

Learning Dynamics With Nonlinear Misspecification [‡]

Christophre Georges
Department of Economics
Hamilton College
Clinton, NY 13323

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

This Version: October 2002

Abstract

This paper considers the behavior of artificial agents who evolve their forecast rules over time in a simple market environment. The forecast rules can be nonlinearly misspecified and thus induce interesting dynamics in the single state variable. In simulations, under the social learning algorithm considered here, learning drives agents' forecast rules toward a weak consistency condition but does not eliminate the misspecification.

JEL Classification: D83, D84, E37, E44.

Keywords: Learning, Expectations, Evolutionary Algorithm

[‡] This paper was prepared for the 8th International Conference of the Society for Computational Economics, June 2002, Aix-en-Provence, France. I would like to thank John Wallace for his research assistance and Hamilton College for financial support.

Introduction

There is a rich and growing literature on learning and adaptation in economic contexts.¹ This paper considers the behavior of artificial agents who evolve their forecast rules over time in a simple market environment. The forecast rules can be nonlinearly misspecified and thus induce interesting dynamics in the single state variable. Learning fails to eliminate this misspecification under the learning algorithm considered here. The paper extends the analysis in Georges (2001) to a more generic market environment.

I consider a very simple market environment in which the temporary equilibrium value of the state variable in each period depends on agents' expectations of the value of this variable in the following period. In this generic model, there is a unique stationary rational expectations equilibrium which is strongly E-Stable in the sense of Evans and Honkapohja (2001), as well as a continuum of bubble equilibria which are not E-Stable.² In simulations under the social learning algorithm considered here, learning drives agents' forecast rules toward a weak consistency condition but does not eliminate the misspecification. The wandering of rules near the weak consistency condition can generate interesting dynamics such as volatility clustering.³

The paper and Georges (2001) illustrate Grandmont's (1998) insight that local instability is likely to arise if agents are allowed to temporarily extrapolate a variety of non-linear trends in the data that are themselves generated by the combination of model fundamentals and expectations.

Structure of the Market

Suppose that the temporary equilibrium value of a single state variable x at each time period t can be represented by the following condition

$$(1) \quad x_t = \sigma_0 + \sigma_1 \bar{F}_t x_{t+1} + \varepsilon_t$$

where $\sigma_0 > 0$ and $\sigma_1 \in (0, 1)$. $\bar{F}_t x_{t+1}$ is the arithmetic *average* of the forecasts of the individual agents at time t of the value of x at time $t + 1$. Thus, I am assuming that, while individual forecasts are heterogeneous, the market equilibrium depends

¹ 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, 2001), Bullard (1994), Chiarella and He (2002a,b), Chen and Yeh (2001), Cross (1983), Evans and Honkapohja (2001), Foster and Young (2001), Fudenberg and Levine (1998), Levy et al. (1994), Lux (1995, 1998), Rubenstein (1998), Sargent (1993, 1999), Sandroni (2000), Simon (1955,1978), Sobel (2000), Tetlow and von zur Muehlen (2001), Vriend (2000), Young (1998).

² Georges (2001) conducts a similar exercise for a stylized currency market. In that model, there is a unique rational expectations equilibrium, and that equilibrium is not E-Stable.

³ The mechanism by which such volatility can be supported by the learning process shares some similarities to those of Youssefmir and Huberman (1997) and Lebaron (2001b).

on the forecast of a representative agent.⁴ I do not address further the question of how this temporary equilibrium is attained.

We can think of (1) as a stylized representation of a very simple asset pricing model or an equation for the evolution of GDP or the aggregate price level in a very simple macroeconomic model.

Rational Expectations Equilibrium

This simple model has a unique stationary rational expectations equilibrium: $x_t = x^* + \varepsilon_t \forall t$, where $x^* \equiv \frac{\sigma_0}{(1-\sigma_1)}$. There is also a continuum of rational bubble equilibria, under which x_t explodes away from x^* at average rate $\frac{(1-\sigma_1)}{\sigma_1}$ over time, as well as non-stationary rational sunspot equilibria. In the case of no shocks ($\varepsilon_t = 0 \forall t$), we can call x^* the stationary perfect foresight equilibrium.

Temporary Equilibrium Under Adaptive Learning

Suppose that agents do not have enough information about the structure of the environment to form rational expectations. Rather they formulate forecasts of this rate of appreciation inductively.

