

Learning Dynamics in an Artificial Currency Market[‡]

Christophre Georges
Department of Economics
Hamilton College
Clinton, NY 13323

cgeorges@hamilton.edu
<http://academics.hamilton.edu/economics/cgeorges/>
315-859-4472

This Version: July 2001

Abstract

This paper considers the behavior of the exchange rate in a very simple artificial currency market with two currencies and artificial agents who evolve their forecast rules over time via a genetic algorithm. The forecast rules are very simple but can be non-linearly misspecified. When this misspecification is suppressed, learning tends to be complete. When it is not suppressed, learning can generate persistent exchange rate dynamics.

Introduction

This paper considers the behavior of the exchange rate in a very simple artificial currency market with two currencies and artificial agents who evolve their forecast rules over time via a genetic algorithm. The forecast rules are very simple but can be non-linearly misspecified. When this misspecification is suppressed, learning tends to be complete. When the term is not suppressed, learning can generate persistent exchange rate dynamics.

Under the non-linear forecast rule, in the absence of learning, there is a continuum of forecast rule parameter values that are consistent with correct expectations in steady state. However, along this continuum, the steady state equilibrium of the exchange rate dynamics is not everywhere stable. Learning may drive the parameter values into a region of the parameter space in which the steady state is locally stable, in which case learning is likely to be complete. However, learning may also

[‡] I would like to thank Jim Bullard, Tom Michl, John Miller, Jeff Pliskin, William Brock, and participants in the 7th Annual Conference of the Society for Computational Economics for their helpful comments. All errors are mine.

drive the parameter values into a region of the parameter space in which the steady state is locally unstable, in which case learning may subsequently fail to converge and instead lead to persistent exchange rate dynamics.

Background

The analysis in this paper follows the recent literature on learning in macroeconomics, finance, and game theory which replaces the assumption that agents are rational and know a great deal about their environments with the assumption that they use rules of thumb and try to improve upon these rules over time through experimentation and imitation. Such boundedly or procedurally rational behavior may arguably be a closer metaphor for the actual behavior of economic agents, since most substantive decisions are made in complex environments under conditions of substantial uncertainty. Learning dynamics can also serve as the mechanism by which a particular equilibrium is selected in environments with multiple equilibria or conversely may introduce persistent disequilibrium dynamics into otherwise well behaved models.¹

Arifovic (1996) and Arifovic and Gencay (2000) have generated chaotic exchange rate dynamics in an overlapping generations model with learning via a GA. In their model, agents live for two periods and adopt behavioral rules of thumb for consumption and currency demand based on the hypothetical performances of these rules of thumb in the previous period.

In the present paper, the GA operates on long lived agents' rules for formulating forecasts of exchange rate movements in a very simple market environment. Interesting dynamics arise from the interaction of learning with the non-linear misspecifications of forecast rules. This approach can be motivated by Grandmont's (1998) insight that local instability is likely to arise if agents are allowed to temporarily extrapolate a variety of non-linear trends that would not be warranted under rational expectations (i.e., under deduction if the agents knew the true model), but appear warranted by the historical path of returns which are generated by the combination of model fundamentals and expectations.

Structure of the Market

There are two currencies. Each yields a certain nominal return denominated in that currency. These returns are based on investment opportunities denominated in these currencies that are not modeled here but rather are taken to be exogenous.

¹ There is a rich and growing literature on learning and adaptation in economic contexts. For some examples of and perspectives on this literature see Arthur et al. (1997), Aumann (1997), Axtell (2000), Axtell et al. (2001), Blume and Easley (1992), Brock and Hommes (1997, 1998), Brock and deFontnouvelle (2000), Bullard (1994), Cross (1983), Evans and Honkapohja (1999, 2001), Fudenberg and Levine (1998), Levy et al. (1994), Lux (1995, 1998), Rubenstein (1998), Sargent (1993, 1999), Sandroni (2000), Simon (1955,1978), Sobel (2000), Young (1998).

There are no other fundamentals, and there are no forward or futures markets, so the exchange rate in the model is grounded solely by uncovered interest parity.

Agents are risk neutral and wish to maximize the expected returns on their portfolios of currencies, valued in their home currencies. The return to converting one unit of currency two to currency one, holding it for one period, and then converting back to currency two is $\frac{e_{t+1} \cdot (1+r_1)}{e_t}$, where e is the value of currency one in terms of currency two and r_1 is the nominal rate of return on currency one.

Perfect Foresight Equilibrium

If all agents had perfect foresight, then for the currency market to be in equilibrium, this return would have to be equal to the nominal return on investments denominated in currency two, $(1 + r_2)$. Define $z \equiv \frac{1+r_2}{1+r_1} - 1$, and $x_t \equiv \frac{e_t}{e_{t-1}} - 1$. Then in perfect foresight equilibrium

$$(1) \quad x_{t+1} = z \quad \forall t.$$

The exchange rate would appreciate continuously at rate z . If for example $r_1 > r_2$, the value of currency one would fall toward zero at a constant rate just sufficient to offset its higher nominal rate of return and thus to leave an agent with perfect foresight indifferent between holding the two currencies. The initial exchange rate is arbitrary, and so there are an infinite number of possible perfect foresight equilibrium paths for the exchange rate e . The rate of appreciation x of the exchange rate, on the other hand, is determinate.

