Macroeconomic Dynamics, **13** (2009), 625–655. Printed in the United States of America.
doi:10.1017/S1365100509080262

LEARNING DYNAMICS AND NONLINEAR MISSPECIFICATION IN AN ARTIFICIAL FINANCIAL MARKET

CHRISTOPHRE GEORGES AND JOHN C. WALLACE
Hamilton College

In this paper, we explore the consequence of learning to forecast in a very simple environment. Agents have bounded memory and incorrectly believe that there is nonlinear structure underlying the aggregate time series dynamics. Under social learning with finite memory, agents may be unable to learn the true structure of the economy and rather may chase spurious trends, destabilizing the actual aggregate dynamics. We explore the degree to which agents' forecasts are drawn toward a minimal state variable learning equilibrium as well as a weaker long-run consistency condition.

Keywords: Learning, Expectations, Agent-Based Model

1. INTRODUCTION

This paper considers the behavior of artificial agents who evolve their forecast rules over time in a simple market environment with over-the-counter trading. Asset price dynamics are driven by a combination of nonlinearly misspecified (overparameterized) forecast rules and learning how to forecast with bounded memory. Learning is modeled by an evolutionary algorithm. Agents are heterogeneous and learn by independently experimenting with novel forecasting rules and by imitating the forecasting behavior of others.

Simulations indicate that when agents are restricted to consider only minimum state variable (MSV) forecast rules, price dynamics converge to a noisy version of the stationary rational expectations equilibrium (REE) and forecasts fluctuate around the rational expectations forecast. This result is consistent with the analytical results of Honkapohja and Mitra (2003), who consider least squares learning by a representative agent with bounded memory.

Christophre Georges is Professor of Economics, Hamilton College. John C. Wallace is a graduate of Hamilton College. We thank Rob Axtell, Jim Bradfield, Jim Bullard, Giovanni Dosi, George Evans, Neil Johnson, Blake LeBaron, Thomas Lux, Tarun Ramadorai, participants at the 9th Annual Conference of the Society for Computational Economics, the Fall 2004 Workshop on Interactions and Markets in Pisa, and seminars at Oxford University and the University of Kiel, and two anonymous referees for their helpful comments on earlier versions of the paper. All errors are ours. Financial support was provided by the Fred L. Emerson Foundation and NSF Grant BCS-0738606. Address correspondence to: Christophre Georges, Department of Economics, Hamilton College, Clinton, NY 13323, USA; e-mail: cgeorges@hamilton.edu; URL: http://academics.hamilton.edu/economics/cgeorges/.

© 2009 Cambridge University Press 1365-1005/09

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

When overparameterized rules are allowed, we find a stronger attraction of forecast rules to a weak consistency condition (WCC) that nests the MSV REE rule than to the MSV REE rule. Deviations from the MSV REE rule can be persistent and can produce nontrivial price dynamics exhibiting clustered volatility and excess kurtosis of returns, as well as explosive bubbles and crashes.

Our focus on overparameterized forecasting rules is motivated by Grandmont's (1998) "uncertainty principle." That principle states that local instability is likely to arise if agents are allowed to temporarily extrapolate a variety of nonlinear trends that would not be warranted under rational expectations, but appear warranted by the recent historical path of state variables generated by the combination of model fundamentals and expectations. This approach stands in contrast to the recent literature on adaptive learning by agents with misspecified forecasting models that are underparameterized (rather than overparameterized) relative to the market fundamentals [Evans and Honkapohja (2001, Ch. 13), Cho et al. (2002), Branch and Evans (2006)].[1] It is closer in spirit to that of Bullard et al. (2007), who consider "add factors" in forecasts.

There is now a large literature on heterogeneous agent and agent-based models in finance [as for example surveyed by Hommes (2006), LeBaron (2006), and Samanidou et al. (2007)]. Whereas many extant models [e.g., Lux (1995), Brock and Hommes (1998), Chiarella and Iori (2002), LeBaron and Yamamoto (2006), Alfarano and Lux (2007), and Boswijk et al. (2007)] focus on the interaction of a small number of trader types, we allow for a continuum of technical forecasting rules. This continuum includes rules that are consistent with a variety of momentum and fundamentalist-type beliefs as well as rational expectations.

In companion papers, Georges (2008b) considers the present model with the evolutionary algorithm replaced by least squares learning, and Georges (2008a) considers a quasi-representative agent version of that model without the explicit trading microstructure. In these papers, agents update their forecast rules via least squares with random wait times between updates. The evolutionary learning algorithm considered in the present paper is highly stylized. However, such a process of groping for better forecasting rules may be a more realistic model of learning than least squares learning when agents face substantial model uncertainty. While the three models display broadly similar behavior, deviations from the MSV REE forecasting rule tend to be more persistent under the evolutionary algorithm than under least squares learning, but, at the same time, the learning dynamics tend to exhibit greater stability under this less directed process of forecast rule discovery than under least squares learning.

2. STRUCTURE OF THE MARKET

The structure of the market is the same as in Georges (2008b). There are two assets, a bond with a fixed return r in each period, and a stock that pays a dividend d_t in each period t. The average price at which the stock is traded within period t is given by P_t. The return on the stock in period $t+1$ is thus $\rho_{t+1} \equiv \frac{d_{t+1}+P_{t+1}-P_t}{P_t}$.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Thus, the return on the stock expected by agent i would be equal to the return on the bond if

$$P_t = \frac{F_t^i[P_{t+1} + d_{t+1}]}{1 + r}, \tag{1}$$

where $F_t^i[P_{t+1} + d_{t+1}]$ is agent i's forecast of $(P_{t+1} + d_{t+1})$ at time t. We will assume behaviorally that agents will desire to incrementally accumulate more of the stock (bond) in the current period if they expect the return on the stock in the next period to be greater (less) than the return on the bond.

Traders engage in bilateral over-the-counter trading. There is no Walrasian auctioneer and there are no market makers. In each trading period t, a randomly selected set of agents update their forecast rules through a process of experimentation and social learning. All agents then construct forecasts of the price and dividend for the next period and post limit orders based on these forecasts. Some agents then sequentially search for trades by sampling the orders posted by other agents and accepting trades that they find appealing. Individual trades may take place out of temporary equilibrium, and there need not be a single price in the market at any moment in time within a trading period. However, the average price at which trades are executed within a trading period becomes common information that is available in all future trading periods.

We will assume throughout the paper that dividends are given by $d_t = \bar{d} + \varepsilon_t$, where ε_t are i.i.d. with zero mean and finite variance.[2]

It will also be useful to define $x_t \equiv P_t + d_t$ and note that (1) can be written

$$x_t = \frac{F_t^i[x_{t+1}]}{1 + r} + d_t.$$

3. RATIONAL EXPECTATIONS EQUILIBRIUM

Given the dividend process assumed above, there is a unique stationary rational expectations equilibrium: $P_t = P^* \ \forall t$, where $P^* \equiv \bar{d}/r$. There is also a continuum of rational bubble equilibria, under which P_t explodes away from P^* over time at rate r, as well as nonstationary rational sunspot equilibria.

Note that even though dividends are stochastic, the stationary REE price is constant over time, so that *any* volatility observed in the simulations below is necessarily excess volatility.

Note also that, if we define $x^* \equiv P^* + \bar{d} = \frac{1+r}{r} \cdot \bar{d}$, then at the stationary REE, we have $x_t = x^* + \varepsilon_t \ \forall t$.

4. FORECAST RULES

Suppose now that agents do not have enough information about the structure of the environment and the other agents' beliefs to form rational expectations. Rather they formulate forecasts of future prices inductively.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Specifically, we suppose that agents use forecast rules of the form[3]

$$\mathrm{F}_t^i[x_{t+1}] = a_{0t}^i + a_{1t}^i x_t + a_{2t}^i \cdot x_{t-1} + a_{3t}^i \cdot x_{t-1}^2 + a_{4t}^i \cdot x_{t-2} + a_{5t}^i \cdot x_{t-2}^2, \quad (\mathbf{2})$$

where $a_{0t}^i, \ldots, a_{5t}^i$ are scalars that can vary across agents i and time t, and as noted above, P_{t-i} is the average price at which trades occurred in period $t-i$. We assume that neither the current period aggregate price P_t (which will emerge from trading during period t) nor the current period dividend d_t is in traders' information sets at time t. Therefore, traders must form iterated forecasts of x_{t+1} by first forecasting the current period's value x_t using the observed values from the preceding three periods.[4]

5. REPRESENTATIVE AGENT BENCHMARK

The quadratic terms in the forecast rules (2) can induce nonlinear dynamics in the price level even in the absence of learning. This can easily be seen if we consider a representative agent version of the model in which agents are homogeneous and prices are set by a Walrasian auctioneer. In that case, the price dynamics without learning would be characterized by (1) and (2). As shown in Appendix A, this system can exhibit a broad range of short-run and long-run dynamics. It is also clear in that case that the forecast rules (2) being used by the agents are misspecified.[5]

6. MINIMAL STATE VARIABLE FORECASTS

Note that the smallest number of state variables required to characterize a rational expectations equilibrium in this model is zero. That is, MSV forecast rules [McCallum (1983, 1999)] are rules of the form

$$\mathrm{F}_t^i[x_{t+1}] = a_{0t}^i \quad (\mathbf{3})$$

and the MVS REE forecast rule is $\mathrm{F}_t^i[x_{t+1}] = x^*$. Below, we will measure the distance of the average forecast rule in the population from the MSV REE rule as the Euclidean distance between the respective vectors of forecast rule parameter values:

$$\mathrm{DMSV} = |(\bar{a}_0, \bar{a}_1, \bar{a}_2, \bar{a}_3, \bar{a}_4, \bar{a}_5), (x^*, 0, 0, 0, 0, 0)|. \quad (\mathbf{4})$$

7. EXPECTATIONAL STABILITY

We can also note that the stationary rational expectations equilibrium $P_t = P^*$ is *expectationally stable* (E-stable) in the sense of Evans and Honkapohja (2001). Consider the representative agent benchmark above and restrict forecast rules to be MSV rules. Then under the notional continuous time learning algorithm employed in the E-stability framework, the representative agent's forecast parameter a_{0t} would converge locally to x^*. Since, for $a_0 = x^*$, $a_1 = a_2 = a_3 = a_4 = a_5 = 0$,

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

the temporary equilibrium, or "actual law of motion" (ALM), characterized by (1) and (2) is exactly the stationary rational expectations equilibrium, the E-stability framework selects that equilibrium as the unique outcome of learning in the MSV forecast rule case.

Further, this rational expectations equilibrium is *strongly expectationally stable* in that it is robust to linear overparameterization, as in the linear part of (2). Thus, in the linear case ($a_3 = a_5 = 0$), the E-stability criterion selects a particular learning equilibrium ($a_0 = x^*$, $a_1 = a_2 = a_4 = 0$) that replicates the stationary rational expectations equilibrium. Some details are given in Appendix B.[6] This result also implies that, at least in the linear case, the stationary rational expectations equilibrium would be selected locally under recursive least squares econometric learning with unbounded memory [Evans and Honkapohja (2001)].

