Exchange Rate Volatility in the Artificial Foreign Exchange Market

Jasmina Arifovic

Simon Fraser University, Burnaby, B.C., Canada and
California Institute of Technology 228-77 Pasadena, CA 91125, U.S.

Abstract. This paper studies co-evolution of different decision rules in an artificial foreign exchange market. The behavior of the exchange rate depends on the type of decision rules that agents use. Evolution of the moving average and least squares forecasting techniques results in a speculative attack on one of the currencies and that currency's eventual collapse. Addition of the rules that evolve the portfolio fractions directly brings in persistent volatility of the exchange rate that resembles the actual exchange rates time series.

1 Introduction

Persistent fluctuations have characterized the behavior of the exchange rates ever since the flexible exchange rate system was introduced. Theoretical models, based on fundamentals like money supplies, real income, interest rates, inflation rates, and current account balances, have not been successful in capturing a high percentage of the variation in the exchange rate at short-or-medium-term frequencies. An alternative way to try to model the behavior exhibited under the flexible exchange rates system is to address explicitly the issue of agents' beliefs and the way in which they change over time.

This paper examines the exchange rate behavior in an artificial foreign exchange market (AFEM). The paper studies co-evolution of different decision rules in a simple general equilibrium monetary model where there are no changes over time in terms of fundamentals. [1] Agents can use different rules in making their portfolio decisions, i.e. how much of their savings to allocate to each currency. Portfolio decisions in turn affect the level of prices. Thus nominal prices, rates of return and exchange rates are endogenously determined. Their dynamics are influenced solely by changes in agents' portfolio decisions.

As a specific application of the AFEM framework, the results of simulations where moving average forecasting techniques evolve are presented. Binary strings encode the size of the sample used in the computation of a forecast and indicate which of the availabe techniques will be used. Rules

[1] In the rational expectations equilibrium, the exchange rate is constant and, due to the perfect substitutability of the two currencies that are traded, the equilibrium exchange rate is also indeterminate.

are updated using the genetic algorithm. These simulations result in the collapse of one of the two currencies and the convergence of the economy to a single-currency equilibrium. The addition of least squares to the population of evolving forecasting techniques brings about even faster collapse. These speculative attacks imply greater volatility of the exchange rate, a feature that also characterizes actual exchange rate time series. However, a complete collapse of one of the currencies that is due purely to shifts in agents' beliefs is not necessarily final outcome of speculative attacks that occur in the actual foreign exchange markets.

The addition of populations of binary strings that encode portfolio fractions generates volatility of the exchange rate that persists over time. Speculative attacks on both currencies occur, but they come to an end before the investment into the currency that is under the attack can be driven down to zero. Thus, the volatility of the exchange rate behavior observed in this AFEM environment captures some of the behavior observed in the actual foreign exchange markets.

It is worthwhile pointing out that, unlike the other artificial stock market models that require a sequence of exogenous shocks to one of the fundamentals (usually a dividend is assumed to follow a stochastic, AR1 process), in order to generate some persistence in the behavior, no such shocks are required in the artificial foreign exchange market described in this paper. Indeterminacy, together with the evolution of beliefs, results in the persistence in volatility of the exchange rate.

Section 2 describes the economic environment under consideration. Section 3 describes the AFEM framework and the results of simulations of two applications. Section 4 concludes and outlines some of the possible extensions of the basic model.

2 Description of the model

The *world* economy consists of two countries. It is an overlapping generations model in which the residents of both countries are identical in terms of their preferences and lifetime endowments. The economy starts at $t = 1$ and lasts for ever. Individual agents live for two periods, and at each $t \geq 1$, $\frac{N}{2}$ new young individuals are born in each country, said to be of generation t. They are young at period t and old at period $t+1$. Each young agent of generation t is endowed with n units of labor and their own production technology for producing a single consumption good with a one-to-one labor/output transformation. When old, agents do not receive any labor endowment. No storage technology is available in the economy. Agents in both countries have the common preferences given by: $u_t[c_t^t, c_{t+1}^t] = c_t^t c_{t+1}^t$, where c_t^t is agent's consumption when young, and c_{t+1}^t is agent's consumption when old.