Equilibrium Under Imperfect Foresight

I assume that agents do not have enough information about the structure of the environment to form rational expectations of the rate of appreciation of currency 1. Rather they formulate forecasts of this rate of appreciation inductively.² If all agents held the same forecast at each time t , then equilibrium in the currency market would require

$$(2) \quad F_t[x_{t+1}] = z \quad \forall t$$

where F_t denotes the universal forecast at time t .

In our model, agents will have heterogeneous forecasts. However, to substantially simplify the model, I will assume that the temporary market equilibrium in

² As noted above, Arifovic (1996) models the learning of behavioral rules rather than forecasts. Bullard and Duffie (1998a) call the former learning how to optimize and the latter learning how to forecast. For a discussion of this distinction and a variety of other issues that arise in agent based modeling in financial markets, see LeBaron (2001a). For evidence on the prevalence of technical analysis in actual foreign exchange markets, see e.g., Taylor and Allen (1992), Menkhoff (1998), and Lui and Mole (1998).

any given period can be functionally related to an average forecast (i.e., to the forecast of a representative agent). Specifically, I assume that market clearing holds when the average forecast in the market satisfies condition (1).³ Then denoting $F_t^i[x_{t+1}]$ as the forecast of x_{t+1} held by agent i at time t , and $\bar{F}_t[x_{t+1}]$ as the arithmetic average of these time t forecasts across the agents in the population, equilibrium is characterized by

$$(3) \quad \bar{F}_t[x_{t+1}] = z \quad \forall t.$$

In equilibrium the *representative* investor expects the rate of appreciation of currency one to just offset the difference in the nominal returns of the two currencies.

Forecast Rules

We suppose that agents use forecast rules of the following form

$$(4) \quad F_t^i[x_{t+1}] = a_t^i + b_t^i \cdot x_t + c_t^i \cdot x_{t-1}^2$$

where a_t^i , b_t^i , and c_t^i are scalars that can vary across agents i and time t . I.e., in each period, agents believe that the rate of appreciation next period will be well approximated by a linear function of the current rate of appreciation and the square of last period's rate of appreciation.⁴ While these forecast rules are misspecified, we hypothesize that this misspecification might appear reasonable to individual agents in the model out of steady state equilibrium.

Given these forecast rules (4), the average forecast is

$$\bar{F}_t[x_{t+1}] = \bar{a}_t + \bar{b}_t \cdot x_t + \bar{c}_t \cdot x_{t-1}^2$$

where \bar{a}_t , \bar{b}_t , and \bar{c}_t are the arithmetic averages of the agents' individual parameter values at time t . Then, according to equilibrium condition (3), the temporary equilibrium at each time t will be given by

$$\bar{a}_t + \bar{b}_t \cdot x_t + \bar{c}_t \cdot x_{t-1}^2 = z.$$

³ Heterogeneity is still important in the model as the average forecast will evolve as the distribution of forecasts across agents evolves. Grandmont (1998) uses a similar device (an "average expectations function") when he considers heterogeneous beliefs, as do Evans and Honkapohja (2001) and others. We do not address further the question of how this temporary equilibrium is attained.

⁴ We assume that agents can condition their expectations on the current period's exchange rate, which itself is determined by the average expectation of the agents. Thus, we are assuming that the agents adjust their forecasts as they witness the current period's exchange rate emerge, and that trade does not take place in each period until this process is complete.

Thus, equilibrium is characterized by⁵

$$(5) \quad x_t = \frac{1}{\bar{b}_t} \cdot ((z - \bar{a}_t) - \bar{c}_t \cdot x_{t-1}^2) \quad \forall t.$$

For arbitrary constant values of \bar{a} , \bar{b} , and \bar{c} , this system can display a variety of different local and global dynamics. For example, holding $\bar{b} = 1$ and $\bar{a} = 0$ we get the perfect foresight equilibrium $x_t = z \forall t$ if $\bar{c} = 0$. However as \bar{c} is progressively lowered, the steady state of (5) diverges from z , and we first have dampened oscillations near this steady state and then pass through bifurcations into ranges with limit cycles, chaos, and explosive dynamics.

Consistent Steady State Expectations

Average expectations are correct in steady state in all periods if the following condition is satisfied:

$$(6) \quad z = \bar{a} + \bar{b} \cdot z + \bar{c} \cdot z^2.$$

This condition (6) is linear in the average forecast rule parameters \bar{a} , \bar{b} , and \bar{c} , and so the values of these parameters that satisfy it lie on a plane, a section of which is represented in Figure 1.




FIGURE 1: Consistent steady state expectations locus.