Specifically, suppose that agents use forecast rules of the form

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

where $F_t^i[x_{t+1}]$ is the forecast of x_{t+1} held by agent i at time t , and a_t^i , b_t^i , c_t^i , d_t^i , e_t^i , and f_t^i are scalars that can vary across agents i and time t .⁵

The average forecast is thus

$$\bar{F}_t[x_{t+1}] = \bar{a}_t + \bar{b}_t \cdot (x_t - \varepsilon_t) + \bar{c}_t \cdot x_{t-1} + \bar{d}_t \cdot x_{t-1}^2 + \bar{e}_t \cdot x_{t-2} + \bar{f}_t \cdot x_{t-2}^2$$

where \bar{a}_t , \bar{b}_t , ... are the arithmetic averages of the agents' individual parameter values at time t . Then, according to equilibrium condition (1), the temporary

⁴ 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 simplifying device (an "average expectations function") when he considers heterogeneous beliefs, as do Evans and Honkapohja (2001) and others.

⁵ We assume that agents can condition their expectations on the current period's equilibrium value of the state variable net of the shock ε_t , 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 deterministic part of the current period's value of the state variable emerge, and that trade does not take place in each period until this process is complete. We assume that agents do not see the shock in period t for reasons noted below.

equilibrium at each time t will be given by⁶

$$(3) \quad x_t = \frac{1}{1 - \sigma_1 \bar{b}_t} \cdot ((\sigma_0 + \sigma_1 \bar{a}_t) + \sigma_1 \bar{c}_t \cdot x_{t-1} + \sigma_1 \bar{d}_t x_{t-1}^2 + \sigma_1 \bar{e}_t x_{t-2} + \sigma_1 \bar{f}_t x_{t-2}^2) + \varepsilon_t.$$

For arbitrary constant values of individual agents' forecast rule parameters a_t^i , b_t^i , ... the average parameter values \bar{a}_t , \bar{b}_t ... will be constant over time, and the deterministic part of the nonlinear system (3) can display a variety of different local and global dynamics including dampened cycles, limit cycles, and chaos. If individuals can learn and change their forecast rule parameters over time in response to systematic errors, the induced changes in the average parameter values will create additional dynamics in this system.

Note that the forecast rules (2) being used by the agents are misspecified.⁷ Nevertheless, we hypothesize that such a misspecification might appear reasonable to individual agents in the model at any given time, given the recent evolution of x . We can then ask whether this misspecification tends to be eliminated over time by learning.

Minimal State Variable Forecasts

The smallest number of state variables required to characterize a rational expectations equilibrium in this model is zero. I.e., the minimal state variable (MSV) forecast rule is

$$F_t^i[x_{t+1}] = a_t^i$$

If all agents use this type of forecast rule, the temporary equilibrium induced by these forecasts would be

$$x_t = \sigma_0 + \sigma_1 \cdot \bar{a}_t + \varepsilon_t$$

Expectational Stability

The stationary rational expectations equilibrium $x_t = x^* + \varepsilon_t$ is *expectationally stable* (E-Stable) in the sense of Evans and Honkapohja (2001), given our assumption that $\sigma_1 \in (0, 1)$. If agents were to use MSV forecast rules, then under the continuous time pseudo learning algorithm used in the E-Stability framework, the

⁶ In the terminology of Evans and Honkapohja (2001), (3) is the “actual law of motion” and (2) represents the “perceived law of motion.” In Grandmont’s (1998) terminology, (3) describes the “actual temporary equilibrium dynamics.” Note that if the time t forecasts were not independent of ε_t , the variance of the disturbance term of (3) would depend on the average forecast rule parameter \bar{b}_t which will vary over time. This would provide an additional mechanism for generating conditional heteroscedasticity.

⁷ For example, agents at time $t - 1$ condition their forecasts of x_t on x_{t-3} , whereas under the actual law of motion (3) induced by these forecasts, x_t does not depend directly on x_{t-3} .

representative forecast parameter \bar{a}_t would converge locally to x^* over time.⁸ Since, for $\bar{a}_t = x^*$, the temporary equilibrium (3) is exactly the stationary rational expectations equilibrium (in this MSV forecast rule case), the E-Stability framework selects that equilibrium as the unique outcome of learning. Further, this rational expectations equilibrium is *strongly expectationally stable* in that it is robust to overparameterization as in (2). Indeed, under the E-Stability condition's continuous learning dynamics, the forecast rule parameters b , c , d , e , and f all converge locally to zero under forecast rules of form (2). By the same token, bubble equilibria are not E-Stable in this model, given $\sigma_1 \in (0, 1)$. Some details are provided in an Appendix.

