

Time Series Properties of an Artificial Stock Market *

Blake LeBaron W. Brian Arthur Richard Palmer
Brandeis University Santa Fe Institute Duke University

September 1998

Forthcoming: *Journal of Economic Dynamics and Control*

Abstract

This paper presents results from an experimental computer simulated stock market. In this market artificial intelligence algorithms take on the role of traders. They make predictions about the future, and buy and sell stock as indicated by their expectations of future risk and return. Prices are set endogenously to clear the market. Time series from this market are analyzed from the standpoint of well known empirical features in real markets. The simulated market is able to replicate several of these phenomenon, including fundamental and technical predictability, volatility persistence, and leptokurtosis. Moreover, agent behavior is shown to be consistent with these features, in that they condition on the variables that are found to be significant in the time series tests. Agents are also able to collectively learn a homogeneous rational expectations equilibrium for certain parameters giving both time series and individual forecast values consistent with the equilibrium parameter values.

JEL Classification: G12, G14, D83

Keywords: Learning, Asset Pricing, Evolution, Financial Time Series

*Blake LeBaron, Brandeis University, Department of Economics, Mailstop 021, Waltham, MA 02453-2728. LeBaron is grateful to the Alfred P. Sloan Foundation and the University of Wisconsin Graduate School for support. Part of this work represents "The Santa Fe Artificial Stock Market," a project begun at the Santa Fe Institute along with J. Holland and P. Tayler. The authors thank the editor and two anonymous referees for helpful suggestions. Also, seminar participants from the NBER asset pricing group, Northwestern and the University of Pennsylvania have provided useful comments.

1 Introduction

The picture of financial markets as groups of interacting agents, continually adapting to new information, and updating their models of the world seems like a very appealing and accurate picture of how real markets operate. However, trying to construct theoretical frameworks for markets can become extremely complicated. The state space quickly explodes as beliefs of all agents become relevant, and solutions become difficult when individual beliefs and decisions depend on those of others. Recently, Arthur, Holland, LeBaron, Palmer & Tayler (1997) have presented a computational platform for analyzing such a market. This paper extends their results, concentrating on the time series features of these artificial markets.

The behavior of traders in this simulated market draws heavily on the literature on convergence to rational expectations under a misspecified learning dynamic.¹ We follow the tradition of assuming that the out of equilibrium learning process would be extremely complex to set up as the result of an expectational equilibrium. Moreover, the complexity of the situation forces agents to act inductively, using simple rules of thumb, in their attempts at optimization. These rules are not static, and are continuously reevaluated and updated according to their performance. No agent will continue to use suboptimal rules when better ones have been discovered.

We make two strong deviations from the earlier research in this area. First, we are not interested simply in equilibrium selection, and convergence properties alone. We are interested in the behavior of the learning and adapting market per se. Most of the reason for this is that we have situations in which the market never really settles down to anything that we could specifically characterize as an equilibrium. Our second major contribution is that we add variable selection to the agent's forecasting problem. They not only update linear prediction models, they must also make some selection as to what information is relevant for their forecasts. They are faced with a large set of information from which to make forecasts, of which only a small part might be relevant. They must decide empirically which series to use. This is a continuation of Sargent's (1993) recommendation for building little econometricians into models. It just moves the level of complexity

¹A good example of this approach is Bray (1982) which is part of the large body of work on bounded rationality introduced by Simon (1969), and recently summarized in Conlisk (1996) and Sargent (1993). There are also connections to the large literature on heterogeneous information rational expectations surveyed in Admati (1991), and with many of the important early papers in Grossman (1989). However, in the experiments performed here differences across agents are due only to differences in interpretation of common public information.

one step closer to reality.

The use of multiagent models for financial markets is driven by a series of empirical puzzles which are still hard to explain using traditional representative agent structures. Among these puzzles are issues of time series predictability using both technical and fundamental information.² Other important features include volatility persistence, and the equity premium puzzle.³ Although some theoretical explanations have been provided for a few of these facts, most of the explanations still remain controversial, and a general theory for all these features remains out of reach.

The literature on building artificial financial markets using computer experiments is steadily growing and new markets are appearing all the time.⁴ Paralleling the computer work have been several new analytic approaches to heterogeneity. They view the market as made up of populations of different strategy types, usually including technical, fundamental, and rational traders.⁵ These analytic papers complement the computer work in many ways. They are able to address some of the same issues, in a tighter analytic framework, but they are slightly restricted in agents' forecasting rules need to be specified ahead of time.