The E-stability criterion has been shown to predict stability of REE under real-time learning more generally when agents have unbounded memory and use decreasing-gain learning algorithms such as recursive least squares learning. Since in the present paper, agents have finite memory, we would not expect learning to be complete. Honkapohja and Mitra (2003) have shown in simple cases similar to that studied here that if the conditions for *iterative E-stability* are satisfied (as they are here as well), then attempts to learn a stochastic steady state via least squares learning with finite memory are incomplete but do lead to asymptotically unbiased forecasts. In other words, the MSV forecast eventually becomes centered on the REE forecast. Further, the variance of the forecast error is decreasing in the length of agents' memory.

However, Honkapohja and Mitra do not consider overparameterized PLM, and they consider least squares learning by a representative agent in a market-clearing environment. Consequently, it is an open question what kinds of persistence we might expect from heterogeneous agents using non-MSV forecast rules and learning via an evolutionary algorithm (which is less directed than least squares learning) in a non-market-clearing environment. We might expect that under conditions of strong E-stability, learning would drive agents' perceived laws of motion (PLMs) to fluctuate around the MSV REE PLM. We do find this for the case of MSV forecast rules. However, for overparameterized rules, we find a stronger attraction to a weaker consistency condition than to the MSV REE rule. There is evidence of an attraction to the MSV REE rule at longer horizons, and that imitation is important in generating that attraction. However, even at long horizons, significant and persistent deviations from the MSV REE PLM generate persistent excess volatility and can lead to global instability.

8. WEAKLY CONSISTENT EXPECTATIONS

Finally, before continuing with the model, notice that there is a continuum of forecast rules of form (2) that are consistent with the stationary rational expectations equilibrium in the following sense: if $x_{t-i} = x^*$ for $i = 1, 2, 3$, then $F_t[x_{t+1}] = x^*$.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

This *weak consistency condition* (WCC) can be written

$$x^* = a_0 + (a_1 + a_2 + a_4) \cdot x^* + (a_3 + a_5) \cdot x^{*2}. \quad \textbf{(5)}$$

Note that (5) is linear in the six forecast rule parameters a_0, a_1, ... and so the values of these parameters that satisfy it lie on a five-dimensional hyperplane. Note also that the MSV REE forecast rule satisfies the WCC.

The WCC can be interpreted as a steady state consistency condition. Consider the representative agent benchmark with no dividend shocks. If the price level is in a steady state equilibrium, and the representative agent's expectations are correct, (5) must be satisfied. That is, the only expectations that can be confirmed by the actual path of prices in steady state [under the deterministic part of system (1)–(2)] are given by forecast rules that satisfy (5).

Further, if the forecast rule parameters satisfy (5), then the steady state [under the deterministic part of (1)–(2)] is the stationary REE $P_t = P^*$, so that if P is initially at this level, it will remain there over time in the absence of dividend shocks. Note, however, that this steady state may or may not be locally asymptotically stable, and that (the deterministic part of) the system (1), (2), (5) admits a rich variety of non-steady state dynamics, even in the absence of learning (again, see Appendix A). Thus, out of steady state, expectations need not be correct, even on average. Consequently, the WCC is weaker than consistency conditions that require that expectations be unbiased, as for example in Hommes and Sorger (1998). The condition is, however, a necessary condition for expectations to be unbiased in the linear case ($a_3 = a_5 = 0$).[7]

Below, we will measure the deviation of the average forecast rule from the WCC as the difference between the actual average value $\bar{a}_{0t}$ and the value that would satisfy (5) given the actual average values of the other parameters ($\bar{a}_{1t}$, $\bar{a}_{2t}$, $\bar{a}_{3t}$, $\bar{a}_{4t}$, $\bar{a}_{5t}$). The formula is

$$\text{DWCC} = |\bar{a}_0 - (1 - \bar{a}_1 - \bar{a}_2 - \bar{a}_4) \cdot x^* + (\bar{a}_3 + \bar{a}_5) \cdot (x^*)^2|. \quad \textbf{(6)}$$

As in Georges (2006), we find below that, in a variety of simulations, there is a strong tendency for agents' forecast parameters to be drawn toward the weak consistency condition (5). As noted above, the tendency to be drawn toward the MSV REE forecast rule appears to be substantially weaker. Consequently, the system can display a variety of stability configurations on a variety of time scales.

9. TRADING

It is common, in the literature on adaptation by heterogeneous agents in financial markets, to bypass the modeling of individual trades by imposing a Walrasian auctioneer to set price in order to clear the market. This is true for artificial agent models [Arthur et al. (1997), LeBaron et al. (1999)] as well as analytical models [Brock and Hommes (1998), Chiarella and He (2002), Brock et al. (2005)]. An alternative but similar approach is to assume a mechanical partial adjustment of

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

prices out of equilibrium [Lux and Marchesi (1999), Chen and Yeh (2001, 2002)]. Market makers who are willing to take positions in the market to absorb the excess demands of other agents can also be introduced while still modeling price adjustment as depending on net orders rather than individual trades, as in Lux (1995), Farmer and Joshi (2002), or Alfarano and Lux (2007). Finally, there is a growing agent-based modeling literature in which the market microstructure is modeled explicitly [Gode and Sunder (1993), Chan et al. (1999), Chakrabarti (2000), Audet et al. (2002), Chiarella and Iori (2002), Yang (2002), Yeh (2003), Chiarella and Iori (2004), Farmer et al. (2005), LeBaron and Yamamoto (2006), and Darley and Outkin (2007)]. In this paper, we follow this line in the literature and model individual trades explicitly under a stylized microstructure.

Trading occurs as follows. In any period, each agent i uses her current forecast rule to generate her forecast $F_t^i[x_{t+1}]$ of the price plus dividend in the next period. She then uses this forecast to set her current bid and ask prices. These prices are set to create a fixed spread around the reservation price that would make the agent indifferent to holding the stock or the bond given her forecast [i.e., the price satisfying (1)]. Agents' desired order quantities are simply fixed fractions of their current holdings of the assets to be sold. Thus, each agent is prepared to shift her portfolio incrementally in the direction of the asset that she believes will have the greater return in the subsequent trading period. Once all limit orders are posted, a set of agents are randomly selected to be active traders. These agents then sequentially search for trades by randomly sampling the order quotes of other traders and executing all trades that they find appealing.[8] When all trades are exhausted, the average price for the period is recorded and becomes public knowledge.

10. LEARNING

We model the evolution of forecast rules as a process of social learning and experimentation using an evolutionary algorithm. Our algorithm is essentially a real coded genetic algorithm modified to serve as a metaphor for social learning in this market context.[9] Agents are long-lived and are able to experiment occasionally with alternative forecast rules and adopt the rule among these alternatives that would have had the smallest forecast error in recent periods. In this way, they grope for rules with better fits to the recent data.

Each agent is endowed with initial real values $a_{0_0}^i, a_{1_0}^i, \ldots, a_{5_0}^i$ of the forecast rule parameters, which are randomly distributed around a set of common values. In each period 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 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

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

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 her own previous rule and may adopt some elements of the comparison agent's rule rather than imitating the rule in its entirety.

More specifically, in a given period, each agent considers updating her strategy with probability *preproduce*, in which case she considers imitating a comparison agent's strategy with probability *pcompare* and conducts crossover with that strategy with probability *pcross*. Crossover results in new forecast rule parameter (gene) values that are affine combinations of the two agents' parameter values. The crossover factor (the weight on the imitating agent's current parameter values) is selected from a uniform distribution on $(0.5 - C, 0.5 + C)$, where $C > 0.5$. Each gene j of the resulting second comparison chromosome is mutated with probability *pmutate* by an additive mutation factor M_j, which is chosen independently from a uniform distribution on $(-M, M)$. There are two supports for this distribution (i.e., two possible values of M), one relatively small and one relatively large, that each agent can adopt in any period. Thus, each agent can experiment either very locally or more broadly at any time. If crossover does not occur, then the original rule is mutated in place of the crossed-over rule.

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 the sum of the absolute forecast errors that the rule would have generated in the past *leval* rounds ("*leval*" stands for the length of the evaluation period):[10]

$$f_t^i = -\Sigma_{s=t-leval}^{t-1} \left| F_t^i[x_s] - x_s \right|. \tag{7}$$

Thus, in period 5, if *leval*=1, the fitness of a rule would depend on how well it would have forecast x_4 given x_2, x_1 and x_0.

We could have alternatively assumed that each trader performs a directed search and accurately fits (2) to the recent data in order to maximize (7) or similarly performs least squares on the last *leval* periods of data. This could be easily accomplished by sophisticated agents in our simple market environment. However, we are treating our less directed search algorithm (a process of individual groping through experimentation and imitation) as a metaphor for learning in more realistic economic environments for which exhaustive search is likely to be impossible.[11] Further, if all agents updated their rules synchronously in the same fashion, there would be no diversity of beliefs, and consequently trading would never occur.[12] Staggered updating on the other hand would permit belief heterogeneity under directed search and could be an alternative to the approach taken here. A companion paper (Georges, 2008b) shows that this alternative approach (staggered least squares learning) yields qualitative outcomes that are similar to those reported below, though with greater instability in the learning dynamics. We will return to the stabilizing properties of social learning in Section 16.

It is also worth emphasizing that in our simulations (following the spirit of Arifovic's election operator), a mutation is only adopted if it yields greater fitness

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

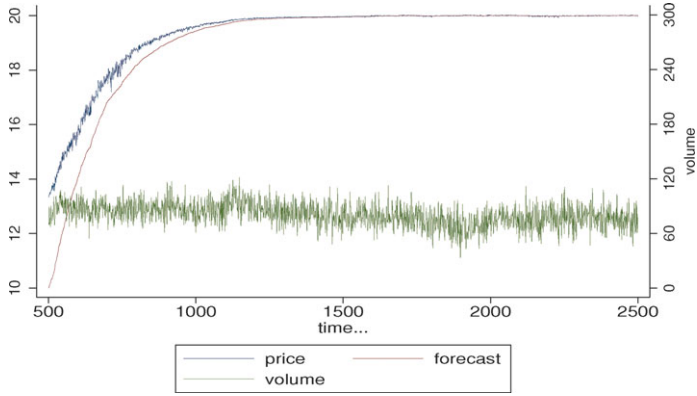


FIGURE 1. The evolution of the state variable P_t and average forecast of P_{t+1} (measured as the average parameter value $\bar{a}_t$ minus $\bar{d}$) over 2,000 rounds with $a_1 = a_2 = a_3 = a_4 = a_5 = 0$ $P^* = 20$ ($\bar{d} = 10, r = 0.5$), and 200 agents whose initial values of a_0^i are dispersed around 20. *Note*: Trading volume is also shown and is measured on the right-hand scale. (The figures in this article can be viewed in color at http://journals.cambridge.org/jid_MDY.)

than the agent's incumbent rule. Thus, there is not ongoing noise in the selection process as in traditional genetic algorithms or standard discrete choice and stochastic stability approaches to evolutionary dynamics in games and financial market modeling.

11. BENCHMARK CASE: ALL AR TERMS SUPPRESSED

As noted above, MSV forecast rules in this model take the form $F_t^i[x_{t+1}] = a_{0t}^i$, which corresponds to a belief that x_t will be constant, or fluctuate unpredictably about a constant, in each period. Note that for any constant ${a_0}^i \neq x^*$, if this belief were confirmed over time by $x_t = {a_0}^i \ \forall t$, then agent i would view the stock as being mispriced (overvalued if ${a_0}^i > x^*$, undervalued if ${a_0}^i < x^*$).[13] Consequently, the only constant belief that can be confirmed at equilibrium will be ${a_0}^i = x^*$.