Thus, the E-Stability criterion selects a particular learning equilibrium ($\bar{a} = x^*$, $\bar{b} = \bar{c} = \bar{d} = \bar{e} = \bar{f} = 0$) which replicates the stationary rational expectations equilibrium. This result also implies that the stationary rational expectations equilibrium would be selected under least squares econometric learning (Evans and Honkapohja, 2001).

Weakly Consistent Expectations

We now ask what are the forecast rules of form (2) that are consistent with the stationary rational expectations equilibrium in the following sense: $\bar{F}_t[x_{t+1}|(x_t - \varepsilon_t) = x_{t-1} = x_{t-2} = x^*] = x^*$. This condition can be written

$$(4) \quad x^* = \bar{a} + (\bar{b} + \bar{c} + \bar{e}) \cdot x^* + (\bar{d} + \bar{f}) \cdot x^{*2}.$$

Note that (4) is linear in the six average forecast rule parameters \bar{a} , \bar{b} , ... and so the values of these parameters that satisfy it lie on a five dimensional hyperplane.

If the average forecast rule parameters satisfy (4), then the perfect foresight equilibrium $x_t = x^*$ is a steady state under the deterministic part of (3), so that if x was initially at this level, it would remain there over time in the absence of shocks ε . However, even if the average forecast parameters are consistent with (4), if x started out of steady state, it will change over time under the deterministic part of (3).

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 (4). Two examples of bifurcation diagrams are given below in order to illustrate the range of possible long run attractors for the deterministic part of the system in the absence of learning. In both cases we set $\sigma_0 = 5$ and $\sigma_1 = 0.5$ so that the stationary perfect foresight equilibrium is

⁸ In the E-Stability framework, the average values of the forecast parameters (the parameters of the perceived law of motion - PLM) are assumed to move toward the corresponding parameter values of the actual law of motion (3) induced by the PLM. This learning process is formalized in a notional continuous time rather than in real time.

$x^* = 10$. In Figure 1A, d is varied constraining (a,d) to satisfy (4) for the case $b = c = e = f = 0$ (in which case (3) is a quadratic map).




FIGURE 1A: Bifurcation diagram as d is varied along the consistent expectations locus holding $b = c = e = f = 0$.

In Figure 1B, f is varied constraining (a,f) to satisfy (4) for the case $b = -0.04$, $c = 0.02$, $d = -0.11$, $e = 0.15$.




FIGURE 1B: Bifurcation diagram as f is varied along the consistent expectations locus holding $b = -0.04$, $c = 0.02$, $d = -0.11$, $e = 0.15$.

We are interested in this paper in exploring how agents' forecast parameters evolve relative to the consistent expectations locus (4) under learning as well as how x evolves relative to the stationary perfect foresight equilibrium x^* .

Learning

Here we model the evolution of forecast rules as a process of social learning and

experimentation using an evolutionary algorithm.⁹ 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 which would have had the smallest forecast error in recent periods.

Specifically, each agent is endowed with arbitrary initial real values of forecast rule parameters $a_t^i, b_t^i, \dots, f_t^i$. In each period the equilibrium value of the state variable x_t is determined according to (3). After an initial set of periods without learning, each agent considers updating her forecast rule (changing the parameter values) with some probability. An agent who is considering updating her forecast rule will generate two new rules to compare to the rule that she used last period. Of the three, she selects the rule with the highest fitness (see below).

The first comparison rule is simply the higher fitness rule of two other agents randomly selected from the population. Thus, the agent may simply imitate another agent in the population. The second comparison rule is generated by applying crossover and mutation operators (each with some probability) to the agent's rule from the last period and the first comparison rule. Thus, the agent may experiment with novel mutations of its own previous rule and may adopt some elements of the comparison agent's rule rather than imitating the rule in its entirety. Further agents can vary in the range of mutations that they consider.¹⁰

⁹ There are numerous approaches in the literature to representing and updating individual agent's forecast rules. For example, Arthur et al. (1997) and LeBaron et al (1999) use a classifier system in conjunction with a genetic algorithm (GA). LeBaron (2001b) specifies forecasting rules as artificial neural nets. Chen and Yeh (2001, 2002) code rules as genetic programs (GP). Brock and Hommes (1997, 1998, 2001) and Chiarella and He (2002a,b) update strategy mixes at the population level based on results from discrete choice theory. Our algorithm is essentially a modified real coded GA. I follow Arifovic (1994, 1996) in modifying the standard GA (introduced by Holland (1975)) by including an election operator by which new chromosomes are evaluated before being admitted into the population and Bullard and Duffie (1998) in allowing long lived agents to retain their own chromosomes over time. These two variations allow the GA to serve as a closer metaphor for social learning, which is more Lamarckian than genetically based evolution (e.g., traditional single population GAs). Michalewicz (1999), Fogel (2000), and others have argued against the use of binary coded GAs and in favor of adopting codings and evolutionary rules to fit each problem being solved or modeled. See Herrera et al. (1998) for a discussion of real GA operators.

