uniform distribution is equal to $\lambda^{m,e}(t+1)$. The old value, the result of imitation, is kept instead and this expected value is equal to $\lambda^{m,e}(t+1)$. Thus, the value of the third part, $\prod_3^{\lambda,1}(t+1)$, is given as

$$\prod\nolimits_3^{\lambda,1}(t+1) = \pi_{ex}(\lambda^{m,e}(t+1))^2. \tag{28}$$

Finally, $\lambda^{p,e}(t+1)$ is given by the sum of $\prod_1^{\lambda,1}(t+1)$, $\prod_2^{\lambda,1}(t+1)$, and $\prod_3^{\lambda,1}(t+1)$:

$$\lambda^{p,e}(t+1) = (1-\pi_{ex})\overline{\lambda}^{m,e}(t+1) + \pi_{ex}\frac{(1-(\lambda^{m,e}(t+1))^2)}{2} + \pi_{ex}(\overline{\lambda}^{m,e}(t+1))^2. \tag{29}$$

(ii) On the other hand, if $R_1(t) < R_2(t)$, the election operator will admit those newly generated rules that are smaller, on average, than $\lambda^{m,e}(t+1)$. The admissible range of values is given by $[0, \lambda^{m,e}(t+1)]$. The probability that the values from the admissible range are drawn is equal to $\lambda^{m,e}(t+1)$, and the expected value of

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

the numbers drawn from this interval is equal to $\lambda^{m,e}(t+1)/2$. Thus, the expected value of the second part, $\prod_2^{\lambda,2}(t+1)$,[11] is given by

$$\prod\nolimits_3^{\lambda,2}(t+1) = \pi_{ex}\lambda^{m,e}(t+1)\frac{\lambda^{m,e}(t+1)}{2}. \tag{30}$$

The third part, $\prod_3^{\lambda,2}(t+1)$, represents the part of the expected value contributed by those rules that do not pass the election operator test. The probability that these values are drawn is equal to $1-\lambda^{m,e}(t+1)$. The value of $\prod_3^{\lambda,2}(t+1)$ is given by multiplying this probability by the probability of experimentation π_{ex}, and the value $\lambda^{m,e}(t+1)$ since the old portfolio values are kept in case of unsuccessful experimentation:

$$\prod\nolimits_3^{\lambda,2}(t+1) = [(1-\lambda^{m,e}(t+1))]\pi_{ex}\lambda^{m,e}(t+1). \tag{31}$$

Again, the value of $\lambda^{p,e}(t+1)$ is given as the sum of $\prod_1^{\lambda,2}(t+1)$, $\prod_2^{\lambda,2}(t+1)$, and $\prod_3^{\lambda,2}(t+1)$:

$$\lambda^{p,e}(t+1) = (1-\pi_{ex})\lambda^{m,e}(t+1) + \pi_{ex}\frac{(\lambda^{m,e}(t+1))^2}{2} + \pi_{ex}(1-\lambda^{m,e}(t+1))\lambda^{m,e}(t+1). \tag{32}$$

The expected value of the portfolio fraction, $\lambda^e(t+1)$, is then given as the weighted average of $\lambda^{m,e}(t+1)$ and $\lambda^{p,e}(t+1)$:

$$\lambda^e(t+1) = \frac{1}{2}\lambda^{m,e}(t+1) + \frac{1}{2}\lambda^{p,e}(t+1). \tag{33}$$

5.1.3. Forecast of rates of return at t+1. Based on these forecasts, the expectations of rate of return on currency 1 and on currency 2, $R_1^e(t+1)$ and $R_2^e(t+1)$, can be derived. The value of $R_1^e(t+1)$ depends on the average value of the amount of savings placed in currency 1 at time t, $\overline{s_1}(t)$, and the expected value of the amount placed in currency 1 at time $t+1$. The value of $\overline{s_1}(t)$, which is known at the beginning of $t+1$, is given as

$$\overline{s_1}(t) = \frac{1}{N}\sum_{i=1}^{N}\lambda_i(t)s_i(t).$$

To derive the expected savings in currency 1, we need to take into account that, once imitation takes place, one half of the savings decisions are not modified further. However, for these same rules whose savings decisions are not changed, the portfolio decisions undergo experimentation and election. The expected value of savings in currency 1 of this fraction of rules is then given by $s^{m,e}(t+1)\lambda^{p,e}(t+1)$. The other half of the rules get their savings decision modified by experimentation and election. These rules also keep the portfolio fractions that resulted from

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

imitation. Putting these two together, the expected value of savings in currency 1 for this fraction of rules is then given by $s^{p,e}(t+1)\lambda^{m,e}(t+1)$. From this, it follows that the expected value of savings in currency 1 at $t+1$, $s_1^e(t+1)$, is given as

$$s_1^e(t+1) = (1/2)[s^{m,e}(t+1)\lambda^{p,e}(t+1) + s^{p,e}(t+1)\lambda^{m,e}(t+1)]. \tag{34}$$