It is not entirely surprising then that in simulations holding $a_1 = a_2 = a_3 = a_4 = a_5 = 0 \ \forall t$, we find that for each agent, learning tends to drive a_{0t}^i to become centered on x^* over time and that consequently, P_t tends to be drawn toward and become centered on P^* over time. Figure 1 is an example of a run for this case with 200 agents.[14]

We can see in this simulation that, even with very short memory (*leval* = 1), the combination of active trading and social learning drives the average price and the implied average expectation of price (measured as $\bar{F}_t[x_{t+1}] - \bar{d}$) toward $P^* = 20$.[15]

Because agents look back only a fixed number of periods (one period, in the current example) when evaluating trading rules, learning is not complete. Dividend

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

shocks create persistent diversity in expectations and thus persistent trading and fluctuations in price, which continue for the duration of the 50,000 period run from which Figure 1 is taken. However, after a transitory period, forecasts become roughly correct on average and the price level becomes roughly centered on P^*. The mean and standard deviation of the price level for periods 10,000–20,000 in this run, for example, are approximately 19.9963 and 0.0112, respectively, and the average forecast error for that period is approximately -0.0014.

Thus, when all agents use MSV forecast rules, they tend to learn a noisy version of the MSV REE forecast rule, which is also the unique rule that satisfies the WCC. This result accords with the results of Honkapohja and Mitra (2003) for least squares learning by a representative agent with bounded memory. Thus, even though our social learning algorithm is less directed than least squares learning, our agents are heterogeneous, and trading takes place out of equilibrium, the model produces results consistent with Honkapohja and Mitra (2003) when agents are restricted to using MSV forecast rules.

12. GENERAL CASE I: a_5 ACTIVE

Adding in variation in the parameters a_1, a_2, a_3, a_4, and a_5 gives agents additional parameters to learn. We find that trading and learning typically drive both the average market price and implied individual price forecasts to be approximately centered on $P = P^*$ but can also lead to the formation of explosive bubbles.

We generally find that there is a strong attraction of individual forecasting rules to the WCC. However, because there is at the same time substantial persistence of heterogeneity within the population, excess volatility tends to persist. Further, this volatility can be clustered due to the interaction of the learning dynamics and the heterogeneity in forecast rules.[16] Interestingly, deviations from the MSV REE forecast rule tend to be persistent both for individuals and for the average rule.

An instructive example can be given by activating only two of the six forecast rule parameters. In the following run, agents start out far from the MSV REE rule with a_0^i randomly distributed around 12 and a_5^i randomly distributed around 0.02. We first consider the case where learning is restricted to individual learning via mutation, and social learning (imitation and crossover) is suppressed. Memory (*leval*) is set to five periods and mutation rates are reduced for the a_5 coefficient. Otherwise, the model parameters are as in Figure 1.[17] In Figure 2, we see that after 10,000 rounds, there has been a strong attraction to the WCC (which is given here by $a_0 = 30 - 900\,a_5$) but essentially no attraction to the MSV REE, which corresponds to the specific point $((a_0, a_5) = (30, 0))$ on the WCC locus.

This result can also be seen by our measurements of the distance of the average rule from the WCC and from the MSV REE rule (DWCC and DMSV, respectively) period by period for the first 2,000 rounds; see Figure 3.

We see here that the average forecast rule (across agents) is attracted quickly to a neighborhood of the WCC, but not necessarily to the MSV. In the left-hand panel, which corresponds to the simulation in Figure 2, the average rule stays near

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

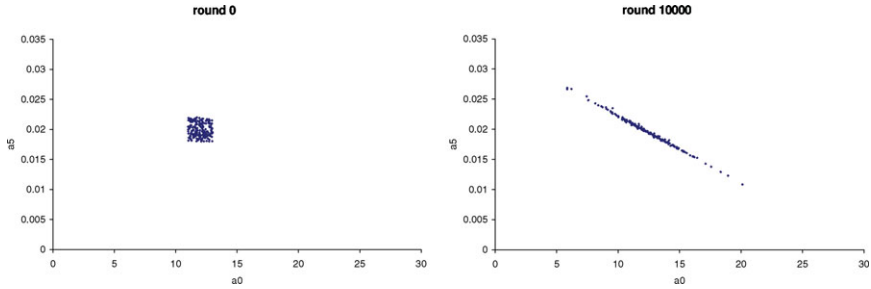


FIGURE 2. The attraction of individual agents to the WCC. The left-hand panel shows the initial dispersion of agents' parameter values a_0 and a_5. The right-hand panel shows the dispersion after 10,000 rounds of trading and individual learning.

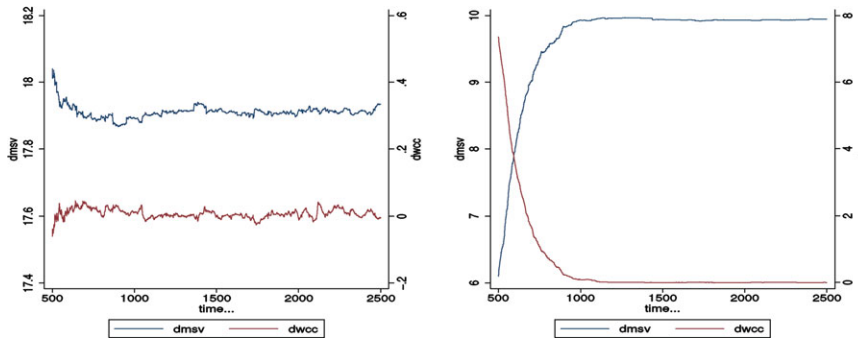


FIGURE 3. The attraction of the average forecast rule to the WCC. The left-hand panel gives the first 2,000 rounds of the simulation in Figure 2. DMSV is measured on the left-hand scale and DWCC on the right-hand scale. Memory (*leval*) is five periods. The right-hand panel shows a different simulation, which is identical to the first except that agents start further from the WCC with initial parameter values distributed around $((a_0, a_5) = (24, 0.015))$.

the WCC and far from the MSV REE rule. In the right-hand panel, agents start further from the WCC but move rapidly toward it in early rounds while diverging further from the MSV REE rule.

There is also continued wandering of both individual agents' rules and the average rule over time due to continued exploration, see Figure 4[18]

Finally, we can see that agents' willingness to deviate from MSV forecast rules yields interesting price dynamics.

The average price per period wanders near the REE price $P^* = 20$ and the log growth rate of the price level displays clustered volatility and modest excess kurtosis (kurtosis is approximately 4.84 for the data displayed in Figure 5).[19]

So we see in this example a clear attraction to the WCC. Any attraction to the MSV REE rule appears to be much slower than the attraction to the WCC. We see

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

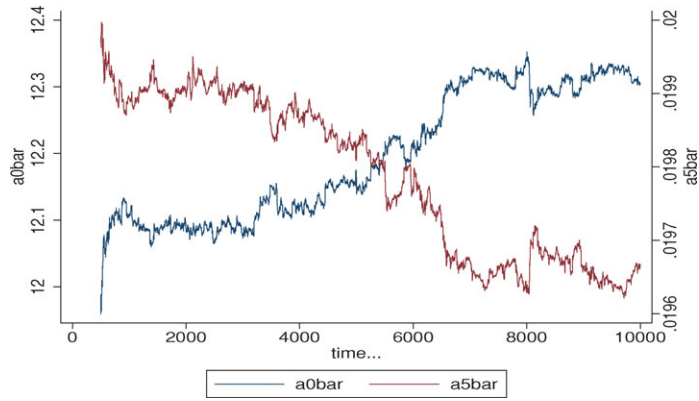


FIGURE 4. The evolution of the average forecast rule parameters $\bar{a}_{0t}$ (left-hand scale) and $\bar{a}_{5t}$ (right-hand scale) over 10,000 periods. Same run as in Figure 2.

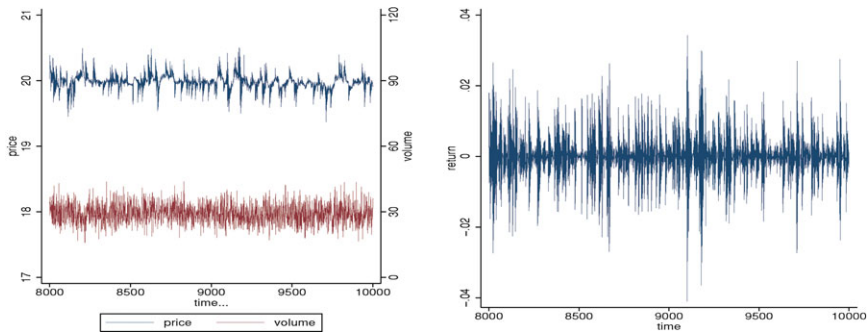


FIGURE 5. The evolution of P_t and volume (left-hand panel) and returns (right-hand panel). $P^* = 20$. Shown are rounds 8,000–10,000. Same run as in Figure 2. Returns are net of dividends and are calculated as the log growth rates of the average price.

in Figure 4, that there is a slow drifting of a_{0t} upward matched by a slow drifting of a_{5t} downward in the first 10,000 trading periods. Continuing the simulation for another 40,000 periods, we find that this slow drifting continues. As agents wander near the WCC in reaction to the dividend shocks and the endogenous price dynamics, there is a weak tendency to move toward rather than away from the MSV. Nevertheless, in 50,000 trading periods, the effect is still modest. By round 50,000, $\bar{a}_0$ has risen to approximately 13.4, $\bar{a}_5$ has fallen to approximately 0.0185, and DMSV has fallen to approximately 16.6.

Consider now the following variations on the simulation above. First, with very short memory ($leval = 1$) and only individual (not social) learning, agents tend to spread out along the WCC increasingly over time. Eventually, this leads to heightened volatility and some dispersion of agents away from the WCC. In some cases, this causes the population to drift far from the MSV over time.[20] However,

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

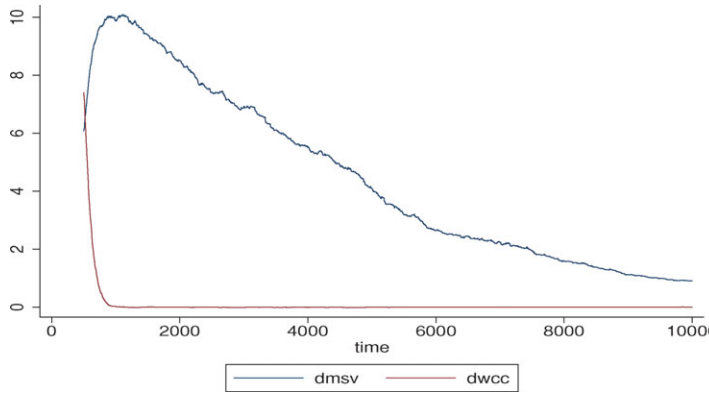


FIGURE 6. The attraction of the average forecast rule to the WCC and MSV. Same as the RH panel of Figure 3, except that *pcompare* has been increased from 0 to 0.01 and *pcross* from 0 to 0.2 in this run.

attraction to the MSV tends to be intensified as memory (*leval*) is increased and as social learning is introduced. The latter effect appears substantially stronger than the former—the ability to imitate others (completely, or partially via crossover) improves agents' abilities to discover and adopt relatively high fitness rules near the MSV REE rule. In either case, this drift appears to generally take place on a longer time scale than the attraction to the WCC. In some cases, convergence to a noisy MSV equilibrium is quite rapid, and within a few thousand trading periods $\bar{a}_5$ is fluctuating around zero. Even with memory of one, introducing social learning frequently drives the population to a neighborhood of the MSV over time. Thus, in many cases, the homogenizing influence of social learning appears to be stabilizing in the long run.[21] In other cases, however, the process is not complete and a_5 continues to wander far from 0.