There is some connection between the results here, and the work on learning in sunspot equilibria such as Evans (1989), and Woodford (1990). The artificial markets in this paper often end up with agents conditioning on variables that should be of no value, but become valuable since others are paying attention to them too. This is similar to the support for a sunspot under a learning dynamic. However, it should be cautioned that sunspot models deal with stable equilibria in which extraneous variables have value. Often the artificial markets do not appear to be converging to anything that looks as stable as a traditional sunspot. It is also the case that emergence of variables as important pieces of forecast information will be endogenous. Traders will be shown certain pieces of information, but they are not encouraged to use them.

²See Campbell & Shiller (1988*a*), Campbell & Shiller (1988*b*) for some of the early work on predictability. Also, Campbell, Lo & MacKinlay (1996) and Fama (1991) provide good surveys of this large research area. The area of technical forecasting, or using past prices alone to forecast future prices has seen a recent resurgence. Papers by Brock, Lakonishok & LeBaron (1992), LeBaron (1998*b*), Levich & Thomas (1993), Sweeney (1988), Sweeney (1986), and Taylor (1992) have reopened the question of the usefulness of simple technical trading rules. Taylor & Allen (1992) show that a majority of traders in foreign exchange markets admit to using technical indicators.

³See Bollerslev, Chou, Jayaraman & Kroner (1990) for a survey of the evidence on volatility persistence in financial markets, and Kocherlakota (1996) for a summary of work on the equity premium puzzle.

⁴Several of these include Arifovic (1996), Lettau (1997), de la Maza & Yuret (1995), Margarita & Beltratti (1993), Marengo & Tordjman (1995), Rieck (1994), Routledge (1994), Steiglitz, Honig & Cohen (1996), and Youssefmir & Huberman (1997). LeBaron (forthcoming 1998*a*) gives a survey of some of this literature.

⁵Examples of this are in Brock & Hommes (1998), Brock & LeBaron (1996), Chiarella (1992), Day & Huang (1990), De Grauwe, Dewachter & Embrechts (1993), Frankel & Froot (1988), Gennotte & Leland (1990), Lux (1998), and Zeeman (1974).

Therefore, their coordination on these variables itself is an interesting dynamic of the market. A related equilibrium concept is Kurz's (1994) rational beliefs equilibrium. In a rational beliefs equilibrium, agents again concentrate on the empirical consistency of their forecasting models. Models are only rejected if they are inconsistent with observed time series. This allows many more possible equilibria to occur, and Kurz & Beltratti (1995) have used this idea to help explain the equity premium.

Another related modeling area is "behavioral finance", or "noise trader" models such as De Long, Shleifer & Summers (1990). In these settings some set of traders make decisions which are less than perfectly rational, and more importantly, drive up volatility, and force prices down inducing rational agents to hold the risky asset. In our model there is a similar empirical feature in that prices fall, and agents' estimated volatilities rise above the rational expectations benchmark values. However, this result is reached through a very different route from the classic noise trader approach. Traders are bounded in their information processing abilities, but subject to these constraints they do not believe they are making systematic mistakes. Also, there is no division of the traders into a specific noise trader group. While some traders may behave in such a way as to increase volatility, they do it because they evolved these behavioral patterns. Persistent psychological biases are interesting in finance and economics, but they are not part of this model.⁶

Section 2 presents the market structure, and details the computerized trading agents. Section 3 gives the empirical results from the computer experiments, and section 4 concludes and discusses the future for this line of research.

2 Market Structure

2.1 Tradable Assets

The market structure is set up to be as simple as possible in terms of its economic components. It also attempts to use ingredients from existing models wherever possible.

There are two assets traded. First, there is a risk free bond, paying a constant interest rate, $r_f = 0.10$, in infinite supply. The second asset is a risky stock, paying a stochastic dividend which is assumed to follow

⁶Also, see Thaler (1992) for further examples of biases in economic decision making.

the following autoregressive process,

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \mu_t, \quad (1)$$

with $\bar{d} = 10$, and $\rho = 0.95$, and $\mu_t \sim N(0, \sigma_\mu^2)$. This process is aimed at providing a large amount of persistence in the dividend process without getting close to nonstationary dividend processes. The total shares of the stock are set to 25, the number of agents who will be trading. The price of a share of stock, p_t , is determined endogenously in the market.