Thus, the value of $R_1^e(t+1)$ is given by

$$R_1^e(t+1) = \frac{s_1^e(t+1)}{\bar{s}_1(t)}. \tag{35}$$

Similarly, the value of $R_2^e(t+1)$ depends on the average savings placed in currency 2 at time t, $\bar{s}_2(t)$, and on the expected value of savings in currency 2 at time $t+1$. The value of $\bar{s}_2(t)$ is given by

$$\bar{s}_2(t) = \frac{1}{N}\sum_{i=1}^{N}(1-\lambda_i(t)s_i(t)) \tag{36}$$

The value of $s_2^e(t+1)$, derived in the same manner as the expected savings in currency 1, is given by

$$s_2^e(t+1) = (1/2)[s^{m,e}(t+1)(1-\lambda^{p,e}(t+1)) + s^{p,e}(t+1)(1-\lambda^{m,e}(t+1))]. \tag{37}$$

The value of $R_2^e(t+1)$ is then computed as the ratio of the two values:

$$R_2^e(t+1) = \frac{s_2^e(t+1)}{\bar{s}_2(t)}. \tag{38}$$

5.1.4. Forecast of average portfolio fraction at $t+2$. With the expected values of rates of return on two currencies, the first stage of the process, that is, derivation of expected values of variables at the end of $t+1$, is completed. However, a rational agent also needs to compute the expected value of average portfolio fraction at $t+2$, $\lambda^e(t+2)$ because she will set her time $t+1$ portfolio decision equal to this value. Thus, the second stage involves derivation of $\lambda^e(t+2)$. The steps are similar to those implemented in derivation of $\lambda^e(t+1)$. Using the expected values of the rates of return at $t+1$, expected fitness values of members of generation t can be computed. Based on these expected fitness values, the probabilities of imitation are computed and used in calculation of the value of $\lambda^{p,e}(t+2)$.

Again, we first have to determine the expected value of a portfolio fraction that results from imitation, $\lambda^{m,e}(t+1)$:

$$\lambda^{m,e}(t+2) = \sum_{i=1}^{N} \Pr_i(t+2)\lambda_i(t). \tag{39}$$

Next, we derive the value of $\lambda^{p,e}(t+2)$ that is the expected value of an outcome of experimentation and election. As before, the value consists of three parts, explained earlier in detail. In addition, the outcome depends on the direction of inequalities of the expected rates of return at $t+1$, $R_1^e(t+1)$ and $R_2^e(t+1)$.

(i) If $R_1^e(t+1) > R_2^e(t+1)$ then the expectation of the average portfolio fraction at $t+2$ will be given by

$$\lambda^{p,e}(t+2) = \pi_{ex}(1-\lambda^{m,e}(t+2))\frac{(1+\lambda^{m,e}(t+2))}{2} + \pi_{ex}\lambda^{m,e}(t+2)\lambda^{m,e}(t+2) + (1-\pi_{ex})\lambda^{m,e}(t+2). \tag{40}$$

(ii) On the other hand, if $R_1^e(t+1) < R_2^e(t+1)$, then the expectation of the average portfolio fraction at $t+1$ will be given by

$$\lambda^{p,e}(t+2) = \pi_{ex}\lambda^{m,e}(t+2)\frac{\lambda^{m,e}(t+2)}{2} + \pi_{ex}(1-\bar{\lambda}^{m,e}(t+2))\lambda^{m,e}(t+2) + (1-\pi_{ex})\lambda^{m,e}(t+2). \tag{41}$$

Note that all of the equations for deriving the value of $\lambda^e(t+2)$ are equivalent to those used for deriving the values of $\lambda^e(t+1)$ except that the time counter is increased by one. So, finally, the value of $\lambda^e(t+2)$ is given as:

$$\lambda^e(t+2) = \frac{1}{2}\lambda^{p,e}(t+2) + \frac{1}{2}\lambda^{m,e}(t+2). \tag{42}$$

Given $\lambda^e(t+2)$, the rational agent makes a decision at the beginning of time $t+1$ to save the amount $\hat{s}$ and to commit a fraction $\lambda^e(t+2)$ of the savings to currency 1 and a fraction $(1-\lambda^e(t+2))$ to currency 2. The utility of rational agent of generation $t+1$ is evaluated at the end of period $t+2$. This can be justified by the atomistic assumption that rational agents do not make up a significant fraction of the economy. In the current setup, decisions made by rational agents do not affect market-clearing prices. If rational agents' decisions had had an impact on prices, forecasts would have had to have taken that into account. Since this complicates the derivation of the forecasts further, we focus our attention to the simpler case.