To illustrate the effect of social learning, consider the same simulation as in Figures 2–5, but add imitation and crossover. Here we set $pcompare = 0.01$ and $pcross = 0.2$.[22] As seen in Figures 6 and 7, the upshot is still a rapid movement of the population toward the WCC early on, but now there is also a slower but persistent movement along the WCC toward a smaller neighborhood of the MSV REE rule.

We can see that after 10,000 trading periods, the population has drifted to the right along the WCC toward the MSV REE forecast rule. This drift is due to imitation, which appears to favor rules of comparison agents that are relatively close to the MSV.

As we increase the rate of social learning in this simulation (by increasing *pcompare*), the movement toward the MSV becomes more rapid. However, this attraction to the MSV is sensitive to the initial conditions. For example, if the initial rules are distributed around $(a_0, a_5) = (10, 0.015)$ rather than $(24, 0.015)$, then DWCC approaches zero in the first 1,000 rounds, whereas DMSV takes

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

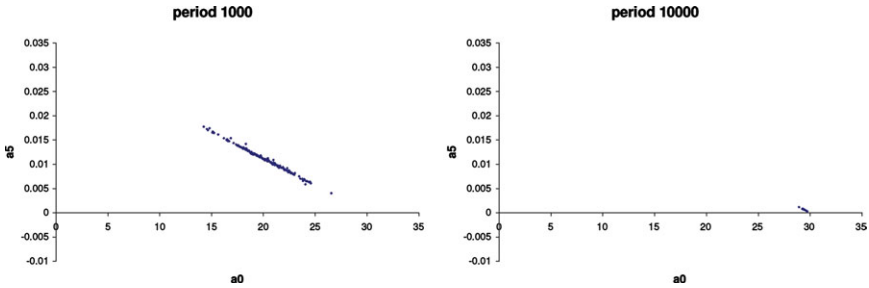


FIGURE 7. The attraction of individual agents to the WCC and MSV. Same simulation as in Figure 6. The left-hand panel shows the dispersion of agents' parameter values a_0 and a_5 after 1,000 rounds of trading and social learning. The right-hand panel shows the dispersion after 10,000 rounds.

10,000 rounds to approach zero even under rapid social learning ($pcompare = 0.1$). Further, in a number of simulations with other initial conditions, convergence to a noisy MSV fails altogether, and DMSV reaches a plateau far from zero.

Thus, to summarize these findings, social learning often causes (a_0, a_5) to be drawn into a noisy MSV equilibrium as predicted by Honkapohja and Mitra (2003). However, this tends to be on a longer time scale than the attraction to the WCC. Further, for some runs, this process is incomplete or diverges altogether.[23] Our preliminary interpretation is that a noisy MSV REE can be a learning equilibrium of our model, but convergence to it is sensitive to parameters of the model and to initial conditions.

This sensitivity of the long-run behavior of the system under learning appears to be due to a near equivalence of forecasting power for rules near the WCC. The true global optimum [the rule that maximizes (7)] is roughly centered on the MSV REE rule (a_0, a_5) = (x^*, 0), but this optimum is noisy, fluctuating broadly near the WCC, and thus presenting a moving target to agents learning with bounded memory. Additional analysis is provided in the next section.[24]

13. FURTHER ANALYSIS OF THE TWO-PARAMETER CASE

To understand the forces operating in the two-parameter case (a_0 and a_5 active), consider the following experiment. We generate a random series $x_t = x^* + \nu_t$, where $x^* = 30$, and the ν_t are i.i.d. uniformly on (−0.5,0.5). We then plot antifitness landscapes in (a_0, a_5) space. That is, for any time in the artificial history, we can plot the sum of absolute forecast errors (7) for various combinations of the forecast rule parameters. Figure 8 is a representative plot.

The landscape will shift and deform over time as the price history evolves. However, this diagram is representative in that we tend to see a strong trough in which antifitness is relatively low (forecasts are relatively accurate). The trough shifts across periods, but is centered, at least roughly, on the WCC over time.[25] For $leval = 1$, the variation is large, but this variation tends to fall as *leval* grows.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

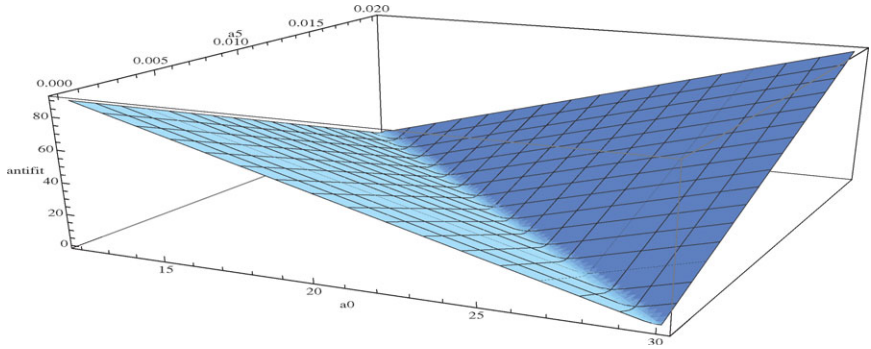


FIGURE 8. Antifitness $(-f)$ for a particular hypothetical price and dividend history as a function of the forecast rule parameters a_0 and a_5. Here, *leval* is set to five periods. The front and rear corners of the domain pictured here [(30,0) and (12,0.02)] both satisfy the WCC. This realization of the fitness function is typical in that it exhibits a trough that lies near the WCC.

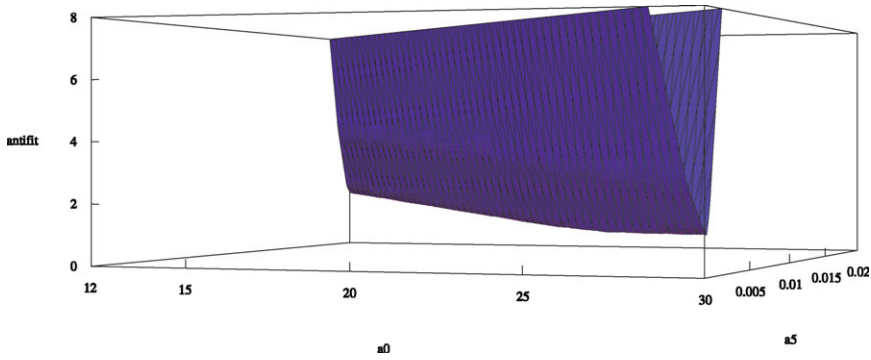


FIGURE 9. Antifitness for a particular hypothetical price and dividend history as a function of the forecast rule parameters a_0 and a_5. Here we focus on the shape of the trough. The lowest point on the trough changes over time. *leval* is set to five periods. $x^* = 30$. The MSV REE rule $(a_0, a_5) = (x^*, 0)$ is in the front right corner.

Further, the valley of the trough is relatively flat, and the minimum value along the trough varies over time (and may exhibit multiple local minima). For $leval = 1$, the trough is perfectly flat—i.e., the global optimum is indeterminate—whereas for $leval > 1$, the global minimum is generally determinate. As memory grows, the variation of the minimum over time also tends to be reduced. Figure 9 is a representative view of the trough with the MSV REE rule again in the front right corner.

For this case of randomly generated data, the relatively flat valley and steep walls of the trough makes this optimization problem (minimizing the forecast error) ill-conditioned. It is not surprising, then, that in our simulations agents

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

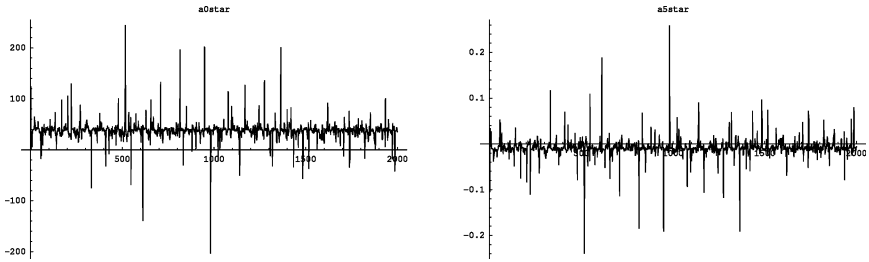


FIGURE 10. Actual optimal values of the forecast rule parameters a_0 and a_5 for the simulation period displayed in Figure 5. Recall that the MSV REE rule $(a_0, a_5) = (x^*, 0) = (30,0)$ and $leval = 5$. The mean and standard deviation of x_t for the period are approximately 29.9779 and 0.1501. The optimal values of (a_0, a_5) have mean (37.4189, −0.00827738) and standard deviation (16.8136, 0.0187044).

rapidly discover the valley but only slowly adjust within the valley. We can also see that social learning (imitation) can make the latter adjustment process much more rapid (though it may also cause premature convergence to spurious metastable learning equilibria).[26] Further, because the global minimum shifts substantially over time as the time series $\{x_t\}$ evolves, there will be ongoing learning of a moving target.

We can take this analysis a step further by calculating the global optima over time for actual data generated by our simulations. When we do this, we discover that the center of the distribution tends to lie near the MSV REE rule, and that the variance is large even for moderately high memory. To continue with the example in Figures 2–5, we calculated the global optimal forecast rules for trading periods 8,000–10,000 (the time period reported in Figure 5), using the values of $x_t \equiv P_t + d_t$ generated in that simulation. Recall that $leval = 5$. The mean and standard deviation of x_t for the period are approximately 29.9779 and 0.1501. The optimal values of (a_0, a_5) have mean (37.4189, −0.00827738) and standard deviation (16.8136, 0.0187044). We see that there is substantial variation in these optima, and that the average values fall well within a standard deviation of the MSV REE values. Further, although there are large transitory deviations of the optimal rules from the MSV rule, there is a strong clustering of these rules near the WCC. See Figures 10 and 11.

Thus, the strong persistence of the center of the distribution of agents' rules displayed in Figure 2 would appear to be due to a combination of the strong relative fitness of rules near the WCC, the relative insensitivity of fitness to movement along the WCC, the variability of the true optimum, and the limitations of boundedly rational learning. If agents were able to come closer to optimizing fully in each period, they would chase the true optimum more closely over time, and rules would be both more variable and centered more closely on the MSV REE. We will return to this observation after reporting on the full model.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

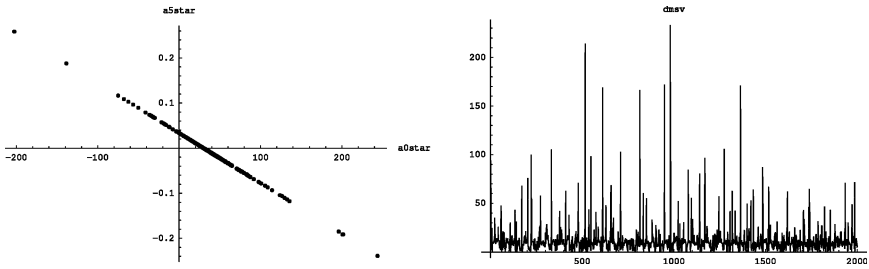


FIGURE 11. Scatter of the optimal values of (a_0, a_5) displayed in Figure 10 and levels of DMSV (distance from the MSV REE rule) for these values. We see substantial clustering near the WCC.