A government of each country issues its own unbacked currency. Supplies of both currencies, $M_{1,t}$ and $M_{2,t}$ are kept constant so that $M_{1,t} = M_1$ and

$M_{2,t} = M_2$. There are no legal restrictions on holdings of foreign currency. Thus the residents of both countries can freely hold both currencies in their portfolios. An agent of generation t solves the following maximization problem at time t:

$$\begin{aligned} \max \quad & c_t^t \ c_{t+1}^t \\ s.t.\ \ c_t^t \leq n - & \frac{m_{1,t}}{p_{1,t}} - \frac{m_{2,t}}{p_{2,t}} \\ c_{t+1}^t \leq & \frac{m_{1,t}}{p_{1,t+1}} + \frac{m_{2,t}}{p_{2,t+1}} \end{aligned}$$

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 defined as $e_t = p_{1,t}/p_{2,t}$. When there is no uncertainty, the return on the two currencies must be equal,

$$R_{1,t} = R_{2,t} = \frac{p_{1,t}}{p_{1,t+1}} = \frac{p_{2,t}}{p_{2,t+1}}, \quad t \geq 1, \tag{1}$$

where $R_{1,t}$ and $R_{1.t}$ are 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:

$$e_{t+1} = e_t = e, \quad t \geq 1 \tag{3}$$

Savings demand derived from agent's maximization problem is given by

$$s_t = \frac{m_{1,t}}{p_{1,t}} + \frac{m_{2,t}}{p_{2,t}} = \frac{n}{2}. \tag{4}$$

Aggregate savings that represent real world money demand are equal to the sum of young agents' savings, i.e. $S_t = Ns_t$. Since the rates of return on the two currencies are identical, the agents are actually indifferent as to which currency they hold. Because of this, equations for individual money demands are not well defined. (There is only one equation for the world real demand.) This fact results in the indeterminacy of the exchange rate.

The indeterminacy of the exchange rate proposition (Kareken and Wallace, 1981) asserts that if there is a monetary equilibrium where savings demand and money supplies are equal for an exchange rate, e, then there exists an equilibrium for any exchange rate $\hat{e} \in (0, \infty)$, $\hat{e} \neq e$. If there is an equilibrium for a price sequence, $\{p_{1,t}, p_{2,t}\}$, for the exchange rate e, we can find a sequence $\{\hat{p}_{1,t}, \hat{p}_{2,t}\}$, for the exchange rate $\hat{e}$ that results in the same sequence of real rates of return as the original price sequence and in turn in the same values of aggregate savings. The reason why this can be accomplished is the equivalence between the two currencies as savings instruments.

3 Description of the artificial foreign exchange market

In this section, we develop the elements of the AFEM. Agents that participate in the market make savings and portfolio decisions. They save half of the amount of the good that they produce in the first period. Given their labor endowment pattern, that is the optimal consumption decision. [2] However, they also have to make their portfolio decision, i.e. how much of their savings to invest in each currency. Let λ_t^i, $\lambda_t^i \in [0,1]$, denote a portfolio fraction decision of agent i of generation t.

In the rational expectations equilibrium, $\lambda_t^i = \lambda$ for all i, and all t. The level of λ determines the (constant) value of the exchange rate. Note that the equilibrium value of λ cannot be deduced from the rational expectations version of the model. However, within the AFEM framework, agents have heterogenous beliefs implying heterogenous values of the portfolio fractions. With heterogenous values of λ, and an assumption that $M_1 = M_2 = M$, the expressions for $p_{1,t}$ and $p_{2,t}$ are given by

$$p_1(t) = 2M / \sum_i^N \lambda_t^i n \quad p_2(t) = 2M / \sum_i^N (1 - \lambda_t^i) n. \tag{5}$$