5.2. Performance of Rational and Boundedly Rational Agents

The performance of rational and boundedly rational agents is measured in terms of the utility earned at the end of every two-period cycle. The utility of the rational agent is compared to the average utility of the population, to the median utility,

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

and to the highest utility in the population of boundedly rational agents of the same generation.

Let $d_t^a = u_t^r - u_t^a$ denote the difference in the utility between the rational agent, u_t^r, and an average utility, u_t^a, of boundedly rational agents of generation t, received at the end of $t+1$. Thus, if the value of this variable is positive, the rational agent did better than the average-performing boundedly rational agent. If it is less than zero, the rational agent did worse than the average-performing boundedly rational agent. Let $\bar{d^a}$ be the average over $T_{\max}$ periods. Similarly, let d_t^m denote the difference between the utility of the rational agent and a median utility of boundedly rational agents of generation t, received at the end of $t+1$. The average over $T_{\max}$ periods is denoted by $\bar{d^m}$.

Finally, let d_t^b denote the difference in utility between the rational agent and the best-performing boundedly rational agent of generation t. It is given by $d_t^b = u_t^r - u_t^b$, where u_t^b is the utility received by the best-performing boundedly rational agent at the end of $t+1$, that is, the agent that received the highest utility in the entire generation. Let $\bar{d^b}$ be the average over $T_{\max}$ periods.

Calculations of the previously described measures were performed for simulations with three different rates of experimentation, 0.0033, 0.033, and 0.33. Reported results are based on five simulations for each rate of experimentation. Each simulation was conducted for 10,000 periods. The results show that, overall, the difference between utilities earned by rational and boundedly rational agents is small. It is greater for larger rates of mutation, indicating that boundedly rational agents do worse if they experiment more. At the same time, on average, the best-performing boundedly rational agents always do better than the rational agents. The difference is the same for the rates of experimentation of 0.033 and 0.33. It is smaller for the really low rateof mutation of 0.003. These results are presented in Table 2.

The overall percentile ranking of the rational agent for the same rates of experimentation are presented in Table 3. For the given rate of experimentation, the frequencies are computed by ranking utilities of the rational and N boundedly rational agents from the lowest to the highest at the end of each t for $T_{\max}$ periods in five simulations. Clearly, the ranking of the rational agent worsens with the increases in the rate of experimentation. Thus, higher rates of experimentation hurt both the rational agent and the average boundedly rational agent in terms of the

Table 2. Differences in utilities between rational and boundedly rational agents

	$\ln c_t(t) + \ln c_t(t+1)$		
π_{ex}	$\bar{d^a}$	$\bar{d^m}$	$\bar{d^b}$
0.0033	0.01	0.01	−0.004
0.033	0.028	0.01	−0.02
0.33	0.09	0.04	−0.024

https://doi.org/10.1017/S1365100501019046 Published online by Cambridge University Press

TABLE 3. Percentile ranking of rational agent

	π_{ex}		
Percentile	0.0033	0.033	0.33
10	0.01	0.00	0.00
20	0.01	0.00	0.00
30	0.02	0.02	0.00
40	0.03	0.05	0.00
50	0.02	0.09	0.01
60	0.03	0.12	0.01
70	0.05	0.14	0.09
80	0.04	0.14	0.33
90	0.06	0.13	0.36
100	0.75	0.26	0.20

utilities earned. However, for higher rates of experimenation, larger fractions of boundedly rational agents perform better than the rational agent.

6. CONCLUDING REMARKS

This paper examined the behavior of the exchange rate in the model with boundedly rational agents who update their savings and portfolio decisions using a version of the stochastic replicator dynamic, an evolutionary algorithm based on imitation and experimentation. While the savings decisions settle to the values in the neighborhood of the steady state, the portfolio fraction values do not settle toward the constant, steady-state values. The fluctuations of portfolio fraction values result in persistent exchange-rate fluctuations. This behavior is similar to the actual exchange-rate time series. In addition, the evolutionary dynamics capture the features of the behavior observed in the experiments with human subjects, mainly, persistent fluctuations of the exchange rate and levels of savings that converge and stay at the values close to the stationary equilibrium values. Moreover, there is similarity in the frequency distributions of portfolio fraction values. Intermediate values have the highest frequency in both simulated data and experiments with human subjects.

In the second part of the paper, a rational agent is added to the economy. The comparison of the long-run performance of the rational and boundedly rational agents shows little difference in terms of the average utilities earned over time. Moreover, as populations of boundedly rational agents remain heterogeneous in terms of their portfolio decisions, there are always boundedly rational agents who receive utilities that are higher than those received by rational agents.