14. GENERAL CASE II: ALL COEFFICIENTS ACTIVE

Now consider the representative run with all coefficients of the forecast rule active shown in Figure 12. Here $\bar{d} = r = 0.05$, which implies that $P^* = 1.0$ and $x^* = 1.05$. Agents start with rule parameters distributed around the MSV REE rule (1.05,0,0,0,0,0). Memory (*leval*) is set to 50 trading periods. *preproduce* = 0.5, *pcompare* = 0.1, *pcross* = 0.2, and *pmutate* = 0.2. Dividend shocks are distributed uniformly on (−0.025, 0.025) and occur with probability 0.8 in any period. Mutations for each forecast rule parameter (gene) are selected from uniform distributions on (−0.05, 0.05) and (−0.002, 0.002). As above, 20 traders are active in each period and each looks at the prices of 4 others.

As in the simulations shown in Figures 2–11, price fluctuates around its constant REE value (which is here 1.0), and the returns exhibit volatility clustering and excess kurtosis (kurtosis is 8.56 for the trading period 10,000–50,000). We see that, although DWCC remains close to zero throughout the 50,000-round simulation, DMSV does not; although the average forecast rule wanders in a broad neighborhood of the MSV REE rule, the attraction to the WCC is substantially stronger. In complementary simulations in which the initial rules are far from both the MSV REE rule and the WCC, rules converge rapidly to a narrow neighborhood of the WCC while again wandering persistently in a substantially broader neighborhood of the MSV REE rule. These qualitative results are fairly robust, though the specific pattern of rule evolution is sensitive to both initial conditions and the parameters of the evolutionary algorithm.

We see in Figure 12 that deviations from the MSV REE rule, although perhaps not growing extreme, can be substantial and persistent. Whereas, if agents learned to forecast using only MSV forecast rules, the average rule would fluctuate closely around the MSV REE rule, with overparameterized rules the average forecast rule can display persistent deviations from the MSV REE rule without generating persistent forecast errors.

Variations in this simulation suggest that the variance of returns appears to be generally falling in memory (*leval*) and rising in the rate of learning (*preproduce*).[27]

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

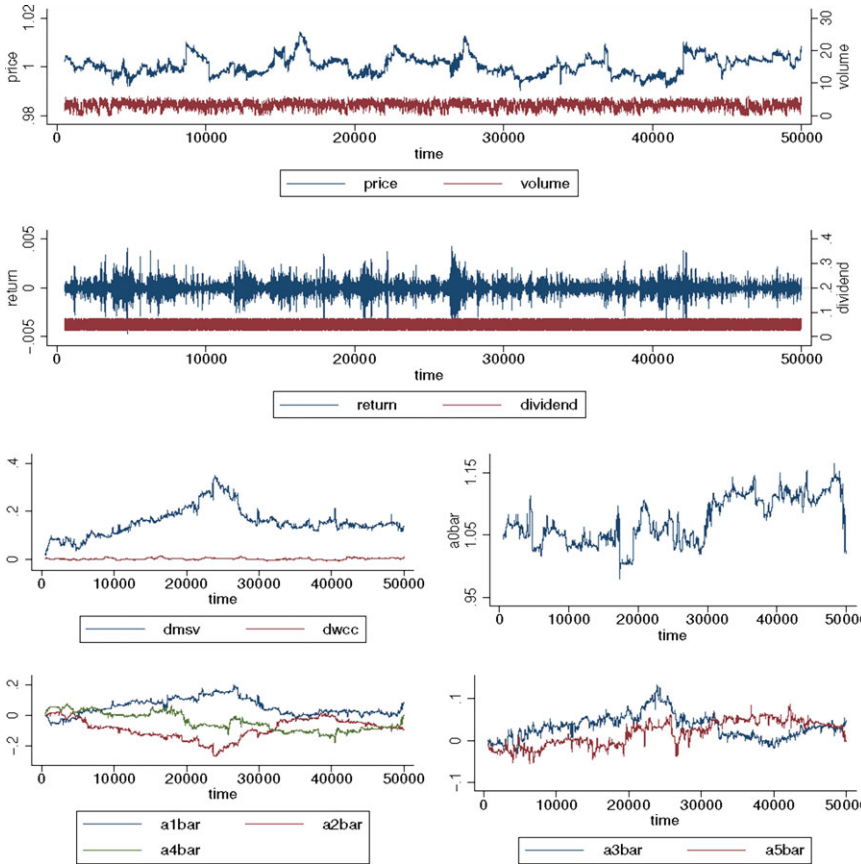


FIGURE 12. Run with all forecast rule parameters active. The initial values $(a_{0_0}^i, a_{1_0}^i, a_{2_0}^i, a_{3_0}^i, a_{4_0}^i, a_{5_0}^i)$ are distributed around the MSV REE rule (1.05,0,0,0,0,0).

The average deviation of agents' rules from the MSV REE rule appears to tend to be increasing in the larger mutation factor—i.e., in the breadth of search.

15. DIVERSITY

We have been focusing on the average forecasting rules. However, in some respects, the diversity of rules is more important than the average.

If we make agents a bit smarter, so that they expect the dynamics of x_t to be centered on x^* and forecast deviations $(x_t - x^*)$ based on recent lagged deviations, then we find that the population tends to maintain greater diversity under social learning. We also find an increased incidence of large bubbles or crashes—i.e., a greater tendency for the price dynamics to become unstable, more pointedly illustrating Grandmont's (1998) "uncertainty principle."

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

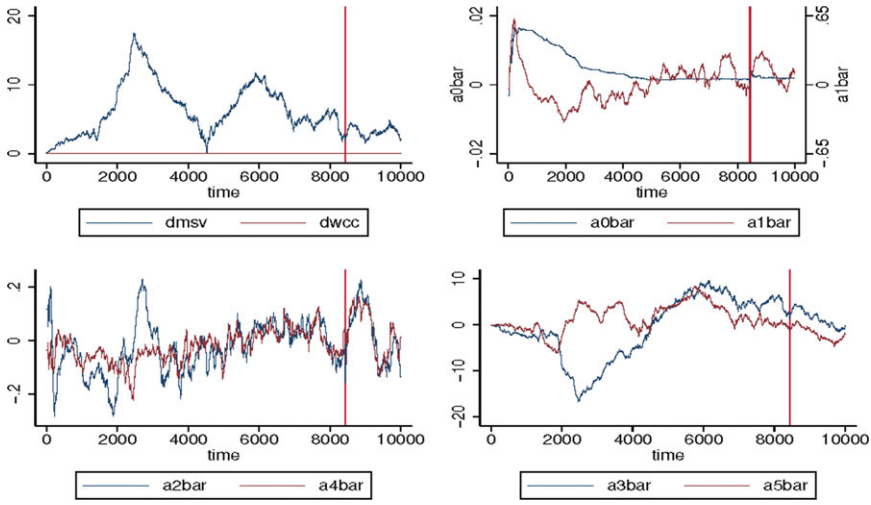


FIGURE 13. Average forecast rule parameters for a run in which agents forecast differences from x^*. $leval = 10$. There is a large burst of volatility starting in trading period 8,434 (that period is indicated in each panel by a vertical line).

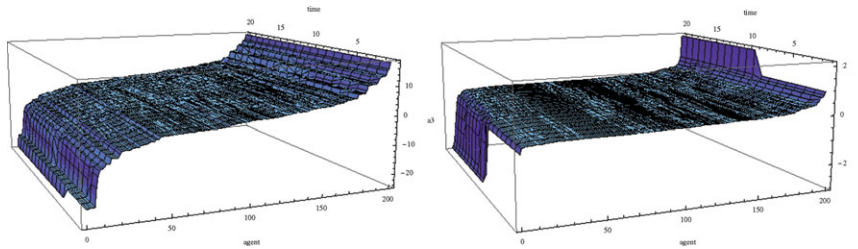


FIGURE 14. a_3 (left-hand panel) and a_4 (right-hand panel) for each of the 200 agents over 20 periods. Here agents are sorted each period by a_3 (left-hand panel) and a_4 (right-hand panel). A burst of volatility begins in period 8,434, which is time 9 in this figure.

In Figure 13, an episode of heightened price volatility occurs beginning at trading period 8,434, apparently triggered by an isolated dividend shock in period 8,433. There is no apparent movement in the average forecast rules around that time that would suggest an increased likelihood of volatility.

There is, however, substantial diversity of forecast rules in the population, and a further spreading of the distributions of several of the forecast rule parameters (a_1, a_2, a_3, a_4) occurs just prior to the episode, as can be seen in Figure 14.

Agents in the outer tails are active traders during the event, and so the discovery of more extreme forecast rules may play an important role in generating the heightened volatility. We also observe some herding as a few agents follow others into the tails, reinforcing this effect. Interestingly, however, traders in the outer

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

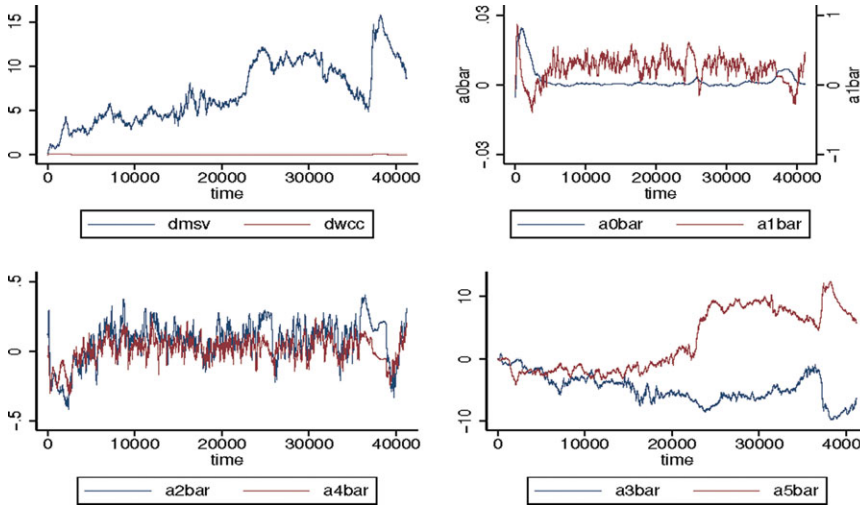


FIGURE 15. A crash occurs starting in period 41,172. Here, average rule parameter values are shown through period 41,185, at which point the crash is in progress.

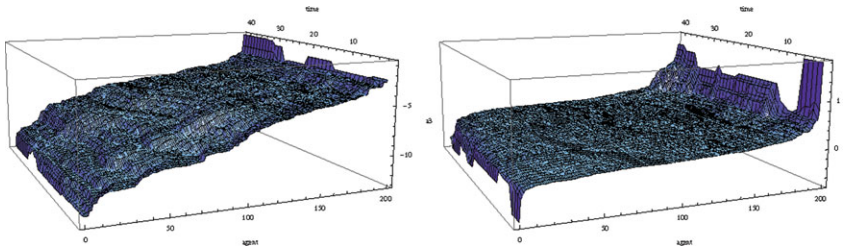


FIGURE 16. a_3 (left-hand panel) and a_4 (right-hand panel) for each of the 200 agents over 41 observations, each 5 periods apart. A crash occurs starting in period 41,172, which is time 34 in this figure.

tails neither initiate the event nor dominate trading during it. There is broad participation in trading across the distribution of agents' rules during the event.