or

$$p_{1,t} = 2M / \bar{\lambda}_t n \quad p_{2,t} = 2M / (1 - \bar{\lambda}_t) n \tag{6}$$

where $\bar{\lambda}_t$ is the average portfolio fraction. The exchange rate is then given by:

$$e_t = \frac{1 - \bar{\lambda}_t}{\bar{\lambda}_t}. \tag{7}$$

However, at the time when AFEM agents make their portfolio decisions, the value of e_t is not known. The agents can choose among different rules in deciding on their portfolio fraction. The choice of a rule is influenced by rule's past performance and by occasional experimentation with new rules.

[2] In the case that agents receive no labor endowment in the second period of their life, the savings decision does not depend on the rates of return on savings.

In the first environment discussed in the paper, agents are endowed with different types of moving-average rules that they use to compute exchange rate forecasts. Let $e_t^{f,i}$ be an exchange rate forecast of agent i of generation t. Then, agent i sets the value of λ_t^i equal to:

$$\lambda_t^i = \frac{1}{1+e_t^{f,i}}. \tag{8}$$

Agents do not use the entire history of the exchange rates. They discard old information and employ a rolling sample. The size of the sample, T, (an even number) differs across agents and evolves over time. Agents can also choose between two types of forecasting procedures, f^1 and f^2. If f^1 is used, every sample observation is included in the computation of the moving average, and the exchange rate forecast is given by:

$$e_t^f = \frac{\sum_{k=1}^{T} e_{t-k}}{T} \tag{9}$$

On the other hand, if f^2 is used, only every second observation is considered in the computation of the forecast of the exchange rate, i.e.:

$$e_t^f = \frac{\sum_{k=0}^{T} e_{t-k-2}}{T/2} \tag{10}$$

The number of f^1 and f^2 rules also evolves over time. The actual exchange rate depends on the way individual forecasts are made and, using (8) and (7), is given by:

$$e_t = \frac{\sum_{i=1}^{n}\left(1-\frac{1}{1+e_t^{f,i}}\right)}{\sum_{i=1}^{n}\frac{1}{1+e^{f,i}}}. \tag{11}$$

At each time period t, there are two populations of forecasting rules, one that represents the rules of the young agents (generation t) and the other that represents the rules of the old agents (generation $t-1$).[3] Only the rules of the young agents play an active role at time t. Each young agent is endowed with a binary string, of length ℓ, that has the following interpretation. The first bit of a binary string indicates whether f^1 or f^2 will be used. The bits $[2\cdots\ell]$ encode the sample size, $T \in [1,\ldots,64]$, i.e. the number of past observations that will be taken into account when computing the moving average of past values of the exchange rate.

[3] See Arifovic (1995, 1996) for a detailed description of the implementation of the genetic algorithm in the overlapping generations environments.

The economy is initialized at the point where forecasting techniques are randomly distributed. An initial set consisting of $T_{max} = 64$ exchange rate observations is generated in the following way. For each observation, a random number between 0 and 1 is drawn from the uniform distribution. This number is interpreted as an average portfolio fraction and is used to compute the exchange rate value.

At each time period t, forecasting rules are decoded and individual forecasts are computed. Then individual portfolio fractions are calculated using these forecasts. Portfolio fractions determine the savings in terms of currency 1 and currency 2. Finally, nominal prices, exchange rate and rates of return in terms of each currency are calculated. Once the rates of return are known, second period consumption values are computed for members of generation $t-1$ and the fitness values for the forecasting rules of generation $t-1$ are calculated. A fitness of string i is given as the utility of agent i of generation $t-1$. The population of forecasting rules of generation $t-1$ is then used to obtain a population of forecasting rules for generation $t+1$.

A population of forecasting rules evolves using the genetic algorithm. Tournament selection is used as the reproduction operator. The one-point crossover takes place with probability 0.6. The probability of mutation is set to 0.033. In addition to these standard genetic operators, the election operator (Arifovic, 1994, 1996) is applied as a local elitist procedure.