If the average forecast rule parameters satisfy (6), then the perfect foresight equilibrium $x = z$ is a steady state under (5), so that if the appreciation rate x is initially at this level, it will remain there over time. However, even if the average

⁵ In the terminology of Evans and Honkapohja (1999, 2001), (5) is the “actual law of motion” and (4) represents the “perceived law of motion.” In Grandmont’s (1998) terminology, (5) describes the “actual temporary equilibrium dynamics.”

forecast parameters are consistent with (6), if the rate of appreciation x of the exchange rate starts out of steady state, the exchange rate will change over time under (5).⁶ Furthermore, these dynamics need not be locally asymptotically stable. Indeed, various out of steady state dynamics (in the absence of learning) including dampened cycles, limit cycles, and chaos are still supported if the average forecast parameters are constrained to satisfy (6). A limit cycle is illustrated in Figure 2.

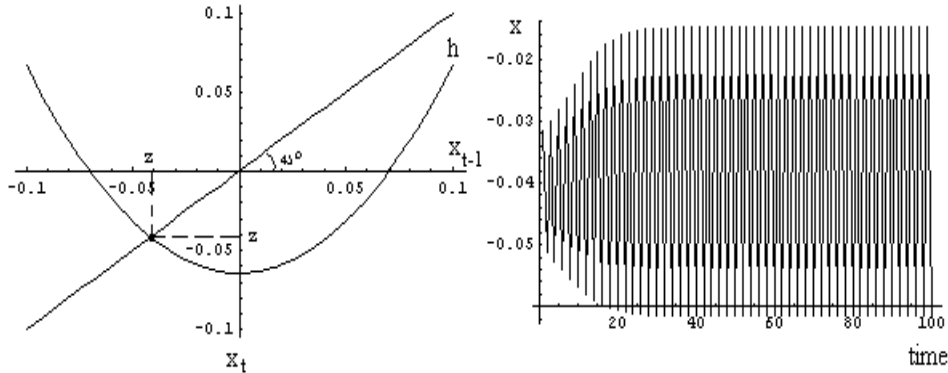


FIGURE 2 The left hand panel shows the equilibrium mapping (5) expressed as $h : x_{t-1} \rightarrow x_t$ for average parameter values $\bar{a} = .1$, $\bar{b} = 2.2$, and $\bar{c} = -28.8$, given rates of return $r_1 = .2$, $r_2 = .15$. These average parameter values satisfy (6), and $x_t = z \approx -0.042$ is a steady state of (5) given these values. The right hand panel gives the evolution of x_t over 100 rounds from initial condition $x_0 = z + .01$.

We are interested in this paper in exploring how agents' forecast parameters evolve relative to the consistent steady state expectations locus (6) under learning as well as how x evolves relative to the perfect foresight equilibrium z .

Learning

If agents were to adopt arbitrary values of a^i , b^i , and c^i which did not change over time, it is highly unlikely that (6) would be satisfied so that z would be a steady state for x under (9). Consequently, the actual appreciation rate would typically not be consistent with either individual expectations (4) or average expectations.⁷ Thus, of particular interest to us will be the way in which agents update their forecast rule parameters a^i , b^i , and c^i over time, and the effect of this learning process on the equilibrium exchange rate under dynamics (5).

⁶ Note that, the average forecast in $t - 1$ is correct if $x_t = z$. However, this condition is not guaranteed by the consistent steady state expectations condition (6) unless x_{t-1} is also equal to z .

⁷ For example, for arbitrary \bar{a} and \bar{b} and $\bar{c} = 0$, the appreciation rate under (5) will be constant over time at a level different from z except on a parameter set of measure zero.

We have assumed that, even though this is a very simple market environment, incomplete information on the part of the agents about the environment makes deductive expectations formation impossible. Nevertheless, this does not mean that agents must be satisfied using purely ad hoc forecasting rules. In such an environment, agents will be interested in modifying their forecast rules over time in response to systematic errors in their forecasts.

Here we model the evolution of forecast rules as a process of social learning via a genetic algorithm.⁸ I follow Arifovic (1994) and Bullard and Duffie (1998b) in modifying the GA to serve as a closer metaphor for social learning, which is more Lamarckian than genetically based evolution.⁹ Here, long lived agents are periodically able to compare their forecast rule parameter values to those of other agents as well as to combinations of these values and to randomly drawn values. Thus, agents are able to experiment with alternative forecast rules on a limited basis in each period and adopt the rule among these alternatives with the smallest forecast error in the last period.

I consider two cases. The first is a benchmark case in which the quadratic term in agents' forecast rules is suppressed. Thus, in Case 1, $c_t^i = 0$ for all i and t . In Case 2, the variation in a_t^i is suppressed and we consider learning on b_t^i and c_t^i . To summarize:

- Case 1: Agents can vary a_t^i and b_t^i .
- Case 2: Agents can vary b_t^i and c_t^i .