Similarly, in the run in Figures 15 and 16, there is a crash in the price level beginning in period 41,172 (price falls from approximately $P^* = 1.0$ to 0.15 over the course of the following 36 trading periods and then explodes toward $+\infty$).

Though the average forecast rule at the onset of the crash is far from the MSV rule (0,0,0,0,0,0), this rule is not extreme relative to previous trading rounds, and in this case it is not clear that there is any dramatic spreading of the distribution of the rules across agents just prior to the crash. However, it appears that this distribution has evolved to be diverse enough so that a serendipitous combination of dividend shocks and specific traders executing orders will set off a bubble.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

16. OBSERVATIONS

In our model of social learning with overparameterized forecast rules, there appears to be a strong attraction of the forecast rule parameters toward the WCC because of learning. At the same time, there appears to be a weaker attraction to the MSV REE rule. The system can remain far from the MSV equilibrium over long periods of time, in terms of both individual forecast rules and average rules, generating nontrivial disequilibrium dynamics. Ongoing heterogeneity in the population far from the MSV REE rule tends to generate persistent and interesting price dynamics even in this simple toy model environment[28] and can evolve to states in which extreme events may be triggered.

Interesting price dynamics, including volatility clustering, fat tails in returns, and extreme events, are by necessity driven in this model by heterogeneity in the population. In turn, that heterogeneity is promoted by large mutation factors (broad rather than local search), limited memory, and relatively large bid–ask spreads (i.e., greater granularity in price movements).

Herding, in the form of social learning (imitation), appears to be alternately stabilizing and destabilizing in this model. Since social learning allows much more rapid learning than individual learning, it tends broadly to pull the distribution of agent rules toward the MSV REE rule (as seen for example in Figures 6 and 7). However, social learning can also allow agents with moderate rules to adopt the extreme rules of other agents, causing spreading in the distribution and thus conditions more favorable to extreme events such as large bursts of volatility, bubbles, and crashes.[29]

Interesting price dynamics are not due exclusively to the nonlinear terms in the forecast rule. For example, some volatility clustering arises even in the simplest MSV cases. Because trading is activated by heterogeneous expectations, dividend shocks occasionally set off a cycle of finite memory learning, expectations dispersion, trading, and price fluctuations. However, as additional linear and quadratic autoregressive terms are added to the forecast rule, the variation over time in the optimal forecast rule increases, leading to an increase in the variation of the actual forecast rules that are discovered and adopted by our boundedly rational agents. This leads to an increase in return volatility, volatility clustering, and the incidence of extreme events. By adopting the functional form (2), agents are essentially overfitting the returns data that they have access to, leading to increased variation in the forecast rule parameters.

Some readers might object to this and argue that, given the simple environment that they operate in, our agents should be more sophisticated in testing this specification. However, in realistic environments, exhaustive specification searches are likely to be infeasible and so forecasters will be required to fall back on heuristic devices.[30] A second possible objection is that the learning algorithm used by our agents to fit (2) to the data is not efficient enough. The tendency for there to be rapid convergence toward the WCC suggests that rapid learning can and does take place under our learning algorithm. However, once the easy work of discovering

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

the antifitness valley of the WCC is accomplished, individual random experimentation (mutation) becomes a blunt tool for finding better forecast rules along that valley. This rapid learning far from equilibrium and sluggish learning near equilibrium can strike a balance that allows diversity to persist in the population, but also allows actual forecast rules to be far from the optimal rule at any time.[31]

Note, however, from Section 13 above that the optimal rule is itself highly variable. Groping and imitation tends to keep adopted rules closer to the MSV REE rule than they might be if agents were more directed in their search but still updated their rules at random intervals. Thus, our undirected learning algorithm may actually be stabilizing in comparison to some forms of more directed search.

Indeed, comparison of our model with a similar model with agents engaged in a more directed process of least squares learning [Georges (2008b)] indicates that social learning leads to greater persistence of deviations from the MSV REE forecast rule, but also greater stability of the resulting price dynamics, in the sense that least squares learning tends to generate a substantially higher incidence of explosive bubbles in this model. The difference is particularly dramatic when forecasting is in levels, rather than deviations from the REE, and when mutation factors in the social learning algorithm are relatively small, i.e., when individual experimentation is relatively local. Groping, by its very inefficiency, tends to prevent agents from discovering and herding on extreme rules with temporarily high fitness that can set off bubbles, whereas least squares learning allows agents to find such rules seamlessly when they arise. Thus, the combination of experimentation and imitation tends to lead agents to wander among the relatively robust rules in a tighter neighborhood of the MSV REE rule, resulting in greater stability of the system.[32] Under both learning models, however, the incidence of explosive bubbles is increasing in the rate of learning and decreasing in memory.

Finally, Georges (2002, 2008a) indicate that the general qualitative results reported here are not merely an artifact of the specific microstructure embedded in our model, but rather also are exhibited when the trading mechanism is replaced by a market-clearing condition that depends on the average expectation in the market.

17. CONCLUSION

Economic environments are complex, and so agents operating in them are generally called upon to construct simplified models, classification schemes, and heuristics to guide their decisions. There are typically many such simplified representations and procedures that can be employed for online decision making. This indeterminacy confounds the optimization problems facing individual agents who must forecast the outcomes of their mutual interactions.

This paper explores the consequences of model uncertainty by starting at a different extreme. Suppose that the environment is quite simple, but that agents do not fully understand its underlying structure and consequently believe that the environment may be more complicated than is required by the true fundamentals.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Will these agents learn their way out of this misspecification? Or will these beliefs become self-fulfilling to some degree?

In the model considered here, there is a unique stationary REE, but a continuum of steady states for the learning process (i.e., for the forecast rule parameters) that are weakly consistent with that equilibrium. The system appears to be drawn toward this continuum, but can wander far from the learning equilibrium that corresponds to the MSV REE over long periods of time.

NOTES

1. Those models often possess "restricted perceptions equilibria" or "misspecification equilibria" that are stable under learning.

2. The stationarity of the dividend process is imposed to allow comparisons with other results on learning.

3. This functional form is clearly arbitrary but provides a simple and tractable case for analysis.

4. We have also experimented with allowing agents to observe both current dividends and the prices at which trades have been made to date during the trading period but have not followed that procedure here to avoid the introduction of an additional source of interperiod price dynamics.

5. Comparing (A.2) and (A.3) in Appendix A, we see that, due to the nonlinear overparameterization of the forecast rule, the PLM does not nest the ALM. Note that this type of specification error is often called overparameterization in the literature, leaving the term misspecification to refer to underparameterization and other types of misspecification for which the PLM does not nest an REE. There is now an extensive literature on learning with underparameterized forecast rules. The standard solution concept used there is "restricted perceptions equilibrium" [Evans and Honkapohja (2001)] or "self-confirming equilibrium" [Sargent (1999) and Cho et al. (2002)]. This has been extended for the case of competing underparameterized models by Adam (2005) and Branch and Evans (2006). Recent exceptions to the focus on underparameterized rules are Bullard et al. (2007), who consider "add-factors," and Adam et al. (2006), who consider escapes on the transition path.

6. In the E-stability framework, the parameters of the agents' PLM implied by the forecast rule (2) are assumed to move toward the corresponding parameter values of the ALM (1), (2) induced by the PLM. This learning process is formalized in a notional continuous time rather than in real time.

7. Suppose that $\mathrm{E}[F[x]] = \mathrm{E}[x]$, where $\mathrm{E}[\cdot]$ denotes unconditional expectation. Then by (1), both expected values must be equal to x^*, and (2) implies the following modified version of the WCC: (5′) $x^* = a_0 + (a_1 + a_2 + a_4) \cdot x^* + (a_3 + a_5) \cdot (x^{*2} + \mathrm{Var}[x])$. Thus, in the linear case, (5) is implied by unbiasedness. In the nonlinear case, (5) approximates (5′) for small variances, so (5) may approximate stronger consistency conditions.

8. We do not assume that traders have access to a single global order book, and so they may not see the best prices in the market. This is meant to represent search across alternative trading platforms, dealers, or private institutions. See, e.g., Harris (2003) for a general discussion of competition between markets and Duffie et al. (2005) or Lagos and Rocheteau (2007) for explicit models of such a search process.

9. Genetic algorithms are algorithms for updating populations of rules via genetics-inspired operations such as mutation, crossover, and selection. See, e.g., Holland (1975), Goldberg (1989), and Mitchell (1996). On applications in economics see, e.g., Holland and Miller (1991), Arifovic (1994), Andreoni and Miller (1995), Bullard and Duffie (1998a, 1998b, 1999), Vriend (2000), and Lux and Schornstein (2005). On real coded GAs see, e.g., Michalewicz (1996), Herrera et al. (1998), and Fogel (2000). The evolutionary algorithm used here is essentially the same as that used in Georges (2002).

10. Because we use a tournament-style selection process for reproduction, only the relative fitness values will be relevant.

11. For example, Aragones et al. (2005) have shown that the relatively straightforward problem of specification search, selecting a parsimonious set of regressors out of a given set of regressors,

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

can be a computationally hard (NP-complete) problem, and so may not be practically solvable for a modest-sized set of regressors.

12. There is a literature [e.g., Diether et al. (2002)] that documents the empirical importance of heterogeneity of opinion in financial markets. Our approach allows for diversity but also for convergence to homogeneous beliefs as a possible outcome under learning.

13. For example, for ${a_0}^i = x^* - \nu$, $\nu > 0$, the reservation price satisfying (1) is $P_t{}' = \frac{x^*-\nu}{1+r} = P^* - \frac{\nu}{1+r} \in ({a_0}^i - \bar{d}, P^*)$.

14. Memory (*leval*) is one period. Twenty traders are active per period. Each looks at the prices of four others and sets desired sell orders at 2% of current holdings of the asset to be sold. The bid-ask spread is 0.001. r is set high (0.5) to create a visible wedge between $\bar{a}_0 - \bar{d}$ and P in Figure 1. The dividend shock ε is uniformly distributed on $(-0.05, 0.05)$ and is active in each period with probability $pshock = 0.1$. Each individual searches for a new rule with probability $preproduce = 0.2$, and if so mutates each gene with probability $pmutate = 0.2$, considers imitating another agent with probability $pcompare = 0.1$, and also crosses over with that individual in that case with probability $pcross = 0.2$. Mutations are drawn from uniform distributions on $(-5.0,5.0)$ and $(-0.5,0.5)$.

15. It is worth noting that without sufficient diversity in the population, trading will die out, and consequently it is possible for the simulations to converge far from the rational expectations equilibrium. Given the absence of a Walrasian auctioneer or other external price adjustment mechanism, prices only adjust when trades are executed. If all agents have very similar forecasts, then no trades will be executed, and so the observed market price will remain unchanged. Forecasts may then evolve to predict this price correctly. The rate of return on stocks and bonds will not be the same, so all agents would prefer to hold more or less of the stock, but none will find it worthwhile to accept currently posted bids or offers. It is also worth noting that static diversity will not necessarily sustain trading, as traders positions are constrained by their wealth. Thus, trading also tends to die out if we suppress learning (evolution). Active trading in combination with learning tends to drive the center of the price distribution toward P^*, as seen in Figure 1.