A *weak* form of the election operator (Franke, 1998) is used in the following way. After the application of the crossover operator on a pair of binary strings takes place, these two binary strings are recorded as parent 1 and parent 2. The resulting offspring strings are recorded as offspring 1 and offspring 2. Once the two offspring undergo mutation, their fitness values are calculated using the last period's rates of return. Then, the fitness of the first offspring is compared to the fitness of parent 1. If it is higher than the parent's fitness, the offspring enters as a member of the new population. However, if the parent's fitness is higher than the offspring's, the parent remains as the member of the new population. Likewise, if the fitness of offspring 2 is higher than or equal to the fitness of parent 2, the offspring 2 enters into the new population. Otherwise, parent 2 becomes a member of the new population.

Simulations of the above described evolutionary process resulted in the convergence of the economies to a single-currency equilibrium. Which of the two currencies is selected depends on a particular sequence of pseudo random numbers. Initially, the rates of return on two currencies fluctuate. There are time intervals during which $R_{1,t}$ is greater than $R_{2,t}$, and those when the direction of inequality changes sign and $R_{2,t} > R_{1,t}$. Eventually, one of the rates of return remains greater than the other long enough that it initiates a steady increase of the holdings of the currency with the higher rate of return. The final result is that agents place all of their savings in the *higher-return* currency.

The populations of forecasting tehniques remain heterogenous. Both f^1 and f^2 moving averages are represented in the populations and binary strings decode to sample lengths of different sizes. However, all of the forecasts result in the same value of λ at the end of the simulation. Figure 1 illustrates the behavior of the average portfolio fraction, and figure 2 the behavior of the rate of return on currency 1 in one of the simulations.

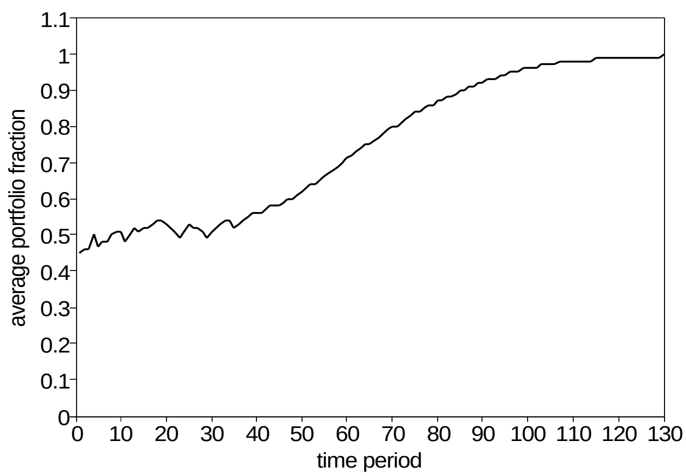

Fig. 1. Average portfolio fraction of f^1 and f^2 rules

After initial fluctuations, $\bar{\lambda}_t$ starts a steady increase towards the value of 1, indicating a speculative attack on currency 2. Once $\bar{\lambda}_t$ reaches the value of 1, currency 2 collapses. Examination of figure 2 reveals that after the initial fluctuations above and below 1, starting with $t = 46$, $R_{1,t}$ takes only values greater than 1. This is the interval during which $R_{1,t} > R_{2,t}$, and exactly the time when $\bar{\lambda}_t$ begins its steady increase. The end of the simulation is characterized by a slow decline of $R_{1,t}$ towards the value of 1 that is its value in the single-currency stationary equilibrium value.