¹⁰ In a given period, each agent considers updating her strategy with probability $p_{reproduce}$, considers imitating a comparison agent's strategy with probability $p_{compare}$, and conducts crossover with that strategy with probability p_{cross} . The crossover factor is selected from a uniform distribution on $(.5 - C, .5 + C)$ where $C > .5$. Each gene of the resulting second comparison chromosome is mutated with probability p_{mutate} by a mutation factor chosen independently from a uniform distribution on $(1 - M, 1 + M)$. There are two values of M , one relatively small (e.g., .1) and one relatively large (e.g., 2), and each agent is assigned one of these values for life (i.e., each experiments either very locally or more broadly at any time). Note that small M agents are able to imitate large M agents over time by adopting their realized chromosomes and vice versa. If crossover does not occur, then the original rule is mutated in place of the crossed-over rule.

Thus, to summarize, the agent will consider three rules:

Rule 1: The rule used by that agent in the previous period.

Rule 2: A rule copied from another agent in the population.

Rule 3: A rule that (may) combine elements from (possibly) mutated versions of rules 1 and 2.

The fitness of a rule is based on its hypothetical forecasting accuracy in recent periods. Specifically, the fitness of a rule at time t is minus a weighted sum of the absolute values of the forecast errors that the rule would have generated in the past $leval$ rounds ($leval$ stands for the length of the evaluation period).¹¹

$$(5) \quad f_t^i = -\sum_{s=t-leval}^{t-1} |F_t^i[x_s : x_{s-1}, x_{s-2}, x_{s-3}] - x_s| \cdot \left(\frac{1}{t-s}\right)^g$$

$g \geq 0$. Thus, in period 4, if $leval = 1$, the fitness of a rule would depend on how well it would have forecasted x_3 given x_2 , x_1 and x_0 .

The experimentation and social learning that takes place in this evolutionary process 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. One drawback is that the range of possible experimentation is limited in any given period. Agents can end up exploring the parameter space widely over time, but only if this comes about incrementally. We have attempted to mitigate this problem by having two agent types and setting a fairly large mutation range for the broad mutation agent type.

Case 1: All AR Terms Suppressed

As noted above, the MSV forecast rule in this model is $F_t^i x_{t+1} = a_t^i$, which corresponds to a belief that x will fluctuate randomly about a constant in each period. This rule is confirmed by the actual economy if $a_t^i = x^* \quad \forall i, t$, in which case $x_t = x^* + \varepsilon_t \quad \forall t$.

In simulations holding $b_t^i = c_t^i = d_t^i = e_t^i = f_t^i = 0 \quad \forall t$, we find that for each agent, a_t^i is drawn toward and becomes centered on x^* over time and that consequently, x_t is drawn toward and becomes centered on x^* over time. It is easy to see why this is the case. According to (1), if the average forecast is above x^* , then the distribution of x will tend to be centered between \bar{a} and x^* . Thus, the average agent's forecast $F^i x = a^i$ is too high, and given a sufficiently small variance of ε will tend, through experimentation and imitation, to be lowered over time. Similarly, an average value $\bar{a} < x^*$ will tend to be driven up toward x^* over time as agents improve their forecasts.

Thus, in this benchmark case, learning selects the stationary rational expectations equilibrium of the model, as predicted by the E-Stability criteria. A difference

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

between the evolutionary process here and both the E-Stability and Least Squared Learning frameworks is that agents only look back a fixed number of periods, and so learning is complete only in the absence of stochastic shocks ($\varepsilon_t = 0$). I.e., the stationary rational expectations equilibrium is selected approximately when the shock term is active.¹²

Here is an example of a run for this case with 100 agents. For all simulations reported in this paper, $\sigma_0 = 5$ and $\sigma_1 = 0.5$ so that $x^* = 10$.¹³

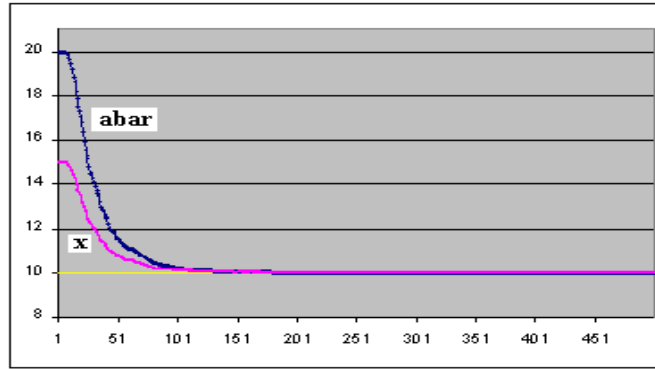


FIGURE 2: The evolution of the state variable x_t and average parameter value \bar{a}_t in a simulation over 500 rounds with $b = c = d = e = f = 0$. $x^* = 10$ and a^i is initially set at 20 for each agent.