The GA

In each of the two cases we consider, agents are able to update two of their forecast rule parameters over time. Thus, to implement the genetic algorithm in each case, we endow each agent at the beginning of the simulation with a chromosome which codes two parameters. The chromosome is a 32 bit binary string with the first 16 bits coding one parameter and the latter 16 bits coding the second parameter. To decode the chromosome, each 16 digit binary number is converted

⁸ There are numerous approaches in the literature to updating individual agent's forecast rules in addition to using GAs. For example, Arthur et al. (1997) and LeBaron et al (1999) use a classifier system in conjunction with a GA to update individual agents' forecast rules. LeBaron (2001b) specifies forecasting rules as artificial neural nets. The GA was introduced by Holland (1975) who argued that it gives a method for searching complex decision spaces in a way that provides a good balance between the benefits and costs of experimentation for on-line decision problems. For general treatments of this method see Goldberg (1989) or Mitchell (1996). For other applications in economics see for example Holland and Miller (1991), Andreoni and Miller (1995), Bullard and Duffie (1998a).

⁹ Arifovic (1994) modifies the standard GA by including an election operator by which new chromosomes are evaluated before being admitted into the population. As noted above, Arifovic (1996) and Arifovic and Gencay (2000) apply such a GA to behavioral rules in an artificial currency market. Bullard and Duffie (1998b) further modify the GA to allow long lived agents to retain their own chromosomes over time.

to base 10 and then normalized so that the parameter value falls within a desired support.¹⁰

In each period $t = 1, 2, \dots$, the equilibrium rate of appreciation of the exchange rate x_t is determined according to (5). The *fitness* or performance of a forecast rule in any period $t = 3, 4, \dots$ is taken to be minus the absolute value of its forecast error in the previous period.¹¹ Thus, in period 3, the fitness of a rule depends on how well it would have forecasted x_2 given x_1 and x_0 .¹²

Starting in the third round, agents' forecast rules are allowed to evolve through individual experimentation and imitation. In any given period, some members of the population are allowed to update their rules. Each of these agent selects at random two other agents from the population. Of those two, the agent with the higher current fitness is retained as the comparison agent. The rules of the agent and the comparison agent are then randomly combined (by crossover: replacing segments of the agent's original bit string with the corresponding segments of the comparison agent's bit string) and/or mutated (by flipping bits of the original or combined string) to produce a third candidate rule. Of these three candidates (the agent's original rule, the rule of the comparison agent, and the combined and/or mutated rule), the agent selects the one with the highest fitness (smallest forecast error in the last period) to use in the current period.¹³

This experimentation and social learning is incremental in the sense that agents retain their most recent rules as candidates in each period and these rules are the basis for mutation and crossover. Further, for low values of the mutation rate, mutation (pure experimentation) will tend to produce candidate rules which are 'close' to the old rules in the sense that few bits will be flipped (i.e. , the rules will be close in Hamming distance). Note however that the parameter values encoded in a new candidate chromosome may be quite different from the agent's original values, since flipping a single bit can correspond to a large change in a parameter

¹⁰ For example, in Case 1, the first 16 bits of each agent i 's initial chromosome codes the first parameter a_1^i , and the latter 16 bits codes the parameter b_1^i . To decode the chromosome, each 16 digit binary number is converted to base 10 and then normalized so that $a^i \in [-.1, .1]$ and $b^i \in [-1, 1]$. Specifically, in the simulations reported below for Case 1, each agent is given the initial chromosome (0,0,1,1,0,...,0|1,0,1,0,...,0), which decodes to approximately $a_1^i = -.0625$, $b_1^i = .2500$.

¹¹ Since we use a tournament style selection process for reproduction, only the relative fitness values will be relevant.

¹² The initial rate of appreciation x_0 is an arbitrary initial condition.

¹³ Specifically, in a given period, each agent considers updating her strategy with probability *pproduce*, includes the pure strategy of the comparison agent (candidate 2) as a candidate strategy with probability *pcompare*, and conducts crossover with probability *pcross*. The two crossover points (one for each of the two 16 bit segments) are selected randomly, and the bits of the agent's chromosome segment to the right of each crossover point are replaced by the corresponding bits of the comparison agent's chromosome. Each gene of the resulting third candidate chromosome is mutated (flipped) with probability *pmutate*.

value.¹⁴ Thus, the genetic operators allow agents to search the parameter space globally, rather than just locally.

Case 1: Quadratic Term Suppressed

As a benchmark, we consider the model when the quadratic term in the forecast rules is suppressed, i.e., in which $c_t^i = 0$ for all i and t . Forecast rules are then simple AR(1) rules. The rate of appreciation x of currency one thus follows (5) with $\bar{c}_t = 0 \forall t$, and so

$$(5') \quad x_t = \frac{1}{\bar{b}_t}(\bar{a}_t - z) \quad \forall t.$$

Then, if agents were to adopt arbitrary values of a^i and b^i which did not change over time, currency one would appreciate at a constant rate over time according to (5'). The average forecasts would be correct in this steady state equilibrium if the average forecast rule parameters \bar{a}_t and \bar{b}_t satisfied

$$(6') \quad z = \bar{a} + \bar{b} \cdot z$$

in which case the equilibrium would be the perfect foresight equilibrium $x_t = z \quad \forall t$.




FIGURE 3: The consistent steady state expectations locus for Case 1.