Simulations of the economies in which only the size of the rolling sample evolved (and all agents used either f^1 or f^2) resulted in the same outcomes, i.e. the convergence to a single currency equilibrium. The addition of the least squares to the pool of forecasting techniques sped up the process of convergence to a single currency equilibrium. [4]

[4] Sargent (1993) applies the stochastic approximation algorithm to a version of this economy described in section 2. His results show convergence to an equilibrium

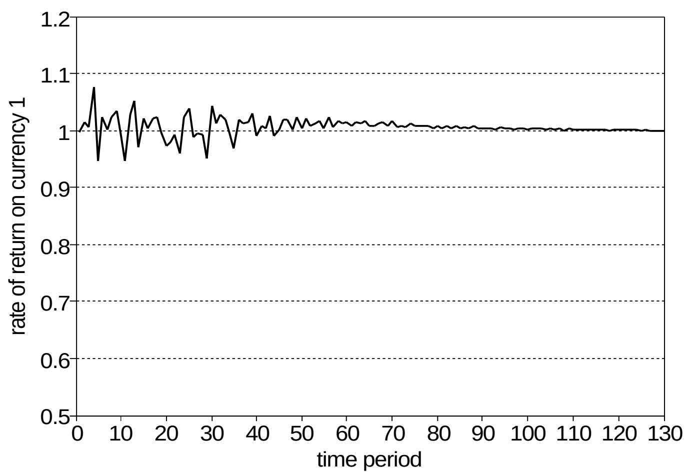

Fig. 2. Rate of return of f^1 and f^2 rules

While this result is interesting in light of the fact that the evolution of beliefs can result in speculative attacks, these types of speculative attacks are not observed in the actual time series. Even though the speculative attacks can occur without any apparent change in fundamentals, they end at some relatively high, but finite value of the exchange rate.

Next, we introduce another class of rules that will be represented by two overlapping populations of binary strings. With this class, a binary string encodes the value of λ_t^i. [5] Two new populations of binary strings are added to the AFEM in order to emulate the model's overlapping generations structure. These two populations that encode the values of λ represent a separate pool of rules that undergo genetic algorithm updating.

Let us denote the first class of rules that consists of f^1 and f^2 rules, the *MA* class, and the second class that consists of strings that encode values of portfolio fraction as the *P* class of rules. Both classes of rules will affect the determination of the price levels through agents' savings decisions. Thus, even though the two classes of rules that are updated separately, the evolution of each class is affected by the make-up of the populations representing the other class of rules through prices and rates of return.

with constant exchange rate. The particular level of the exchange rate selected by the adaptive algorithm depends on the initial conditions.

[5] Arifovic, 1996, showed that when this is the only class of rules used by agents, evolution results in persistent fluctuations of the exchange rate. Arifovic and Gencay, 1998 showed that the time series generated in this environment exhibit chaotic behavior.