Case 2: Nonlinear Terms Suppressed

Adding in variation in the parameters c and e give agents two additional parameters to learn. Both the forecast rule and induced law of motion for x are linear in this case.

We find again that agents learn over time to forecast $x = x^*$ on average. However, there is now a continuum of parameter values (a, b, c, e) that are weakly consistent in the sense defined above. I.e., the consistent expectations locus (4) is three dimensional. We find that learning tends to drive individual agents' forecast rules to this locus, but not toward any particular point on the locus. Particularly,

¹² Under least squares learning with infinite memory, the gain of the updating rule for the least squares estimates of the forecast rule would decline over time at a rate fast enough to make each agent's rule converge. Here, due to finite memory, the gain does not decrease over time. See e.g., Honkapohja and Mitra (2002).

¹³ In this run, the distribution of ε is uniform on $(-0.01, 0.01)$. Each agent evolves in any round with probability 0.2, in which case she considers direct imitation of another agent with probability 0.4, considers crossover with probability 0.2, and mutates each of her six genes (parameter values) with probability 0.3. Half of the population are low range mutators who consider mutations of no more than 10%, half are high range mutators who consider changing by a factor between -1 and 3. The memory *level* in the fitness function is set to 1.

whereas only the point $(x^*, 0, 0, 0)$ is supported by the E-Stability criterion, and so is also locally stable under LSL, the evolutionary algorithm here does not appear to select that forecasting rule globally. Rather, in the absence of shocks, rules appear to tend to converge to arbitrary locations on this locus. With shocks, the rules tend to continue to wander near this locus. If there is a general attraction to $(x^*, 0, 0, 0)$ it is apparently quite weak, whereas the attraction to (4) appears quite strong.

Here is an example of a run for this case with 100 agents. The parameters for the model and learning algorithm are the same as in the simulation in the last section. We see that the distribution of the state variable becomes centered (at least roughly) on $x^* = 10$.




FIGURE 3a: The evolution of x_t and average parameter values \bar{a}_t , \bar{b}_t , \bar{c}_t , and \bar{e}_t in a simulation over 500 rounds with $d = f = 0$. The stationary perfect foresight equilibrium value x^* is 10. \bar{b} , \bar{c}_t and \bar{e}_t are initially 0.2 and are measured on the right hand scale. \bar{a} is initially 20 and is measured along with x on the left hand scale.

We can also see (Figure 3b) that the average forecast rule parameter values are attracted to the consistent expectations locus.¹⁴

¹⁴ In Figure 3b, the distance of the average rule from the locus is measured as the difference between the actual value of \bar{a}_t and the value that would satisfy (4) given the actual values of the other parameters $(\bar{b}_t, \bar{c}_t, \bar{e}_t)$.

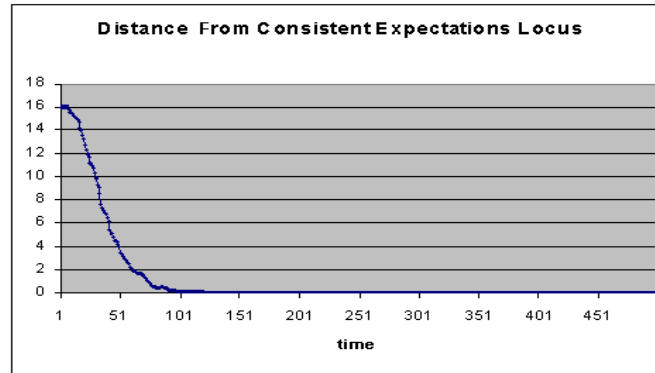


FIGURE 3b: The distance from the consistent expectations locus (4) in the simulation shown in Figure 3a. By the 500th round, the value of the distance is fluctuating in the vicinity of 0.001.