2.2 Preferences

There are 25 agents, who are assumed to be myopic 1 period constant absolute risk aversion (CARA) investors maximizing an expected utility function of the following form,

$$\hat{E}_t^i(-e^{-\gamma W_{t+1}^i}) \quad (2)$$

subject to,

$$W_{t+1}^i = x_t^i(p_{t+1} + d_{t+1}) + (1 + r_f)(W_t^i - p_t x_t^i). \quad (3)$$

The notation \hat{E}_t^i indicates the best forecast of agent i at time t . It is not the conditional expectation, but agents i 's perceived expectation. If stock prices and dividends were gaussian then it is well known that the share demand for these preferences is given by,

$$x_t^i = \frac{\hat{E}_t^i(p_{t+1} + d_{t+1}) - (1 + r_f)p_t}{\gamma \hat{\sigma}_{p+d,i}^2}, \quad (4)$$

where $\hat{\sigma}_{p+d,i}^2$ is agents i 's forecast of the conditional variance of $p + d$. This relation only holds under gaussianity of stock prices. This will hold in the linear rational expectations equilibrium given below, but outside of this it is not clear what the distribution of stock prices will be. In these cases the demand function should be taken as given, but the connection to a CARA utility maximizer is broken.

Given agents are identical with the same coefficient of absolute risk aversion, γ , then it is easy to solve

for a homogeneous linear rational expectations equilibrium (REE) by conjecturing a linear function mapping the current state into a price,

$$p_t = f d_t + e. \quad (5)$$

Plugging this into the demand for shares and forcing each agent to optimally hold one share at all times gives,

$$f = \frac{\rho}{1 + r_f - \rho} \quad (6)$$

$$e = \frac{\bar{d}(f + 1)(1 - \rho) - \gamma\sigma_{p+d}^2}{r_f}. \quad (7)$$

Given this direct connection between the price and the dividend, and the dividend dynamics, it is an easy task to find optimal forecasts in the homogeneous rational expectations equilibrium. One form of these is

$$E(p_{t+1} + d_{t+1}) = \rho(p_t + d_t) + (1 - \rho)((1 + f)\bar{d} + e). \quad (8)$$

This forecasting equation is one of the crucial things that agents will be estimating, and it will often be referred to with the following parameters,

$$\hat{E}(p_{t+1} + d_{t+1}) = a(p_t + d_t) + b. \quad (9)$$

This equilibrium will be used as a benchmark for many of the experiments. It will be interesting to see how close the agents parameters come to these, and if the time series properties in some of the simulated markets are close to what we would expect in the converged market. For easy reference the simulation parameters are summarized in table 1. The allowable ranges for a and b are given in the lines labeled, “a range” and “b range” which are centered around their REE values.

2.3 Forecasting

The goal of the artificial agents is to build forecasts of the future price and dividend which they will use in their demand functions above. They do this by maintaining a list of several hypothesis, or candidate

forecasting rules. These rules may apply at all times, or in certain specific states of the world.⁷ They are monitored for forecast accuracy, and only the best rules will be used in forecasting. After several trading periods have gone by these rules are grouped together, and the worst performing rules will be eliminated, and new ones are added using a genetic algorithm which tries to take useful pieces of good rules and build them into even better rules.

Each agent contains a table of 100 of their own rules mapping states into forecasts. These rules govern behavior while the agent is trading through their forecast of future prices and dividends. The rules are modified in a learning process that occurs at a lower frequency than trading. There is no interaction among the agents' rule books. In other words, there is no role here for imitative behavior. Reactions to other agents' behavior only occurs directly through prices.

Agents build forecasts using what are called "condition-forecast" rules. This is a modification of Holland's "condition-action" classifier system.⁸ The basic idea is that the rules will match certain states of the world which are defined endogenously. These states map into a forecast for the future price and dividend which is then converted to share demand through the agent's demand function.⁹

The classifier rules will be used to match binary conditions in the market which requires predefining a set of binary states that can be used for forecasting. This limits our agents to forecasts built on the state variables we give them, and creates a kind of focal point. However, the classifiers allow agents to ignore any of these state variables. The set of states includes both "technical" and "fundamental" information. The fundamental information will be based on dividend price ratios, and the technical information will use moving average types of trading rules. The states are summarized in a binary state vector 12 bits long. Each element corresponds to whether the conditions in table 2 are true or false.