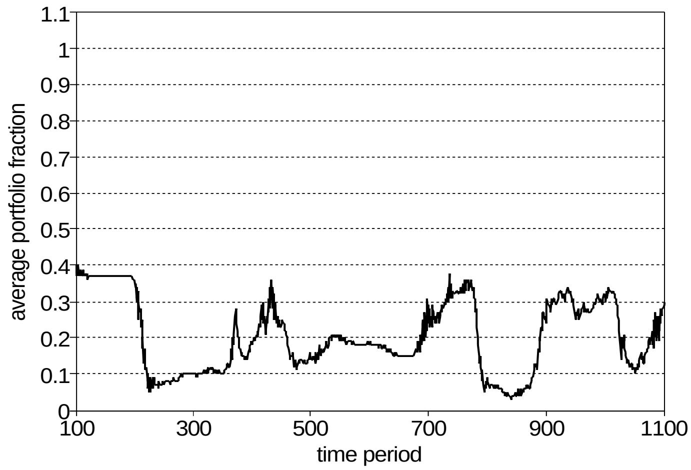

Fig. 3. Average portfolio fraction of MA rules

How does the addition of this class of rules affect the behavior of the economy? Figures 3, 4 and 5 present behavior observed in one of the simulations. Figure 3 shows the behavior of $\bar{\lambda}_t$ of the first class of agents, $\bar{\lambda}_t^{ma}$, and figure 4 shows the behavior of $\bar{\lambda}_t$ of the second class of agents, $\bar{\lambda}_t^{p}$. The difference between the two is noticeable. While both exhibit wide and persistent fluctuations, the behavior of $\bar{\lambda}_t^{ma}$ is less erratic, the amplitude of fluctuations is smaller, periods of upward and downward movements are longer, and except for one instance where both $\bar{\lambda}_t^{ma}$ and $\bar{\lambda}_t^{p}$ take values very close to 1, $\bar{\lambda}_t^{ma}$ generally takes lower values than λ_t^{p}. The main impact of the addition of P class of rules on the behavior of λ_t^{ma} is that it does not converge to 1 or to 0. Instead it exhibits persistent fluctuations that do not die out over time. Fluctuations of λ_t^{ma} and λ_t^{p} result in continuing fluctuations of $R_{1,t}$, $R_{2,t}$ and e_t.

The co-evolution of the two classes of rules is quite interesting and is the subject of investigation. The issues being examined are: exact make-up of each of the two classes of populations, the impact of each class of rules on the behavior of the other class, the welfare implications for agents using different classes of rules, and finally the time-series properties of the simulated data.

4 Further Research

The paper develops a framework for studying the artificial foreign exchange market within the context of the general equilibrium monetary model with endogenous price determination. The version of the model in which agents

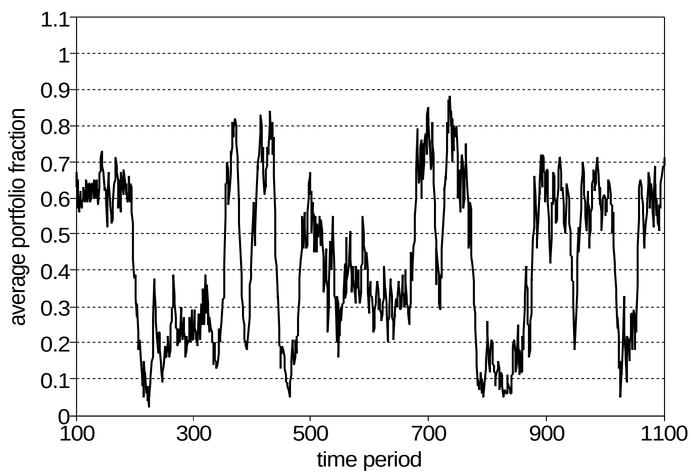

Fig. 4. Average portfolio fraction of P rules

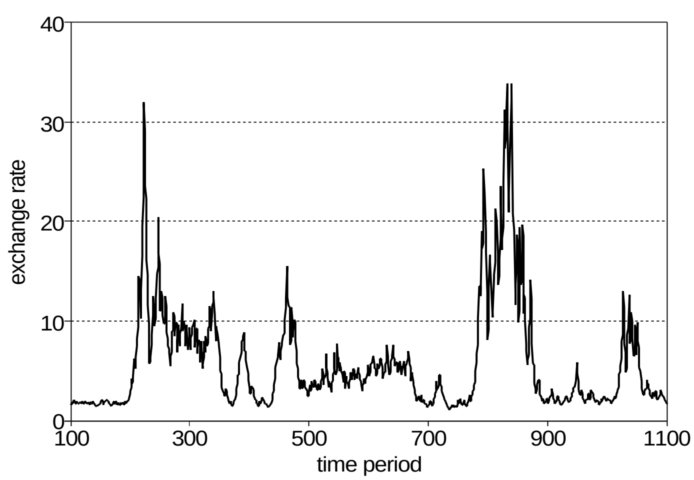

Fig. 5. Exchange rate

are rational does not provide a way to determine the portfolio fraction value. The reason is that agents are indifferent between the currencies that have the same rates of return in the homogenous-expectations equilibrium. The model described in this paper establishes the link between the exchange rate forecast and the portfolio decision and thus provides a way to model and examine the co-evolution of different forecasting rules.