Since these values do not, however, converge to the MSV values (even approximately), linear propagation of shocks (under (3)) persists under learning. It is also worth noting that the average parameter values tend to continue to wander near the locus (4) in response to the shocks ε . While, in the simulation above, the wandering settles down within a few hundred rounds, this is often not the case. The degree of wandering depends in part on the mutation ranges considered by agents. The presence of agents with large mutation ranges tends to cause rapid convergence toward the locus (4) early on if the simulation is started far from equilibrium. The presence of agents with small mutation ranges will then cause substantial wandering near this locus, as agents pursue small fluctuations in x near \bar{x} . Reducing the high and low mutation ranges appears to promote persistent wandering.

An important but unresolved question is whether, for sufficiently small mutation ranges, learning would eventually drive forecasts to the MSV forecast (i.e., the learning equilibrium predicted by the E-Stability condition). I.e., the lack of attraction seen in the simulations reported here could, in principle be attributable to the broad nature of the search process followed by agents (i.e., the agents' inability to explore the search space sufficiently) rather than to the absence of an attractor. It may then be the case that the MSV forecast rule is a global attractor, but its attraction is weak, whereas the attraction toward (4) is strong far from equilibrium – a kind of turnpike result.

Case Three: Quadratic Terms 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 (nonlinear) dynamics will be due to the learning process (i.e., the updating of the forecast rules over time). In contrast to this, the quadratic terms in (2) add a nonlinearity to the

equilibrium dynamics (3) in an otherwise linear model, and so, as noted above, can produce irregular dynamics even in the absence of learning.¹⁵ Thus, increasing the complexity of the forecasting rules has two interesting effects. First, the dynamics of the state variable in the absence of learning can become more complex. Second, learning how to update these rules may become a more challenging problem for agents in the model for any given history of the state variable.

In simulations, we find that, again, there is a tendency for forecast rules to evolve toward the consistent expectations locus (4) and for evolution of the state variable x to become centered on x^* . However, again, if there is a tendency for the system to be attracted to the particular learning equilibrium (set of forecast parameters) supported by the E-Stability criterion, this tendency is not apparent in a variety of simulations. Thus, the nonlinear endogenous dynamics and propagation of shocks persist under learning. Further, as the system wanders near the locus, it can display varied dynamics such as the clustered volatility exhibited below.




FIGURE 4a: Evolution of x_t in a simulation over 1000 rounds with initial forecast parameter values $b = 0$, $c = f = 0.05$, $d = e = -0.05$.

¹⁵ 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 4b: Evolution of the average parameter values $\bar{a}_t, \bar{b}_t, \bar{c}_t, \dots$ \bar{a} is measured on the right hand scale, all others on left.




FIGURE 4c: Evolution of the distance of the average forecast rule from the consistent expectations locus (4).

Here is a run with 50 agents for 1000 periods. All agents start on the consistent expectations locus with $a = 10, b = .5, c = .5, d = -.14, e = 1.8, f = -.14$. Without learning this forecast rule generates the following series for the state variable:




FIGURE 5A: The evolution of x over 1000 periods in the absence of learning.

With learning, the volatility magnification and clustering persists.




FIGURE 5B: The evolution of x over 1000 periods with learning. $level$ is set to 20 periods, $preproduce = .4$, the maximum mutation ranges are 20% for high mutators and 5% for low mutators.




FIGURE 5C: The evolution of the quadratic term parameters \bar{d} and \bar{f} in the simulation shown in Figure 5B.

These simulations display similar features to those in Georges (2001). First, there is clustered volatility in the state variable, with fluctuations centered broadly on the unique stationary rational expectations equilibrium. While that equilibrium is unique, it is supported by a continuum of forecast rule parameter values, which, in the absence of shocks (ε) and out of steady state volatility, appears to be absorbing under the learning dynamics. Bursts of volatility can be set off by the average forecast parameters drifting into the region of this continuum in which x^* is locally unstable under the deterministic part of (3).

Second, the heterogeneity of the forecast rules appears to increase in response to these bursts of volatility, and this heterogeneity can lead, through social learning to the stabilization of the dynamics of the state variable.¹⁶ As volatility increases, agents tend to experiment more broadly and scatter in the parameter space, causing the average values to adjust quickly. From observing a number of simulation runs, there does appear to be a tendency for the system to get kicked into more stable parts of the parameter space following periods of intense volatility. Consider for example the following response of a small population ($N=20$) to a large burst of volatility.

¹⁶ 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.




FIGURE 6A: The evolution of x over 5000 rounds with 20 agents, level=1. Mutation factors are .02 and .005. ε in this run is uniform on $(-0.1, 0.1)$.