⁷In many ways this is very close to behavioral rules of thumb such as those considered in Cochrane (1989) or Campbell & Mankiw (1989). One important difference here is that the rules of thumb will change over time if better rules can be found. In other words while there isn't explicit optimization, learning and improvement is allowed.

⁸The structure of classifier systems is developed in Holland, Holyoak, Nisbett & Thagard (1986). An example of another use in economics is Marimon, McGrattan & Sargent (1990). There has been some controversy about whether classifiers are able to solve dynamic optimization problems. This is stated most clearly in Lettau & Uhlig (1997), who show that they often fail to find dynamic solutions. Since this market consists only of myopic traders, the debate over dynamic optimization is not relevant here.

⁹Holland's "condition-action" classifiers map directly into actions. This could be done here too, but this would force the payoffs, or strengths of the different rules to be based on how much payoff is generated for the agent, either direct, or utility based. This would cause an unfair competition between rules that are good for good states of the world, and rules that are good for bad states of the world. Only the former would survive. The forecasting rules will be based on forecast accuracy which lessens this problem a little.

The choice of these bits is arbitrary, but some tests for robustness have been performed. Removing small sets of the bits do not change the results. However, making big changes to this list of information will cause big changes. For example, removing 3 out of the 4 technical trading bits can have a big impact, but removing only one will not change things. The cutoff points for the dividend/price ratios were determined by examining the region in which the price was moving for many different runs. It was desired to have good coverage over the range of observed prices. This coverage does mean that the tails will be visited infrequently. The extreme events at $1/4$ and $9/8$ are visited with probabilities less than 0.01. The moving average values are lengths that are commonly used by traders, but these should be viewed with some caution since the time horizon has not been calibrated. At this time they can be taken as values which cover a wide range of persistence in the price series.

Classifier rules contain two parts. The first is a bit string matching the state vector, and the second is a forecast connected with that condition. The bit string part of the classifier contains 12 positions corresponding to the states, and each contains one of three elements, a 1, 0, or a #. The 1 and 0 match corresponding bits in the state vector, while the # symbol is a wildcard symbol matching either. A shorter example rule is given in table 3 along with the bit strings that it would match. The wildcard character is a crucial part of the classifier system for experiments of this kind, since it allows agents to dynamically decide which pieces of information are relevant, and to completely ignore others. These are critical to the kinds of information selection issues discussed in the introduction.

The second part of the rule table helps convert the matched set of bits into a price-dividend forecast. Forecasts are built as linear functions of current prices and dividends. For each matched bit string there is a corresponding real valued vector of length three corresponding to linear forecast parameters, and a conditional variance estimate. Let this vector be given by (a_j, b_j, σ_j^2) . This is mapped into the expectation formation as follows,

$$\hat{E}_t(p_{t+1} + d_{t+1}) = a_j(p_t + d_t) + b_j \tag{10}$$

$$\hat{\sigma}_{p+d}^2 = \sigma_j^2. \tag{11}$$

Comparing equation 10 with equation 8 shows that the linear REE is in the set of things that the agent can

learn. In this case all the bits from table 2 are irrelevant since expected $p + d$ will be a linear function of current $p + d$.

The parameters a and b are initially set to random values distributed uniformly about the REE equilibrium forecast values. For each rule they then remain fixed for the rest of the life of that forecasting rule. New values are set through evolutionary procedures described in later sections.¹⁰

2.4 Trading

Once each agent has settled on a rule, they then substitute the forecast parameters into their demand for shares. For agent i the demand for shares would look like,

$$x_t^i(p_t) = \frac{\hat{E}_t^i(p_{t+1} + d_{t+1}) - (1 + r_f)p_t}{\gamma \hat{\sigma}_{p+d}^2} \quad (12)$$

$$x_t^i(p_t) = \frac{a_{i,j}(p_t + d_t) + b_{i,j} - (1 + r_f)p_t}{\gamma \hat{\sigma}_{i,j}^2}. \quad (13)$$