However, if condition (6') was not satisfied, then the equilibrium exchange rate (5') would appreciate at a different rate than the perfect foresight gross rate z , and average expectations would be incorrect in each period. By (3), the average

¹⁴ E.g., in Case 1, flipping the first bit of a chromosome causes a^i to jump by half of its range, whereas flipping the 16th bit causes a very small change in a^i .

expected appreciation rate is driven to z in each period, but for combinations of parameters not satisfying (6'), the actual appreciation rate is different from z .¹⁵

Results From Case One

For a variety of runs, I find that learning is complete. A sample run is given in Figure 4 below.¹⁶ We see in Figure 4A that the rate of appreciation of the exchange rate fluctuates away from the perfect foresight equilibrium early on, but moves rapidly toward the perfect foresight value over time, as agents update their forecast rules in an effort to reduce their forecast errors. We see that in figure 4B, the forecast rule parameters a and b continue to vary substantially even as the rate of appreciation of the exchange rate becomes close to its equilibrium value. The average values approach the parameter subspace defined by (6) and then wander near that subspace for some time.




FIGURE 4A: The equilibrium appreciation rate (x_t) in a simulation over 250 rounds. The lighter line is the perfect foresight equilibrium appreciation rate.

¹⁵ The forecast error for agent i in period $t-1$ (as revealed in period t) is $F_{t-1}^i[x_t] - x_t = a_{t-1}^i + (b_{t-1}^i - 1) \cdot (z - \bar{a}_t) / \bar{b}_t$. The average forecast error is $\bar{F}_{t-1}[x_t] - x_t = \bar{a}_{t-1} + (\bar{b}_{t-1} - 1) \cdot (z - \bar{a}_t) / \bar{b}_t$.

¹⁶ Parameter values are as follows. The rates of return are set to $r_1 = .2$, $r_2 = .15$, so that $z = -0.0416667$. The GA parameter values are $preproduce=.2$, $pcross=.2$, $pmutate=1/16$, and $pcompare=.2$. The population size is 40.




FIGURE 4B: The evolution of \bar{a} (darker line) and \bar{b} , the average values of the forecast rule parameters. Simulation is the same as above.

The evolution of individual agents' forecast parameters is given in Figure 4C. We see that their values at first disperse from their common initial value, then tend to move toward the consistent expectation locus (6), and ultimately converge to a single (arbitrary) point on that locus.




FIGURE 4C: The evolution of the population of parameter values a^i and b^i .

It is worth noting that if we start with a common value of $b^i < 0$, then learning

appears to be initially destabilizing. However, as agents experiment with new rules, the average value of b^i tends to become positive and agents are then drawn toward (6').

It is also worth emphasizing that learning under the GA above apparently admits a continuum (characterized by (6')) of steady states for the forecast rule parameters, even though the steady state value of the appreciation rate is unique ($x = z$). Agents' parameter values tend to wander toward this locus over time, but do not tend to converge to any particular point on the locus. This is in contrast for example with the pseudo learning rule given by Evans and Honkapohja's expectational stability test. If the average values of the forecast parameters (the parameters of the perceived law of motion) are assumed to move toward the corresponding parameter values of the actual law of motion (5'), then we would have dynamics under which $\bar{b} \rightarrow 0$ and $\bar{a} \rightarrow \frac{1}{b}(\bar{a} - z)$ in continuous time. The unique steady state of this learning process would be $b = 0$ and $a = z$ (a particular point on the locus (6')), which is locally unstable.

Case Two: Quadratic Term Active

With any linear autoregressive forecast rule, the (temporary) equilibrium of our model will be characterized by a linear AR process, so that any irregular dynamics will be due to the learning process (i.e., the updating of the forecast rules over time). In contrast to this, the quadratic term in (4) adds a non-linearity to the equilibrium dynamics (5) in an otherwise linear model, and so, as noted above, can produce irregular dynamics even in the absence of learning.¹⁷

In Case 2, we hold $a = 0.1$ for all agents and time periods, and allow agents to update their b and c values over time. Given $\bar{a} = 0.1$, the locus of \bar{b} and \bar{c} values that are consistent with correct average expectations in steady state (6) is indicated in the left panel of Figure 5. As with (6') in Case 1, this is a slice of the plane shown in Figure 1.

¹⁷ In Arifovic (1996), non-linearity is introduced not through forecast rules but rather through structure of the underlying OLG model. Persistent fluctuations in Brock and Hommes (1997, 1998) are driven by irregular switching between forecast types with different costs of forecasting. In that model, when the economy is near the rational expectations equilibrium, the costs of generating rational expectations outweigh the benefits to the agents, who then switch to less costly rules of thumb which cause local instability.




FIGURE 5: The locus of \bar{b} and \bar{c} values that are consistent with correct average expectations in steady state (i.e., consistent with equation (6)) given $\bar{a} = 0.1$ is given in the left panel. A bifurcation diagram is given in the right panel. This diagram illustrates the changing nature of the long run attractor of the system (5) as \bar{c} and \bar{b} are varied along the consistent expectations locus given in the left panel.

