Performance of rational and boundedly rational agents in a model with persistent exchange rate volatility

Jasmina Arifovic

Simon Fraser University
and
California Institute of Technology

ABSTRACT

The model studied in this paper is a two-country overlapping generations model with boundedly rational agents who update their decision rules using a version of the stochastic replicator dynamic. The results presented in the paper 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 a rational agent who has a full knowledge of the evolutionary economy. The performance of this agent is compared to the performance of boundedly rational agents based on the average utility that each type of agent receives over time. Results show that overall the difference between utilities earned by rational and boundedly rational agents is small. In addition, best performing boundedly rational agents perform better than the rational agent.

JEL classification: D84; E42; F31

Keywords: bounded rationality; exchange rate fluctuations;

** 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.

1. Introduction

The persistent fluctuations of exchange rates have proven to be an empirical phenomenon which is difficult to explain theoretically. The study by Meese and Rogoff (1983a, 1983b), for instance, showed that the structural models aimed at explaining the behavior of the exchange rate in terms of the fundamentals failed to improve on the random walk out-of-sample forecasting accuracy. It seems that no model based on fundamentals like money supplies, real income, interest rates, inflation rates, and current account balances has been able to explain or predict a high percentage of the variation in the exchange rate, at least at short-or medium-term frequencies. Non-structural 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.)

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 showed 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 in Arifovic (1996). This is not surprising since 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

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

differ in the specific implementation of these principles such as the representation of beliefs and the types of operators that are used for updating.

The stochastic replicator dynamic used in this paper allows a degree of analytical tractability that was not possible to obtain with a genetic algorithm model. Mainly, a two-period ahead conditional forecast of the rates of return can be derived. Generally, applications of evolutionary algorithms in macroeconomic models are simulation-based because the environment is usually too complex to be characterized analytically. The possibility to derive a two-period ahead forecast is exploited in order to add to the model a rational agent who uses the forecast to make optimal savings and portfolio decisions. The performance of this rational agent is compared to the performance f boundedly rational agents.

A number of recent papers that examine the statistical learning models have compared the performance of boundedly rational and rational agents in terms of their forecasting errors (Marcet and Nicolini, 1999, Bullard and Duffy, 2000, and Hommes and Sorger, 1998). This research tries 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 also compares the performance of the rational and boundedly rational agents. However, given that this is a model of bounded rationality with heterogenous beliefs where agents do not make forecasts of the next period's values of prices, the same methodology which is based on the comparison of the forecasting errors cannot be used in this set up to assess the performance of boundedly rational agents. Instead, the comparison is based on the utilities of the two types of agents earned over the long horizons. As in the other above cited models, decisions of the rational agents do not affect the determination of prices.

The performance is evaluated based on the utility that each of the two types of agents earn. The results show that the best within-a-generation boundedly rational agent performs better than the rational agent. In addition, while the rational agent performs better that the average and the median boundedly rational agent, the differences in performance, measured in terms of average utility, is small.

The paper proceeds as follows. The description of the economic model is given in section 2. The updating of agents' rules is described in section 3. The dynamics of adaptation are examined in section 4. The derivation of the rational agent's forecasting rule is provided in section 5. This 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 the 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 freely to borrow 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:

$$\max\ \ln\ c_t(t) + \ln\ c_t(t+1)$$
$$s.t.\ c_t(t) \leq w^1 - \frac{m_1(t)}{p_1(t)} - \frac{m_2(t)}{p_2(t)}$$
$$c_t(t+1) \leq w^2 + \frac{m_1(t)}{p_1(t+1)} + \frac{m_2(t)}{p_2(t+1)}$$

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. 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, $m_2(t)/p_2(t)$.

The exchange rate $e(t)$ between the two currencies is given by $e(t) = p_1(t)/p_2(t)$. Since 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 have to be equal to:

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

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

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

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

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

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, i.e. that

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

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 indeterminacy of exchange rate 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

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

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:

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

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 denoted by $\lambda_i(t) \in [0,1]$. The portfolio fraction defines the fraction of 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 $p_2(t)$:

$$p_1(t) = H_1 / \sum_i^N \lambda_i(t)s_i(t) \quad p_2(t) = H_2 / \sum_i^N (1-\lambda_i(t))s_i(t). \tag{9}$$

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:

$$m_{i,1}(t) = \lambda_i(t)s_i(t)p_1(t) \qquad m_{i,2}(t) = (1-\lambda_i(t))s_i(t)p_2(t). \tag{10}$$

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:

$$c_{i,t}(t+1) = w^2 + \frac{m_{i,1}(t)}{p_1(t+1)} + \frac{m_{i,2}(t)}{p_2(t+1)} \tag{11}$$

At the end of period $t+1$, agents utilities are computed based on 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:

$$\mu_{i,t} = \ln\ c_{i,t}(t) + \ln\ c_{i,t}(t+1). \tag{12}$$

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 occasional experimentation 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.