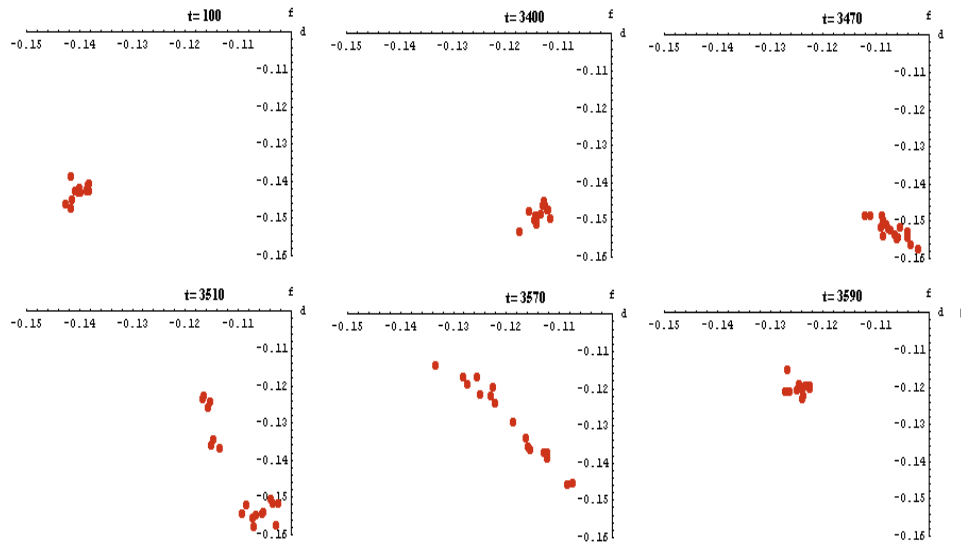


FIGURE 6B: Evolution of d_t^i and f_t^i for the population.

In this run, we see the population drifting in (\bar{e}, \bar{f}) space until the parameters induce a very large burst of volatility. This volatility causes the population to scatter and then converge (via social learning) to a new configuration which, in this case, generates less volatility in the system.




FIGURE 6C: The evolution of the quadratic term parameters \bar{d} and \bar{f} in the simulation shown in Figure 6A,B.

Comments

Increasing the memory in the fitness functions does not appear to promote greater stability in general. Increasing the population size on the other hand does appear to promote stability, and in particular the secular tendency toward stability (lower and more constant volatility).¹⁷

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. In a variety of runs, the time series of x does display leptokurtosis (high incidence of extreme events), though the degree of these fat tails is modest relative to actual market data.¹⁸

Conclusion

I have illustrated the dynamics of a very simple artificial market environment under a simple non-linear forecast rule with learning. In the baseline case with the non-linearity suppressed, learning tends to be near-complete, with the learning selecting the stationary rational expectations equilibrium. However, when the forecast rule is quadratic, additional volatility may arise from the interaction of learning and the model. These simulations are interesting as an illustration of how learning

¹⁷ This is a somewhat unfortunate result found in much of the simulation literature in finance. Many models that generate fat tails in returns from small populations do not do so for large populations for example.

¹⁸ See, e.g., Mantegna and Stanley (2000)

can fail to be complete and produce interesting dynamics in a very simple market environment.

As with Georges (2001), the model presented here 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.

Appendix: E-Stability

Consider the linear case $d = f = 0$.

The *perceived law of motion* (PLM) in this case is

$$\text{A1} \quad x + t = 1 + b x_{t-1} + c x_{t-2} + e x_{t-3}$$

By (1), the induced *actual law of motion* (ALM) is

$$\text{A2} \quad x + t = \left(\frac{1}{1 - \sigma_1 b}\right)((\sigma_0 + \sigma_1 a) + \sigma_1 \cdot (c x_{t-1} + e x_{t-2})) + \varepsilon_t$$

We suppose that the (forecast rule) parameters in the PLM adjust in notional time toward the corresponding values of the induced ALM as follows

$$\begin{aligned} \text{A3} \quad \dot{a}_t &= \frac{\sigma_0 + \sigma_1 a_t}{1 - \sigma_1 b_t} - a_t \\ \dot{b}_t &= \frac{\sigma_1 c_t}{1 - \sigma_1 b_t} - b_t \\ \dot{c}_t &= \frac{\sigma_1 e_t}{1 - \sigma_1 b_t} - c_t \\ \dot{e}_t &= 0 - e_t \end{aligned}$$

This system of equations has a steady state at $a = x^*$, $b = c = e = 0$. The Jacobian of the system linearized about this fixed point is

$$\text{A4} \quad J = \begin{pmatrix} (\sigma_1 - 1) & x^* \sigma_1 & 0 & 0 \\ 0 & -1 & \sigma_1 & 0 \\ 0 & 0 & -1 & \sigma_1 \\ 0 & 0 & 0 & -1 \end{pmatrix}$$

The eigenvalues of this Jacobian are the elements on the main diagonal. Thus, for $\sigma_1 < 1$, these eigenvalues are all (real and) negative, and so the steady state is locally asymptotically stable.