The local dynamics (near $x = z$) for parameter values on this locus are stable and monotonic for $\bar{c} > 0$ and cyclical for $\bar{c} < 0$. For $\bar{c} < 0$ close to zero, cycles are damped, but as \bar{c} becomes progressively more negative (i.e., as we move down the consistent expectations locus shown in the left panel of Figure 5), the equilibrium passes through a series of bifurcations into limit cycles and chaos before the system becomes globally unstable.¹⁸ These bifurcations are illustrated in the right panel of Figure 3.

Learning adds dynamics to the system during transitional learning periods (by changing \bar{b} , and \bar{c} in (5) over time). If the learning process converges so that the average parameter values (\bar{b} , and \bar{c}) converge in the long run, then the values that are selected determine the long run dynamics of the appreciation rate x . If the learning process does not converge, then it will contribute directly to the long run dynamics of the system.

We find that the learning system tends to converge, with $x_t \rightarrow z$ over time, if it finds itself in a region of the learning-parameter space for which the dynamics under (5) are fairly well behaved. However, if the dynamics under (5) are sufficiently volatile, then learning breaks down and contributes to the volatility of the system. This can lead the system to blow up, to wander into a more stable configuration, or to fluctuate irregularly and persistently.¹⁹ While many of the simulations that we

¹⁸ By (5) we see that there is a negative and unstable root at $x = z$ for $\bar{c} < \frac{-\bar{b}}{2|z|}$.

¹⁹ The outcome depends on 1. initial conditions, 2. learning speeds, 3. size of search space (range of b^i and c^i allowed).

ran led the system to either converge or blow up, we will focus on cases of persistent dynamics below.²⁰

The following simulation is instructive. What we see here is all agents starting out with the same forecast parameters, which are approximately $a_0^i = .1$, $b_0^i = 2$, $c_0^i \approx -22.5$.²¹ In the absence of learning, (5) would produce damped cycles in this case. However, with learning, cycles start damped and then vary in amplitude.




FIGURE 6A: The evolution of the appreciation rate x over 200 periods.

²⁰ The fairly high incidence of the system blowing up would be removed if we were to place reasonable bounds on the forecasts.

²¹ Specifically, each agent starts with the chromosome $(1,1,0,0,\dots,0|0,0,1,0,\dots,0)$ which codes b^i and c^i . The parameter a^i is fixed at 0.1 in all rounds. The decoding of the chromosome is as before, except that the normalization places $b^i \in [-4, 4]$ and $c^i \in [-30, 30]$. The GA parameters are $preproduce=.4$, $pcompare=.2$, $pcross=.2$, $pmutate=1/8$. The chromosome length is 32, and the population size is 40.




FIGURE 6B: The evolution of the average parameter values \bar{b} (darker line) and \bar{c} . Same simulation as in Figure 6A.




FIGURE 6C: Evolution of the population's parameter values b^i and c^i . Same simulation as in Figure 6A.

Figure 6 shows that in this particular run of the simulation learning initially causes the exchange rate to converge toward an appreciation rate of z . During this period, individual agents' rules converge toward the locus of parameter values given by (6) illustrated in Figures 1 and 5. However, as the average value \bar{c} falls, z becomes an unstable equilibrium, and volatility increases. As x fluctuates more wildly, learning breaks down and the population scatters in the parameter space

(i.e., forecast rules become more heterogeneous). However, this scattering causes the average value of \bar{c} to rise, leading the system back into a stable regime, and individuals begin to cluster again around the consistent steady state expectation locus. When the simulation was run for an additional 500 rounds, the cycle of convergence and divergence continued.

This simulation displays some interesting features. First, there is clustered volatility in the appreciation rate, with fluctuations centered broadly on the unique perfect foresight equilibrium $x = z$. While that equilibrium is unique, it is supported by a continuum of forecast rule parameter values, which, in the absence of volatility appears to be absorbing under the learning dynamics. Bursts of volatility are apparently set off by the average forecast parameters drifting into the region of this continuum in which z is a locally unstable steady state equilibrium. Second, heterogeneity (of the forecast rules) appears to increase in response to these bursts of volatility, subsequently stabilizing the appreciation rate dynamics.²²

Here are the results from another 400 round run:




FIGURE 7: The evolution of the appreciation rate x over 400 periods. Same as in Figure 6 except $preproduce=.6$.

Longer Memories

Increasing the memory in the fitness functions does not appear to promote greater stability. We considered weighted sums of the absolute forecast errors from a finite number of past rounds, and considered various weights (of the form $(1/s)^\rho$, where s is the number of periods in the past and $\rho \geq 0$).²³ Indeed, the tendency

²² This is a similar result to that of LeBaron (2001b) where homogeneity leads to low liquidity in the market – common expectations can make it difficult for agents to unwind their positions, leading to large price movements. Above, homogeneity does not lead to volatility per se, but volatility does cause homogeneity to disappear.

²³ E.g., for $\rho = 0$, past forecast errors receive equal weights. For $\rho > 0$, past forecast errors are discounted.

to either converge or blow up appears to be heightened by increasing the memory in the fitness functions in this way.