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:

$$Pr_i(t+2) = \frac{\mu_{i,t}}{\sum_{i=1}^{N} \mu_{i,t}}. \tag{13}$$

Rules of generation t are lined up on the interval $[0,1]$. Each rule occupies a slot that is equal to its relative fitness. 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 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. Denote this copy by $C_j(t+2)$.

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

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 probabilities 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 $E_j(t+2)$.

Election operator

Prior to final determination of the population of rules of generation $t+1$, 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 or equal, the new rule $E_j(t+2)$ replaces a 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 the 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 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. In a stationary equilibrium, all agents make the same savings and portfolio fraction decisions. The nominal 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 take place. Imitation has no effect as all the rules are identical. Experimentation brings in new rules, different from the equilibrium ones into the population, 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 were 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 since the rates of return were 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 evolutionary stable with respect to invading new rules $E_j(t+1)$s 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 fractio). However, it is not evolutionary stable with respect to invading rules $E_j(t+1)$s 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 3 values of the rate of experimentation. For each π_{ex}, 5 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 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.

[Table 1 around here]

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 stochastic replicator dynamic Presented in this paper and the genetic algorithm dynamics 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 of experimentation. However, crossover, the operator that performs recombination of the parts of existing binary strings has no counter part in the current set up. Crossover which plays important role in the genetic algorithm adaptation could be implemented in this framework as well.[2] However, it has not been implemented in the framework described in this paper for reasons of analytical tractability, i.e. the use would have

[2] See Michalewicz (1996) for details on how to use crossover with real numbers.

prevented the derivation of the forecasting rule of a rational agent presented in the next section.

This paper shows that the main dynamics are preserved under different type of rules' representation and application of different operators. We are not aware of any other study that examines the robustness of the results of evolutionary adaptation in economic environments under different rules' representation and different updating schemes.

5. A 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 decision, and at the end of the old age, her utility is evaluated. Thus in every t, there are 2 rational agents, young and old. 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 types 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 the given generation. Finally, the third is the agent that received the highest utility in a given generation. This analysis is done in order 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 than the rational agent, then it can be argued that the model can be used as a reasonable model of agents' adaptation. In addition, the best agent of a generation should perform better than the rational agent. Thus, a possibility that a rule can do better than the rational rule provides additional incentive for agents to follow the evolutionary rules in their decision making. On average, they might not do much worse than the rational agent and, in addition, there is a chance of actually doing better. Of course, there are always agents who perform below the average.

Next, we describe the derivation of a two-period ahead forecast of the rational agent. She will use this forecast to make savings and portfolio decisions. In order to simplify things, the analysis will be carried out for the endowment pattern where $w^1 > 0$ and $w^2 = 0$. For this endowment pattern and the above defined preferences, the optimal savings decision, $\hat{s}(t)$, is always equal to $w^1/2$ regardless of the rate of return. Thus the only issue that remains to be resolved is what fraction of these savings to place in each currency. This, of course, will depend on the expected values of the rates of return on the two currencies at the end of period $t+1$. These in turn depend on the expected values of portfolio fractions at $t+1$ and $t+2$.

Since the optimal savings decision is independent of the rate of return, it is also independent of the fractions placed in each currency. This allows for separating the analysis of the evolution of savings decision and the evolution of portfolio decisions. A rational agent has to make a savings and a portfolio decision based on her knowledge of the distribution of decision rules of populations of generation $t-1$ and t. She uses this information and the knowledge of the updating processes to make his/her savings and portfolio decision. These decisions depend on the rational agent's forecasts of the average savings decision at $t+1$ and on the average portfolio fractions at $t+1$ and $t+2$. The rational agent's utility is evaluated at the end of time $t+2$.