Thus, as the analysis shows, boundedly rational agents do not do much worse than rational agents. In addition, the fact that best-performing boundedly rational agents perform better than rational agents might give some incentive for agents to follow the evolutionary rules in their decisionmaking. These results can be used to argue that the model of bounded rationality can be used as a reasonable model of agents' adaptation.

A critisicm of the models of bounded rationality has been that, because agents' behavior depends on the specific way in which their beliefs are modeled, these models can produce so many different outcomes that they become useless as instruments for generating predictions. Developing criteria for evaluating the models of bounded rationality in macroeconomics addresses this criticism. Thus, it represents an important part of research in this area. Overall, the evolutionary model of the exchange-rate behavior meets five criteria important for the evaluation of models of bounded rationality. First, it captures the persistence in the exchange-rate volatility observed in actual exchange-rate time series, but not explained by rational expectations models of the exchange rate. Second, it does well when compared to the evidence from the experiments with human subjects. Third, the model is robust to changes in the representation of decision rules and the way in which the evolutionary adaptation is implemented. Finally, given that performance of boundedly rational agents is comparable to the performance of rational agents, it does well when evaluated using this criterion as well.

NOTES

1. A number of recent papers have demonstrated that the introduction of learning behavior may generate endogenous fluctuations: e.g., Bullard (1994), Timmerman (1996), Arthur et al. (1997), Brock and Hommes (1998), and Hommes and Sorger (1998).

2. The argument that results of models with boundedly rational agents are highly sensitive to changes in parameter values has often been used as a criticism of this type of research.

3. In each period, a rational agent who lives for two periods is born and her utility is evaluated at the end of the second period of life.

4. This is an environment with perfect currency substitution. It is convenient for studies of the issues related to international capital flows. The way to think about these agents is that they are traders in the international capital markets. The model can be extended to include three types of agents, where type 1 is restricted to using currency 1, type 2 is restricted to using currency 2, and type 3 can hold either currency. It is these type-3 agents on which we concentrate.

5. See Michalewicz (1996) for details on how to use crossover with real numbers. The crossover was not implemented for the reasons of analytical tractability; i.e., its implementation would have prevented the derivation of the forecasts that are used by the rational agent.

6. At the begining of $t+1$, the values of endogenous variables are not yet known.

7. This just implies that a random number that determines whether or not experimentation is performed will be drawn. Thus, it does not guarantee that the experimentation will actually take place.

8. Superscript s denotes the parts of the expected value of $s^{p,e}(t+1)$.

9. Superscript λ denotes the parts of $\lambda^{p,e}(t+1)$.

10. Superscript 1 refers to case 1 of $R_1(t) > R_2(t)$.

11. Superscript 2 denotes case 2 of $R_1(t) < R_2(t)$.

REFERENCES

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

Arthur, B., J. Holland, B. LeBaron, R. Palmer & P. Tayler (1997) Asset pricing under endogenous expectations in an artificial stock market. In B. Arthur, S. Durlauf & D. Lane (eds.), *The Economy* *as an Evolving Complex System*, Vol. II, pp. 15–43. SFI Studies in the Sciences of Complexity, Vol. 27. Menlo Park, CA: Addison-Wesley.

Brock, W.A. & C.H. Hommes (1998) Rational route to randomness. *Econometrica* 65, 1059–1095.

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

Bullard, J. & J. Duffy (2001) Learning and excess volatility. *Macroeconomic Dynamics* 5, 272–302.

Frankel, G. & A. Rose (1995) Empirical research on nominal exchange rates. In G. Grossman & K. Rogoff (eds.), *Handbook of International Economics*, Vol. 3, pp. 1689–1729. Elsevier Science.

Hommes, C. & G. Sorger (1998) Consistent expectations equilibira. *Macroeconomic Dynamics 2*, 287–320.

Kareken, J. & N. Wallace (1981) On the indeterminacy of equilibrium exchange rates. *Quarterly Journal of Economics* 96, 207–222.

Marcet, A. & J.P. Nicollini (1997) Recurrent Hyperinflations and Learning. Economics working paper 244, Universitat Pompeu Fabra.

Meese, R. & K. Rogoff (1983a) Empirical exchange rate models of the seventies. *Journals of International Economics* 14, 3–24.

Meese, R. & K. Rogoff (1983b) The out-of-sample failure of empirical exchange rate models. In J. Frenkel (ed.), *Exchange Rates and International Macroeconomics*. Chicago: University of Chicago Press.

Michalewicz, Z. (1996) *Genetic Algorithms + Data Structures = Evolution Programs*, 3rd ed. New York: Springer-Verlag.

Timmerman, A.G. (1996) Excess volatility and predictabiity of stock prices in autoregeressive divident models with learning. *Review of Economics Studies* 63, 523–557.