Shocks

When random shocks were added to the equilibrium dynamics (added linearly to (5)), the qualitative results above are preserved except that the population appears to perpetually scatter and churn along the locus (6). Consequently, for the small population sizes that I have considered (typically 40), the average parameter values fluctuate considerably over time. Consequently, rather than converging, \bar{c} continues to wander and occasionally becomes sufficiently negative so as to produce instability and cause learning to break down. Thus, the presence of noise in the model increases the likelihood that the steady state will become unstable eventually under learning. I.e., in terms of Grandmont's insight, agents tend in this model to eventually observe and extrapolate transitory non-linear trends in destabilizing ways. A sample run is shown in Figure 8:




FIGURE 8: The evolution of the appreciation rate x over 1500 periods with noise.

Not surprisingly, adding noise also makes the dynamics of the appreciation rate more irregular, more closely mimicking the behavior of actual forex markets. Nevertheless, the model presented here is a toy model which is too simple to simulate actual financial market patterns with any accuracy. For example, the simulated appreciation rate shown in Figure 8 displays a high degree of first order serial correlation and close to normal kurtosis. Agents are leaving even very basic linear structure in the time series unexploited,²⁴ and the time series does not display the

²⁴ How much systematic forecast error is acceptable? Hommes and Sorger (1998) argue that we should expect agents to learn to uncover some but not all of the structure underlying the equilibrium dynamics. For example, non-linear structure would be missed by agents using linear forecasting tools. They introduce the notion of a Consistent Expectations Equilibrium under which

leptokurtosis (high incidence of extreme events) displayed in actual market data.²⁵

Conclusion

I have illustrated the exchange rate dynamics for a very simple artificial currency market under a simple non-linear forecast rule with learning. In the baseline case with the non-linearity suppressed, learning tends to be complete, with the rate of appreciation of the exchange rate tending to converge to its rational expectations equilibrium value. However, when the forecast rule is quadratic, persistent out of steady state fluctuations may arise from the interaction of learning and the model.

These simulations are interesting as an illustration of how learning can fail to be complete and produce interesting dynamics in a very simple market environment. However, while the fluctuations that arise in some of the simulations do exhibit clustered volatility, they are unrealistically regular. Given the set of forecast rules at their disposal, investors in the model are unable to identify the regularity that persists in these fluctuations. This is due to the simplicity of the model and particularly the limited memory both in the forecast rules themselves and in the fitness functions that agents use to evaluate these rules. However, the model illustrates the volatility that can be produced by a simple non-linear misspecification in the forecast rules used by boundedly rational agents in an otherwise very simple and linear model.

Literature

- Andreoni, James, and John H. Miller, "Auctions with Adaptive Artificial Agents," *Journal of Games and Economic Behavior*, **10**, 1995, 39–64.
- Arifovic, Jasmina, "Genetic Algorithm Learning and the Cobweb Model," *Journal of Economic Dynamics and Control*, **18**, 1994, 3-28.
- Arifovic, Jasmina, "The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies" *Journal of Political Economy*, **104**:3, June 1996, 510-41.
- Arifovic, Jasmina, "Evolutionary Dynamics of Currency Substitution," *Journal of Economic Dynamics and Control*, **25**, 2001, 395–417.
- Arifovic, Jasmina and Ramazan Gencay, "Statistical Properties of Genetic Learning in a Model of Exchange Rate," *Journal of Economic Dynamics and Control*, **24**, 2000, 981–1005.

agents' expectations are consistent with the actual behavior of the economy in terms of a limited number of linear sample statistics.

²⁵ See, e.g., Baillie and Bollerslev (1990).

- Aumann, Robert J., "Rationality and Bounded Rationality," *Games and Economic Behavior*, **21**, 1997, 2–14.
- Axtell, Robert, "Effects of Interaction Topology and Activation Regime in Several Multiagent Systems," Working Paper, Santa Fe Institute, 2000.
- Axtell, Robert, Joshua Epstein and Peyton Young, "The Emergence of Economic Classes in an Agent-Based Bargaining Model," in Steven Durlauf and Peyton Young, editors, **Social Dynamics**, MIT Press, 2001.
- Baillie, Richard T., and Tim Bollerslev, "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates," *Review of Economic Studies*, **58**, 1990, 565–585.
- Blume, L.E. and D. Easley, "Evolution and Market Behavior," *Journal of Economic Theory*, **58**, 1992, 9–40.
- Bray, M. and N. Savin, "Rational Expectations Equilibria, Learning, and Model Specification," *Econometrica*, **54**, 1986, 1129–1160.
- Brock, William A., and Cars H. Hommes, "A Rational Route to Randomness," *Econometrica*, **65**:5, September 1997, 1059–1095.
- Brock, William A., and Cars H. Hommes, "Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model," *Journal of Economic Dynamics and Control*, **22**, 1998, 1235–1274.
- Brock, William A., and Patrick deFontnouvelle, "Expectational Diversity in Monetary Economics," *Journal of Economic Dynamics and Control*, **24**, June 2000, 725–759.
- Bullard, James, "Learning Equilibria," *Journal of Economic Theory*, **64**, 1994, 468–485.
- Bullard, James and John Duffie, "Using Genetic Algorithms to Model the Evolution of Heterogeneous Beliefs," *Computational Economics*, **00**, 1998a, 1–20.
- Bullard, James and John Duffie, "Model of Learning and Emulation with Artificial Agents," *Journal of Economic Dynamics and Control*, **22**, 1998b, 179–207.
- Cross, John G., **A Theory of Adaptive Economic Behavior**, Cambridge University Press, 1983.
- Evans, George W. and Seppo Honkapohja, "Learning Dynamics," in Taylor, John B. and Michael Woodford, **Handbook of Macroeconomics**, vol. 1A, North-Holland, 1999, pp. 449–542.
- Evans, George W. and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton U. Press, 2001.
- Farmer, J. Doyne, "Physicists Attempt to Scale the Ivory Towers of Finance," *Computing in Science and Engineering*, November/December 1999, 26–39.