16. It is worth reiterating that price dynamics depend critically on ongoing heterogeneity in the population, as, given the absence of an auctioneer or market makers, there must be heterogeneous expectations for trades to occur.

17. Mutations for a_0 are selected from uniform distributions on $(-5.0,5.0)$ and $(-0.5,0.5)$ as above. For a_5, mutations are selected from uniform distributions on $(-0.05,0.05)$ and $(-0.005, 0.005)$.

18. Recall that we are using a kind of election operator, so exploration (the mutation operator) only leads to changes in active rules if the newly discovered rules appear to be superior to the existing rules. Thus, there is no a priori guarantee that mutation will cause rules to change over time.

19. The first-order stationarity of the price level is an artifact of the assumption that dividends are stationary.

20. Interestingly, however, because prices are only affected by executed orders, price dynamics may stay relatively stable even when expectations are extreme or highly volatile and when price dynamics would be explosive in a representative agent version of the model. That is, in these simulations, price dynamics may remain bounded even if initial forecast rules make the system depicted in Appendix A globally unstable.

21. That is, social learning agents may herd in destabilizing ways in the very short run, but over longer periods of time tend to herd on relatively stable rules near the MSV rather than on less stable rules far from the MSV.

22. Thus, on average, in any period, an individual searches for a new rule by random experimentation (mutation) with probability 0.2 as above, but now, if she does so, she also considers imitating the better of two randomly selected other agents' rules with probability 0.01 and crossing over with that rule as well with probability 0.2.

23. In cases in which there is not convergence to a noisy MSV equilibrium, we can typically engineer convergence by altering the mutation rates, memory length, or other learning parameters, the dividend shock process, or initial conditions.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

24. As seen below, if agents were more directed in their search, their rules would be both more variable and more closely centered on the MSV REE rule, because they would more closely track the true global optimum (for the given fitness function), which displays such variation given the limited memory in the fitness function. Note, however, that in our simulations, learning is rapid for agents who are far from the WCC. Further, if all agents were able to search the entire parameter space with an equal degree of accuracy, then there would be no diversity, and trading would not occur. Given the diversity of expectations observed in actual financial markets and the rapid learning observed in our simulations among agents with very poor forecasting rules, we posit that the degree of inefficiency of search generated by our learning algorithm is sensible.

25. Intuitively, suppose that x_t is close to x^* over time. Then a good forecast will take x_t histories that are close to x^* into forecast values that are close to x^*. The WCC is therefore potentially a good first approximation to the optimal forecast rule at any time.

26. The fact that social and individual learning can lead to different outcomes appears to be due to the greater rapidity of learning under social learning along the lines of Allen and Carroll (2001), rather than to a fundamental difference in the selection process as in Vriend (2000).

27. These observations are based on 50 repetitions of the simulation over 20,000 trading periods for each of eight sets of parameters.

28. This model is highly stylized and we make no attempt here to replicate time series properties of actual financial markets with any accuracy. For example, while we do tend to see clustered volatility and excess kurtosis of returns, we make no claims to be able to replicate the empirical distributions for returns or for the serial dependence of volatility that are well documented for actual financial markets [see, e.g., Lux and Ausloos (2002)]. Even more fundamentally, there is substantial serial correlation in the returns generated in these simulations that is not being exploited by our technical traders and consequently arbitraged away, and so prices do not exhibit the standard martingale property.

29. The latter effect is similar to Lux (1995) and Lux and Marchesi (1999), in which agents herd on optimistic or pessimistic expectations.

30. See Aragones et al. (2005) for a formal argument. Given finite computational capabilities, we can think of specification (2) as a "fast and frugal heuristic" in the spirit of Gigerenzer and Selten (2001). Georges (2008b) finds that a simple form of specification testing reduces but does not eliminate excess volatility in a similar model with least squares learning.

31. There is some similarity to Youssefmir and Huberman (1997), where an indeterminacy leads to a wandering in the strategy space, which in turn triggers intermittent volatility. However, given the stochastic dividend in our model, there is only a near indeterminacy of optimal rules close to the WCC.

32. A complementary result is reported by Arifovic et al. (2007) in the context of a New Keynesian monetary model. They find that convergence to an REE occurs for a wider range of monetary policy rules when agents with unbounded memory learn via an evolutionary algorithm than when they learn by least squares.

33. Under alternative timing assumptions, this need not be the case. If, for example, agents were able to condition their expectations of x_{t+1} directly on x_t, then x_t would not depend directly on x_{t-2} under the ALM even though it would under the PLM (i.e., agents would incorrectly believe that x_t depends directly on x_{t-2}.

REFERENCES

Adam, Klaus (2005) Learning to forecast and cyclical behavior of output and inflation. *Macroeconomic Dynamics* 9, 1–27.

Adam, Klaus, Albert Marcet, and Juan Pablo Nicolini (2006) Stock Market Volatility and Learning. Manuscript, European Central Bank.

Alfarano, Simone and Thomas Lux (2007) A noise trader model as a generator of apparent financial power laws and long memory. *Macroeconomic Dynamics* 11 (Supplement 1), 80–101.

Allen, Todd M. and Christopher Carroll (2001) Individual learning about consumption. *Macroeconomic Dynamics* 5(4), 255–271.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Andreoni, James and John H. Miller (1995) Auctions with adaptive artificial agents. *Journal of Games and Economic Behavior* 10, 39–64.

Aragones, Enriqueta, Itzhak Gilboa, Andrew Postlewaite, and David Schmeidler (2005) Fact free learning. *American Economic Review* 95(9), 1355–1368.

Arifovic, Jasmina (1994) Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control* 18, 3–28.

Arifovic, Jasmina, James Bullard, and Olena Kostyshyna (2007) Social Learning and Monetary Policy Rules. Manuscript, Simon Fraser University.

Arthur, W. Brian, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Taylor (1997) Asset pricing under endogenous expectations in an artificial stock market. In W. Brian Arthur, Steven N. Durlauf, and David A. Lane (eds.), *The Economy as Evolving Complex System II*, pp. 15–44. Reading, MA: Addison–Wesley.

Audet, Nicolas, Toni Gravelle, and Jing Yang (2002) Alternative Trading Systems: Does One Shoe Fit All? Bank of Canada Working Paper 2002-33.

Boswijk, H. Peter, Cars H. Hommes, and Sebastiano Manzan (2007) Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control* 31, 1938–1970.

Branch, William A. and George W. Evans (2006) Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory* 127(1), 264–295.

William A. and Cars H. Hommes (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235–1274.

Brock, William A., Cars H. Hommes, and Florian Wagener (2005) Evolutionary dynamics in financial markets with many trader types. *Journal of Mathematical Economics* 41, 7–42.

Bullard, James and John Duffie (1999) Using genetic algorithms to model the evolution of heterogeneous beliefs. *Computational Economics* 13, 41–60.

Bullard, James and John Duffie (1998a) Learning and the stability of cycles. *Macroeconomic Dynamics* 2, 22–48.

Bullard, James and John Duffie (1998b) Model of learning and emulation with artificial agents. *Journal of Economic Dynamics and Control* 22, 179–207.

Bullard, James, George W. Evans, and Seppo Honkapohja (2007) A Model of Near-Rational Exuberance. Manuscript, Federal Reserve Bank of St. Louis.

Chakrabarti, Rajesh (2000) Just another day in the inter-bank foreign exchange market. *Journal of Financial Economics* 56, 29–64.

Chan, Nicholas T., Blake LeBaron, Andrew W. Lo, and Tomaso Poggio (1999) Agent-Based Models of Financial Markets: A Comparison with Experimental Markets. Manuscript, Brandeis University.

Chen, Shu-Heng and Chia-Hsuan Yeh (2001) 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), 363–393.

Chen, Shu-Heng and Chia-Hsuan Yeh (2002) 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), 217–239.

Chiarella, Carl and Xue-Zhong He (2002) Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics* 19, 95–132.

Chiarella, Carl and Giulia Iori (2002) A simulation analysis of the microstructure of double auction markets. *Quantitative Finance* 2, 346–353.

Chiarella, Carl and Giulia Iori (2004) The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows. Manuscript, University of Technology, Sydney.

Cho, In-Koo, Noah Williams, and Thomas J. Sargent (2002) Escaping Nash inflation. *Review of Economic Studies* 69, 1–40.

Darley, Vincent and Alexander V. Outkin (2007) *NASDAQ Market Simulation: Insights on a Major Market From the Science of Complex Adaptive Systems*. Singapore: World Scientific Publishing.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina (2002) Differences of opinion and the cross-section of stock returns. *Journal of Finance* 57, 2113–2141.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Duffie, Darrell, Nicolae Garleanu, and Lasse Heje Pedersen (2005) Over-the-counter-markets. *Econometrica* 73(6), 1815–1847.
Evans, George W. and Seppo Honkapohja (2001) *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
Farmer, J. Doyne and Shareen Joshi (2002) The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization* 49(1), 149–171.
Farmer, J. Doyne, Paolo Patelli, and Ilija I. Zovko (2005) The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences* 102, 2254–2259.
Fogel, David B. (2000) *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*, 2nd ed. New York: IEEE Press.
Georges, Christophre (2002) Learning Dynamics with Nonlinear Misspecification. Manuscript, Hamilton College.
Georges, Christophre (2006) Learning with misspecification in an artificial currency market. *Journal of Economic Behavior and Organization* 60, 70–84.
Georges, Christophre (2008a) Bounded memory, overparameterized forecast rules, and instability. *Economics Letters* 98(2), 129–135.
Georges, Christophre (2008b) Staggered updating in an artificial financial market. *Journal of Economic Dynamics and Control* 32(9), 2809–2825.
Gigerenzer, Gerd and Reinhard Selten (eds.) (2001) *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
Gode, Dhananjay K. and Shyam Sunder (1993) Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economics* 101(1), 119–137.
Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley.
Grandmont, Jean-Michel (1998) Expectations formation and stability of large socioeconomic systems. *Econometrica* 66(4), 741–781.
Harris, Larry (2003) *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford, UK: Oxford University Press.
Herrera, F., M. Lozano, and J.L. Verdegay (1998) Tackling real-coded GAs: Operators and tools for behavioural analysis. *Artificial Intelligence Review* 12, 265–319.
Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press [2nd ed.: MIT Press, 1992].
Holland, John H. and John H. Miller (1991) Artificial adaptive agents in economic theory. *American Economic Review Papers and Proceedings* 81, 365–370.
Hommes, Cars (2006) Heterogeneous agent models in economics and finance. In Leigh Tesfatsion and Kenneth L. Judd (eds.), *Handbook of Computational Economics: Vol. 2, Agent-Based Computational Economics*, pp. 1109–1186. Amsterdam: Elsevier.
Hommes, Cars and Gerhard Sorger (1998) Consistent expectations equilibria. *Macroeconomic Dynamics* 2, 287–321.
Honkapohja, Seppo and Kaushik Mitra (2003) Learning with bounded memory in stochastic models. *Journal of Economic Dynamics and Control* 27(8), 1437–1457.
Lagos, Ricardo and Guillaume Rocheteau (2007) Search in asset markets: Market structure, liquidity, and welfare. *American Economic Review* 97(2), 198–202.
LeBaron, Blake (2006) Agent based computational finance. In Leigh Tesfatsion and Kenneth L. Judd (eds.), *Handbook of Computational Economics: Vol. 2*, *Agent-Based Computational Economics*, pp. 1187–1233. Amsterdam: Elsevier.
LeBaron, Blake, W. Brian Arthur, and Richard Palmer (1999) Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487–1516.
LeBaron, Blake and Ryuichi Yamamoto (2006) Long-Memory in an Order-Driven Market. Manuscript, Brandeis University.
Lux, Thomas (1995) Herd behavior, bubbles, and crashes. *Economic Journal* 105, 881–896.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

Lux, Thomas and Marcel Ausloos (2002) Market fluctuations I: Scaling, multi-scaling and their possible origins. In A. Bunde and H.J. Schellnhuber (eds.), *Theories of Disasters: Scaling Laws Governing Weather, Body and Stock Market Dynamics*, pp. 372–409. Berlin: Springer-Verlag.