5.1. The Forecasting Rule

We first derive the expected value of savings at $t+1$. The expected value of savings that is the outcome of the process of imitation at $t+1$, $s^{m,e}(t+1)$, is given by:

$$s^{m,e}(t+1) = \sum_{i=1}^{N} Pr_i(t+1) s_i(t-1). \tag{14}$$

Next, we have to consider the impact of experimentation and of the election operator in order to derive the expression for the expected value of savings at $t+1$, $s^e(t+1)$. The expected value of savings decisions that undergo experimentation, $s^{p,e}(t+1)$, is derived in the following way. First, if $s^{m,e}(t+1) > \hat{s}(t)$, this value is given by:

$$s^{p,e}(t+1) = \pi_{ex} \frac{2|\hat{s}(t) - s^{m,e}(t+1)|}{w^1} \frac{p_1(t+1) + p_2(t+1)}{2} \tag{15}$$

$$+\pi_{ex}\Big[\frac{(w^1 - s^{m,e}(t+1))}{w^1} + \frac{p_2(t+1)}{w^1}\Big]\bar{s}^{m,e}(t+1)$$
$$+(1-\pi_{ex})s^{m,e}(t+1)$$

where

$$p_1(t+1) = \hat{s} + |s^{m,e}(t+1) - \hat{s}|$$

$$p_2(t+1) = \hat{s} - |s^{m,e}(t+1) - \hat{s}|.$$

the first term of the above expression gives the contribution to the expected value of the result of successful experimentation, i.e. experimentation that passes the election operator test. The term $2|\hat{s}(t) - s^{m,e}(t+1)|/w^1$ gives the probability that acceptable values will be chosen and the term $(p_1(t+1) + p_2(t+1))/2$ gives the expected value of the admissible range.

The second term is contribution due to experimenting with values of savings that do not pass the election operator test. Thus, the term in squared brackets represents probability that values in that inadmissible range are selected and then multiplied by $s^{m,e}(t+1)$ because that is the expected value of the old rule that remains in the population.

Finally, the third term is the contribution due to the value of those savings rules that do not undergo experimentation.

Second, if $s^{m,e}(t+1) < \hat{s}(t)$, the value of $s^{p,e}(t+1)$, is given by:

$$s^{p,e}(t+1) = \pi_{ex}\frac{2|\hat{s}(t) - s^{m,e}(t+1)|}{w^1}\frac{p_1(t+1) + p_2(t+1)}{2} \tag{16}$$
$$+\pi_{ex}\Big[\frac{s^{m,e}(t+1)}{w^1} + \frac{w^1 - p_1(t+1)}{w^1}\Big]\bar{s}^{m,e}(t+1)$$
$$+(1-\pi_{ex})s^{m,e}(t+1).$$

The first, second and third term have similar interpretation as in equation (15). Finally, since savings part is selected for experimentation with probability 1/2, the expected value of savings $s^e(t+1)$ is the weighted average of $s^{m,e}(t+1)$ and $s^{p,e}(t+1)$:

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

Next, we derive the expected value of $\lambda(t+2)$. We first have to determine the expected value of the portfolio fraction that is the outcome of the process of imitation, $\bar{\lambda}^{m,e}(t+1)$. It is given by:

$$\bar{\lambda}^{m,e}(t+1) = \sum_{i=1}^{N} Pr_i(t+1)\lambda_i(t-1) \tag{17}$$

If $R_1(t-1) > R_2(t-1)$, $\lambda^{p,e}(t+1)$, the the expected value of the result of experimentation is given by:

$$\bar{\lambda}^{p,e}(t+1) = \pi_{ex}(1-\bar{\lambda}^{m,e}(t+1))\frac{(1+\bar{\lambda}^{m,e}(t+1))}{2} + \tag{18}$$
$$\pi_{ex}\bar{\lambda}^{m,e}(t+1)\bar{\lambda}^{m,e}(t+1) + (1-\pi_{ex})\bar{\lambda}^{m,e}(t+1).$$

Note that in this case, the election operator will on average accept only those new rules generated through experimentation that prescribe fractions greater than $\bar{\lambda}^{m,e}(t+1)$. The probability of this event is given by the first term in the above expression. The first part of the second term is related to the part of the expected value contributed by the rules that undergo experimentation, but that are inadmissible according to the election operator (these will be λs smaller than $\lambda^{m,e}(t+2)$). The second part of the second term is the expected value contributed by the rules that do not undergo experimentation. If $R_1(t-1) < R_2(t-1)$, $\lambda^{p,e}(t+1)$, the the expected value of the result of experimentation is given by:

$$\lambda^{p,e}(t+1) = \pi_{ex}\lambda^{m,e}(t+1)\frac{\lambda^{m,e}(t+1)}{2} + \tag{19}$$
$$\pi_{ex}(1-\lambda^{m,e}(t+1))\lambda^{m,e}(t+1) + (1-\pi_{ex})\lambda^{m,e}(t+1) \tag{20}$$

.

In this case, the election operator will admit those newly generated rules that are smaller, on average, than $\lambda^{m,e}(t+1)$. Thus the first part of the expression gives the amount contributed to the expected value by experimentation. The second part is again the amount due to the rules that undergo experimentation but do not pass the election operator test, and those rules that do not undergo experimentation.