- Farmer, J. Doyne, and Shareen Joshi, "The Price Dynamics of Common Trading Strategies," Working Paper, Santa Fe Institute, 2000.
- Fudenberg, Drew and David K. Levine, **The Theory of Learning in Games**, MIT Press, 1998.
- Gaunersdorfer, Andrea, and Cars Hommes, "A Nonlinear Structural Model for Volatility Clustering," manuscript, CeNDEF, 2000.
- Goldberg, D.E., **Genetic Algorithms in Search, Optimization, and Machine Learning**, Addison-Wesley, 1989.
- Grandmont, Jean-Michel, "Expectations Formation and Stability of Large Socio-economic Systems," *Econometrica*, **66**:4, July 1998, 741–781.
- Holland, J.H., **Adaptation in Natural and Artificial Systems**, U. Michigan Press, 1975 (2nd. ed. MIT Press 1992).
- Holland, John H., and John H. Miller, "Artificial Adaptive Agents in Economic Theory," *American Economic Review*, Papers and Proceedings, **81**, May 1991, 365–70
- Hommes, Cars and Sorger, Gerhard, "Consistent Expectations Equilibria," *Macroeconomic Dynamics*, **2**, 1998, 287–321.
- Kandori, Michihiro, George J. Mailath, and Rafael Rob, "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, **61**:1, January 1993, 29–56.
- LeBaron, Blake, "A Builder's Guide to Agent Based Financial Markets," manuscript, Brandeis U., 2001a.
- LeBaron, Blake, "Volatility Magnification and Persistence in an Agent Based Financial Market," manuscript, Brandeis U., 2001b.
- LeBaron, Blake, W. Brian Arthur, and Richard Palmer, "Time Series Properties of an Artificial Stock Market," *Journal of Economic Dynamics and Control*, **23**, 1999, 1487–1516.
- Levy, Moshe, Haim Levy and Sorin Solomon, "A Microscopic Model of the Stock Market: Cycles, Booms, and Crashes," *Economics Letters*, **45**, 1994, 103-111.
- Lui, Yu-Hon, and David Mole, "The Use of Fundamental and Technical Analysis by Foreign Exchange Dealers: Hong Kong Evidence," *Journal of International Money and Finance*, **17**, 1998, 535–545.
- Lux, Thomas, "Herd Behavior, Bubbles, and Crashes," *Economic Journal*, **105**, July 1995, 881–896.
- Lux, Thomas, "The Socio-Economic Dynamics of Speculative Markets: Interacting Agents, Chaos, and the Fat Tails of Return Distributions," *Journal of Economic Behavior and Organization*, **33**, 1998, 143–165.

- Lux, Thomas, and Michele Marchesi, "Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market," *Nature*, **397**, 11 February 1999, 498–500.
- Menkhoff, L., "The Noise Trading Approach - Questionnaire Evidence From Foreign Exchange," *Journal of International Money and Finance*, **17**, 1998, 547–564.
- Mitchell, Melanie, **An Introduction of Genetic Algorithms**, MIT Press, 1996.
- Rubenstein, Ariel, **Modeling Bounded Rationality**, MIT Press, 1998.
- Sandroni, Alvaro, "Do Markets Favor Agents Able to Make Accurate Predictions?" *Econometrica*, **68**:6, November 2000, 1303–1341.
- Sargent, Thomas J., **Bounded Rationality in Macroeconomics**, Clarendon Press, 1993.
- Sargent, Thomas J., **The Conquest of American Inflation**, Princeton U. Press, 1999.
- Simon, Herbert A., "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, **69**, February 1955, 99–118.
- Simon, Herbert A., "Rationality as Process and Product of Thought," *American Economic Review*, **68**:2, May 1978, 1–16.
- Sobel, Joel, "Economists' Models of Learning," *Journal of Economic Theory*, **94**, 2000, 241–261.
- Taylor, Mark P., and Helen Allen, "The Use of Technical Analysis in the Foreign Exchange Market," *Journal of International Money and Finance*, **11**, 1992, 304–314.
- Young, H. Peyton, **An Evolutionary Theory of Institutions**, Princeton University Press, 1998.