Lux, Thomas and Michele Marchesi (1999) Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature* 397, 498–500.

Lux, Thomas and Sascha Schornstein (2005) Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics* 41, 169–196.

McCallum, Bennett T. (1983) On non-uniqueness in rational expectations models: An attempt at perspective. *Journal of Monetary Economics* 11, 139–168.

McCallum, Bennett T. (1999) Role of the minimal state variable criterion in rational expectations models. *International Tax and Public Finance* 6, 621–639.

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

Mitchell, Melanie (1996) *An Introduction of Genetic Algorithms*. Cambridge, MA: MIT Press.

Samanidou, E.E. Zschischang, D. Stauffer, and T. Lux (2007) Agent-based models of financial markets. *Reports on Progress in Physics* 70, 409–450.

Sargent, Thomas J. (1999) *The Conquest of American Inflation*. Princeton, NJ: Princeton University Press.

Vriend, Nicholas J. (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analysis. *Journal of Economic Dynamics and Control* 24, 1–19.

Yang, Jing (2002) The efficiency of an artificial double auction stock market with neural learning agents. In S.H. Chen (ed.), *Evolutionary Computation in Economics and Finance*, pp. 85–105. New York: Springer-Verlag.

Yeh, Chia-Hsuan (2003) Tick Size and Market Performance. Manuscript, Department of Information Management, Yuan Ze University.

Youssefmir, Michael and Bernardo Huberman (1997) Clustered volatility in multiagent dynamics. *Journal of Economic Behavior and Organization* 32, 101–118.

APPENDIX A: NONLINEAR DYNAMICS UNDER A REPRESENTATIVE AGENT

Consider the following representative agent version of the model. The price is governed by the indifference condition [(1) above]. Recalling the definition $x_t \equiv P_t + d_t$ and the assumption that the dividend process follows $d_t = \bar{d} + \varepsilon_t$, we have

$$x_t = \frac{F_t[x_{t+1}]}{1+r} + \bar{d} + \varepsilon_t. \tag{A.1}$$

The PLM according to the agent's forecast rules [(2) above] is (suppressing the time subscripts on the forecast rule parameters)

$$x_t = a_0 + a_1 \cdot x_{t-1} + a_2 \cdot x_{t-2} + a_3 \cdot x_{t-2}^2 + a_4 \cdot x_{t-3} + a_5 \cdot x_{t-3}^2. \tag{A.2}$$

The ALM induced by this PLM is

$$\begin{aligned} x_t = \left(\frac{1}{1+r}\right) \{ & [a_0 \cdot (1+a_1) + \bar{d} \cdot (1+r)] + \left(a_1^2 + a_2\right) \cdot x_{t-1} \\ & + a_3 \cdot x_{t-1}^2 + (a_1 a_2 + a_4) \cdot x_{t-2} + (a_1 a_3 + a_5) \cdot x_{t-2}^2 \\ & + a_1 a_4 \cdot x_{t-3} + a_1 a_5 \cdot x_{t-3}^2 \} + \varepsilon_t. \end{aligned} \tag{A.3}$$

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

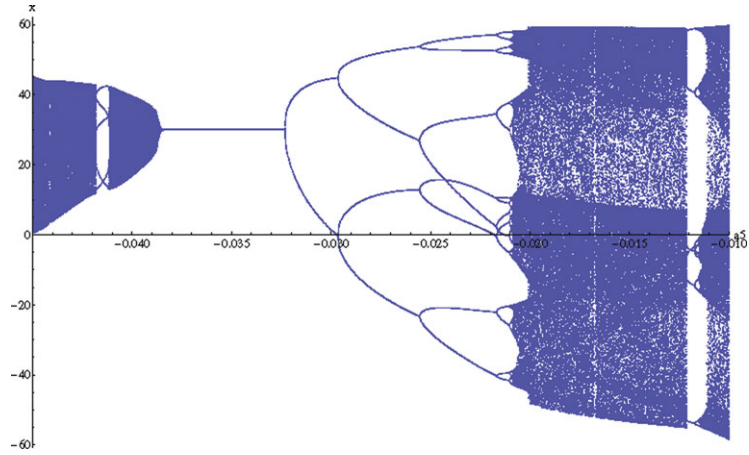


FIGURE A.1. Bifurcation diagram as a_5 and a_0 are varied subject to the weak consistency condition for $a_1 = 0$, $a_2 = -0.24$, $a_3 = -0.04$, $a_4 = 0.8$. Here, $r = 0.5$ and $\bar{d} = 10$, so $x^* = 30$.

The nonlinear terms in the representative agent's forecast function permit a broad range of global dynamics even in the absence of learning (i.e., even if the forecast rule parameters are fixed). This is true even if the weak consistency condition (5) is also satisfied. In that case, we can write the ALM (A.3) as follows, indicating that the deviations of x from the steady state expected value x^* at time t depend nonlinearly on past deviations:

$$\begin{aligned} x_t = x^* + \left(\frac{1}{1+r}\right)\big[&\left(a_1^2 + a_2\right)\cdot\left(x_{t-1} - x^*\right) + a_3\cdot\left(x_{t-1}^2 - x^{*2}\right) \\ &+ \left(a_1 a_2 + a_4\right)\cdot\left(x_{t-2} - x^*\right) + \left(a_1 a_3 + a_5\right)\cdot\left(x_{t-2}^2 - x^{*2}\right) \\ &+ a_1 a_4\cdot\left(x_{t-3} - x^*\right) + a_1 a_5\cdot\left(x_{t-3}^2 - x^{*2}\right)\big] + \varepsilon_t. \end{aligned} \tag{A.4}$$

The bifurcation diagram in Figure A.1 illustrates the variety of long run attractors permitted by the deterministic part of (A.4), i.e., by the deterministic part of the ALM (A.3) subject to the WCC (5).

Because the average REE value $x_t = x^*$ is a steady state of the deterministic part of (A.4), we can think of (5) as constraining agents' forecasts to be broadly centered on x^* and thus to be correct on average in a limited sense.

APPENDIX B: E-STABILITY

Consider the representative agent version of the model discussed in Appendix A. First consider the linear version for which $a_3 = a_5 = 0$. Thus the PLM (A.2) is reduced to

$$x_t = a_0 + a_1\cdot x_{t-1} + a_2\cdot x_{t-2} + a_4\cdot x_{t-3}$$

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

and the ALM (3) induced by the PLM is reduced to

$$x_t = \left(\frac{1}{1+r}\right)\left\{[a_0 \cdot (1+a_1) + \bar{d} \cdot (1+r)] + \left(a_1^2 + a_2\right) \cdot x_{t-1} + (a_1 a_2 + a_4) \cdot x_{t-2} + a_1 a_4 \cdot x_{t-3}\right\} + \varepsilon_t.$$

Following Evans and Honkapohja (2001), we assume that the (forecast rule) parameters in the PLM adjust in notional time toward the corresponding values of the induced ALM as follows:

$$\begin{aligned} \dot{a}_{0t} &= \frac{a_{0t} \cdot (1+a_{1t}) + \bar{d} \cdot (1+r)}{1+r} - a_{0t} \\ \dot{a}_{1t} &= \frac{a_{1t}^2 + a_{2t}}{1+r} - a_{1t} \\ \dot{a}_{2t} &= \frac{a_{1t} \cdot a_{2t} + a_{4t}}{1+r} - a_{2t} \\ \dot{a}_{4t} &= \frac{a_{1t} a_{4t}}{1+r} - a_{4t}. \end{aligned} \tag{B.1}$$

This system of equations has a steady state at $a_0 = x^*$, $a_1 = a_2 = a_4 = 0$. The Jacobian of the system linearized about this fixed point is

$$J = \begin{bmatrix} \left(\frac{1}{1+r} - 1\right) & \frac{x^*}{1+r} & 0 & 0 \\ 0 & -1 & \frac{1}{1+r} & 0 \\ 0 & 0 & -1 & \frac{1}{1+r} \\ 0 & 0 & 0 & -1 \end{bmatrix}. \tag{B.2}$$

The eigenvalues of this Jacobian are the elements on the main diagonal. Thus, given $r > 0$, these eigenvalues are all (real and) negative, and so the steady state is locally asymptotically stable. Consequently, the stationary REE is strongly E-stable in the linear case.

Further, the stationary equilibrium is also strongly iteratively E-stable. That is replacing the system of differential equations above with the discrete time analog

$$\begin{aligned} a_{0t} &= \frac{a_{0t-1} \cdot (1 - a_{1t-1}) + \bar{d} \cdot (1+r)}{1+r} \\ a_{1t} &= \frac{a_{1t-1}^2 + a_{2t-1}}{1+r} \\ a_{2t} &= \frac{a_{1t-1} a_{2t-1} + a_{4t-1}}{1+r} \\ a_{4t} &= \frac{a_{1t-1} a_{4t-1}}{1+r}, \end{aligned} \tag{B.1′}$$

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press

we have the same steady state, and the corresponding Jacobian

$$J = \begin{pmatrix} \frac{1}{1+r} & -\frac{\bar{d}}{r} & 0 & 0 \\ 0 & 0 & \frac{1}{1+r} & 0 \\ 0 & 0 & 0 & \frac{1}{1+r} \\ 0 & 0 & 0 & 0 \end{pmatrix}. \qquad \textbf{(B.2}') $$

The eigenvalues of this matrix are the elements on the main diagonal, which are all $\in (-1, 1)$ for any $r > 0$.

A second steady state of (B.1) ($a_0 = \frac{\bar{d}\cdot(1+r)}{1+2r}$, $a_1 = 1+r$, $a_2 = a_4 = 0$) represents bubble REE. However, that steady state is not locally asymptotically stable, and so bubble REE are not E-stable. They would not, for example, be learnable under least squares learning.

Note that, even though the PLM is overparameterized, it induces an ALM with the same structural form.[33] If, however, we include the additional nonlinear terms in the forecast rules (i.e., $a_3, a_5 \neq 0$), then this is no longer the case. The PLM (A.2) has no term to match the term $a_3 \cdot x_{t-1}^2$ in the ALM (A.3) and so no longer nests the ALM.

https://doi.org/10.1017/S1365100509080262 Published online by Cambridge University Press