Literature

- 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.
- 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.
- 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 Cars H. Hommes, "Evolutionary Dynamics in Financial Markets With Many Trader Types," Manuscript, April 2001.
- Bullard, James, "Learning Equilibria," *Journal of Economic Theory*, **64**, 1994, 468-485.
- Bullard, James and John Duffie, "Model of Learning and Emulation with Artificial Agents," *Journal of Economic Dynamics and Control*, **22**, 1998, 179-207.
- Chen,-Shu-Heng and Chia-Hsuan Yeh, "Evolving Traders and the Business School with Genetic Programming: A New Architecture of the Agent-Based Artificial Stock Market," *Journal of Economic Dynamics and Control*, **25**:3-4, March 2001, 363-93.
- Chen,-Shu-Heng and Chia-Hsuan Yeh, "On the Emergent Properties of Artificial Stock Markets: the Efficient Market Hypothesis and the Rational Expectations Hypothesis," *Journal of Economic Behavior and Organization*, **49**:1, September 2002, 217-239.

- Chiarella, Carl, and Xue-Zhong He, "Heterogeneous Beliefs, Risk and Learning in a Simple Asset Pricing Model," *Computational Economics*, **19**, 2002a, 95–132.
- Chiarella, Carl, and Xue-Zhong He, "An Adaptive Model of Asset Pricing and Wealth Dynamics With Heterogeneous Trading Strategies," Manuscript, U. Technology, Sydney, May 2002b.
- Cross, John G., **A Theory of Adaptive Economic Behavior**, Cambridge University Press, 1983.
- Evans, George W. and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton U. Press, 2001.
- Farmer, J. Doyne, and Shareen Joshi, "The Price Dynamics of Common Trading Strategies," *Journal of Economic Behavior and Organization*, **49**:1, September 2002, 149–171.
- Fogel, David B., *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*, Second Ed., IEEE Press, 2000.
- Foster, Dean P. and H. Peyton Young, "On the Impossibility of Predicting the Behavior of Rational Agents," *Proceedings of the National Academy of Sciences*, **98**:22, 2001, 12848–12853.
- Fudenberg, Drew and David K. Levine, **The Theory of Learning in Games**, MIT Press, 1998.
- Georges, Christophre, "Learning With Misspecification in an Artificial Currency Market," Manuscript, Hamilton College, 2001.
- Grandmont, Jean-Michel, "Expectations Formation and Stability of Large Socioeconomic Systems," *Econometrica*, **66**:4, July 1998, 741–781.
- Herrera, Lozano, and Verdegay, "Tackling Real-Coded GAs: Operators and Tools for Behavioural Analysis" *AI Review*, **12**, 1998, 265–319.
- Holland, J.H., **Adaptation in Natural and Artificial Systems**, U. Michigan Press, 1975 (2nd. ed. MIT Press 1992).
- Hommes, Cars, "Financial Markets as Nonlinear Adaptive Evolutionary Systems," *Quantitative Finance*, **1**, 2001, 149–167.
- Honkapohja, Seppo, and Kaushik Mitra, "Learning with Bounded Memory in Stochastic Models," forthcoming *Journal of Economic Dynamics and Control*, 2002.
- 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.
- 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.
- Mantegna, Rosario N. and H. Eugene Stanley, **An Introduction to Econophysics: Correlations and Complexity in Finance**, Cambridge U. Press, 2000.
- Michalewicz, Zbigniew, *Genetic Algorithms + Data Structures = Evolution Programs*, 3rd Ed., Springer, 2000.
- 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.
- Tetlow, Robert and Peter von zur Muehlen, "Avoiding Nash Inflation: Bayesian and Robust Responses to Model Uncertainty," manuscript, Board of Governors, December 2001.
- Vriend, Nicholas J., "An Illustration of the Essential Difference Between Individual and Social Learning, and its Consequences for Computational Analysis" *Journal of Economic Dynamics and Control*, **24**, 2000, 1–19.
- Young, H. Peyton, **An Evolutionary Theory of Institutions**, Princeton University Press, 1998.

Youssefmir, Michael and Bernardo Huberman, "Clustered Volatility in Multiagent Dynamics," *Journal of Economic Behavior and Organization*, **32**, 1997, 101–118.