In the model, money is the only available asset and its only role is that of the store of value. In addition, there are no restrictions on foreign currency holdings. [6] Finally, agents adopt different decision rules, and thus make heterogenous portfolio decisions.

These features of the AFEM make it quite appropriate and conveninient for examination of the exchange rate behavior under the flexible exchange rates system. Trading in foreign exchange markets that results in observed volatility is based on differences in the expected rates of return on different currencies. Thus, the main role of money in these transactions is that of the store of value. In addition, in the world of heterogenous beliefs, the rates of return on currencies need not be the same, and this inequality becomes the crucial driving force of the dynamics. In this respect, the AFEM captures the features of trading in real world foreign exchange markets, where rates of return on different currencies are not equalized despite a great degree of mobility and absence of restrictions on foreign currency holdings. A number of extensions of the basic framework are currently under consideration.

First, an environment that is more interesting in terms of the fundamentals will be developed, e.g. specification of different monetary and fiscal policies, definition of a stochastic process that governs the shocks to the production technology, addition of capital to the production technology, endogenous labor supply etc. Thus, the AFEM framework will allow examination of the impact of the shocks to the fundamentals and of their interaction with the dynamics that are driven by changes in agents' beliefs on the exchange rate behavior.

Second, a number of different forecasting rules will be added. We can then examine the impact of different forecasting techniques on the behavior of the exchange rate. Agents will be given an opportunity to choose among different rules and techniques, i.e. all the rules will be subjected to the evolutionary pressure. It will be interesting to examine what rules and techniques survive the selection pressure and whether the evolution results in the selection of a single rule or in the continuous extinction and reappearance of different decision rules. (The framework can also be extended to include, for example, classifier-system type of predictor rules, similar to those used in Arthur et al., 1997, and neural networks).

Third, the model presented in this paper can be used as the basis for developing a framework that can address the question of the impact of technical trading rules on foreign exchange markets. Since the era of floating exchange rates began in the early 1970s, technical analysis has been widely adopted by foreign currency traders.[7] This is partly due to the poor predictive

[6] Obviously, restrictions on foreign currency holdings can be added to the model.

[7] Taylor and Allen (1992) present the results on the issue of technical analysis by major dealers in the foreign exchange market in London.

(out-of-sample) performance of both the structural and the non-structural, time-series exchange rate models. The AFEM framework will provide an environment in which to examine the impact of different trading rules and their performance in competition with alternative forecasting techniques.

References

1. Arifovic, J. (1994) Genetic Algorithm Learning and the Cobweb Model. Journal of Economic Dynamics and Control 18: 3-28.
2. Arifovic, J., (1995) Genetic Algorithms and Inflationary Economies, Journal of Monetary Economics 36, 219-243.
3. Arifovic, J. (1996) The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies, Journal of Political Economy 104: 510-541.
4. Arifovic, J. and R. Gencay (1998) Statistical Properties of Genetic Learning in a Model of Exchange Rate, Journal of Economic Dynamics and Control, forthcoming.
5. Arthur, B., 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, Addison-Wesley.
6. Franke, R. (1998), Behavioral Heterogeneity and Genetic Algorithm Learning in the Cobweb Model, Journal of Evolutionary Economics 8, 383-406.
7. Kareken, J., and Wallace, N. (1981). “On the Indeterminacy of Equilibrium Exchange Rates”, Quarterly Journal of Economics 96:207-222.
8. LeBaron, B., Arthur, W. and Palmer, R. (1999) “Time series properties of an aritifical stock market”, Journal of Economic Dynamics and Control 23, 1487-1516.
9. Sargent, T.J. (1993) Bounded Rationality in Macroeconomics, Claredon Press, Oxford.
10. Taylor, M.P. and H. Allen (1992) The Use of Technical Analysis in the Foreign Exchange Market, Journal of International Money and Finance 11, 304-14.