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 $\lambda^{p,e}(t+1)$:

$$\lambda^{e}(t+1) = \frac{1}{2}\lambda^{m,e}(t+1) + \frac{1}{2}\lambda^{p,e}(t+1). \tag{21}$$

The second stage involves the formation of expectations of $\lambda^{e}(t+2)$. In order to do this, we first have to derive the expression for the expected values of the rates of return on the two currencies at the end of $t+1$ as these values are used in the derivation of the expected outcomes of imitation and experimentation at $t+2$. These values are given by:

$$R_1^e(t+1) = \frac{(1/2)(s^{m,e}(t+1)\lambda^{p,e}(t+1) + s^{p,e}(t+1)\lambda^{m,e}(t+1))}{\bar{s}_1(t)} \tag{22}$$

$$R_2^e(t+1) = \frac{(1/2)[s^{m,e}(t+1)(1-\lambda^{p,e}(t+1)) + s^{p,e}(t+1)(1-\lambda^{m,e}(t+1))]}{\bar{s}_2(t)} \tag{23}$$

where $\bar{s}_1(t)$ is the average savings placed in currency 1 at time t and $\bar{s}_2(t)$ is the average savings placed in currency 2 at time t. These values of expected rates of return are used to determine expected fitness values of rules of generation t. 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:

$$\lambda^{p,e}(t+2) = \sum_{i=1}^{N} Pr_i(t+2)\lambda_i(t)) \tag{24}$$

If $R_1^e(t) > R_2^e(t)$, implying that $\lambda^e(t+1) > \bar{\lambda}(t)$, the expectation of the average portfolio fraction at $t+2$ will be given by:

$$\lambda^e(t+2) = \pi_{ex}(1-\lambda^{p,e}(t+2))\frac{(1+\lambda^{p,e}(t+2))}{2} + \tag{25}$$
$$\pi_{ex}\lambda^{p,e}(t+2)\lambda^{p,e}(t+2) + (1-\pi_{ex})\lambda^{p,e}(t+2)$$

.

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

$$\lambda^e(t+2) = \pi_{ex}\lambda^{p,e}(t+2)\frac{\lambda^{p,e}(t+2)}{2} + \tag{26}$$
$$\pi_{ex}(1-\bar{\lambda}^{p,e}(t+2))\lambda^{p,e}(t+2) + (1-\pi_{ex})\lambda^{p,e}(t+2)$$

.

Given this forecast of $\lambda(t+2)$, the rational agent makes a decision at the beginning of time t 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.

5.2. Performance of rational and boundedly rational agents

This section compares the performance of the boundedly rational agents and rational agents. The discussion focuses on the environment where only decisions of boundedly rational agents affect the price levels. The performance will be 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, and to the highest utility in the population of boundedly rational agents.

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 best performing boundedly rational agent. Let $\bar{d}^a$ be the average over T_{max} periods. Similarly, let $\bar{d}_t^m$ 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 a 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 $t+1$, i.e. the agent that received the highest utility in the entire generation. Let $\bar{d}^b$ be the average over T_{max} periods.

The calculations of the above described measures were done for simulations with 3 different rates of experimentation, 0.0033, 0.033, and 0.33. Reported results are based on 5 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. Rates of mutation that are too high result in overall worse performance of the boundedly rational agents. At the same time the best within a generation boundedly rational agent always does better than the rational agent. The difference is the same for the rates of experimentation of 0.033 and 0.33. It is smaller for the really low rate of mutation of 0.003. These results are presented in table 2.

[Table 2 around here]

The overall 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 5 simulations. Clearly, the ranking of the rational agent worsens with the increases in the rate of experimentation.

[Table 3 around here]

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 simple evolutionary algorithm based on imitation and experimentation. While the savings decisions settle to the values in the neighborhood of the steady state, the portfolio fractions values do not settle towards the constant, steady state values. The fluctuations of portfolio fraction values results in the persistence in the exchange rate fluctuations.

In the second part of the paper, a rational agent is added to the economy. This agent makes the same types of decisions as boundedly rational agents and receives utility based on these decisions. A rational agent makes savings and portfolio decisions using two-period ahead forecasts of the rates of return on two currencies. The comparison of the long-run performance of the rational and boundedly rational agents show little difference in terms of the average utilities earned over time. Moreover, as populations of boundedly rational agents remain heterogenous in terms of their portfolio decisions, there are always boundedly rational agents who receive utilities that are higher than those received by rational agents.

References

Arthur, B, J. Holland, B. LeBaron, R. Palmer, and P. Tayler (1997) Asset Pricing Under Endogenous Expectations in an Artificial Stock Market. In B. Arthur, S. Durlauf, and D. Lane (eds.), *The Economy as an Evolving Complex System II*, pp. 15-43. SFI Studies in the Sciences of Complexity, Vol. XXVII, Addison-Wesley.

Arifovic, J. and R. Gencay (in press) Statistical Properties of Genetic Learning in a Model of Exchange Rate. *Journal of Economic Dynamics and Control.*

Arifovic, J. (1996) The Behavior of Exchange Rate in the Genetic Algorithm and Experimental Economies, *Journal of Political Economy* 104, 510-541.

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

Bullard, J. and J. Duffy (1999) Learning and Excess Volatility, manuscript.

Frankel and Rose (1995) Empirical Research on Nominal Exchange Rates, in G. Grossman and K. Rogoff (eds.), *Handbook of International economics*, vol. III, pp. 1689-1729. Elsevier Science, B.V.

Hommes, C. and G. Sorger (1998) Consistent Expectations Equilibira, *Macroeconomic Dynamics* 2, 287-320.

Kareken, J., and N. Wallace (1981) On the Indeterminacy of Equilibrium Exchange Rates. *Quarterly Journal of Economics* 96, 207-222.

Marcet, A. and J.P. Nicollini (1997) Recurrent Hyperinflations and Learning, Economics Working Paper 244, Universitat Pompeu Fabra.

Meese, R. and K. Rogoff (1983a) Empirical Exchange Rate Models of the Seventies. *Journal of International Economics* 14, 3-24.

Meese, R. and 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 edition. New York: Springer Verlag.

Timmerman, A.G. (1996) Excess Volatility and Predictabiity of Stock Prices in Autoregeressive Dividend Models with Learning. *Review of Economic Studies* 63, 523-557.

Table 1

Frequency distribution of portfolio fraction values

π_{ex}	0.0033	0.033	0.33
λ values			
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

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

Table 3

Ranking of rational agent

π_{ex}	0.0033	0.033	0.33
percentile rank %			
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

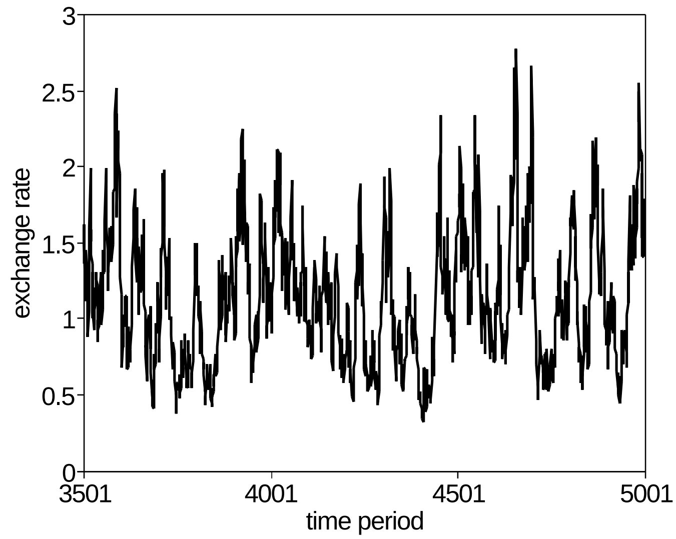

Figure 1 - Exchange Rate

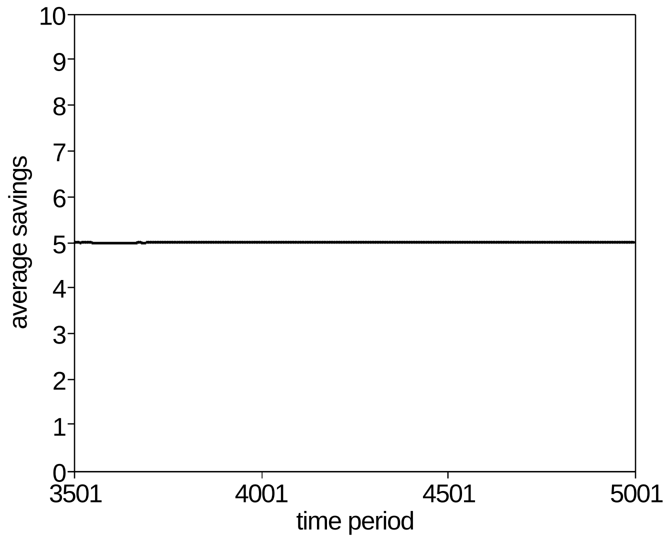

Figure 2 - Average Savings