Now it is a trivial exercise for an auctioneer to find a price which clears the market by balancing the demand to the fixed supply of shares.¹¹

$$\sum_{i=1}^N x_t^i(p_t) = \sum_{i=1}^N \frac{a_{i,j}(p_t + d_t) + b_{i,j} - (1 + r_f)p_t}{\gamma \hat{\sigma}_{i,j}^2} = N \quad (14)$$

At the end of trading agents update the accuracy of all the matched forecasting rules according to an exponentially weighted average of squared forecast error,

$$\nu_{t,i,j}^2 = \left(1 - \frac{1}{\tau}\right) \nu_{t-1,i,j}^2 + \frac{1}{\tau} \left((p_t + d_t) - (a_{i,j}(p_{t-1} + d_{t-1}) + b_{i,j}) \right)^2. \quad (15)$$

¹⁰It would be possible to easily estimate these parameters using ordinary least squares (OLS), since the eventual forecast criteria will be based on squared error, and once the difficult problem of what should go on the right hand side of the regression has been determined by the bitstrings. In order to keep the learning mechanism consistent across both the real and bitstring components of the rules it was decided that the real parts should be “learned” in the same way as the bit strings through the genetic algorithm.

¹¹This is not a perfect expectational equilibrium since the rule used is not conditioned on the current price information. Doing this causes the demand function to be a complicated nonlinear function for which the equilibrium price would be hard to find.

For the experiments in this paper τ is fixed at 75. Also, the squared error inputs into the above estimate are not allowed to exceed 500. This bound is almost never reached after the first few periods have gone by.¹²

Fixing τ at 75 is a crucial design question. The value of τ determines the horizon length that the agent considers relevant for forecasting purposes. It does this in a smooth exponentially weighted fashion, but it does set an arbitrary cutoff on information. If agents used all past data, then they would be making the implicit assumption that the world they live in is stationary. On the other hand if $\tau = 1$, then the rules use only the last periods forecast error, and are essentially evaluated based on noise. The value of τ is critical since it is one of the two important economic variables related to learning, and speeds of measurement and adjustment. A complete study of the changes in dynamics related to τ is beyond the scope of this paper, but a few robustness checks were performed on the results presented in the next section. Values from $\tau = 25$ to $\tau = 250$ do not greatly affect the results. However, values outside this range may give different outcomes. In the future it will be desirable to allow agents to adjust this parameter endogenously.

Trades are subject to some limitations. First, agents are restricted to trading a maximum of 10 shares each period. Also, they must stay inside a budget constraint. Their budgets are set up so that the constraints are not binding in most cases. Finally, there is a short sale restriction of 5 shares. This means that no one agent can go short more than 5 shares which would be 1/5 of the entire market. All of these conditions appear to only be binding in the early stages of the market, when agents' random initial conditions set up very unstable strategies for a short time.

2.5 Learning and rule evolution

The rules described so far describe a static rule set without learning. The agent is allowed to learn and change behavior by altering this set of rules. This adaptation is designed to eliminate poorly performing rules, and to add new ones through the use of a genetic algorithm (GA).¹³ The genetic algorithm is designed to combine old rules into new, novel hypotheses, which will then be allowed to compete with the existing population of forecasting rules.

¹²The current estimated volatility, $\nu_{t,i,j}^2$, is kept track of for each rule, but it is not used to update the variance estimate, σ_j^2 , until the rules are updated with a genetic algorithm. This is described in section 2.5.

¹³The genetic algorithm was originally developed by Holland (1992). Goldberg (1989) provides a useful introduction, and Mitchell (1996) provides a good overview of the literature. Further references to work on general evolutionary optimization procedures can be found in Bäck (1996) and Fogel (1995).

On average, the GA is implemented every k periods for each of the agents, asynchronously, so that learning does not occur for all agents in certain periods. k will be the one crucial parameter that is allowed to change across the experimental runs.

Rules are selected for rejection and persistence into future generations based on a strength measure. First, the variance estimate is updated to the current active variance for all rules.¹⁴

$$\sigma_j^2 = \nu_{t,j}^2 \tag{16}$$

Then the strength of each rule is given by

$$s_j = -(\sigma_j^2 + cB_j), \tag{17}$$

where B_j is the number of bits not equal to $\#$ in a rule, and c is a cost per bit which is set equal to 0.005 for all experiments. The purpose of this bit cost is to make sure that each bit is actually serving a useful purpose in terms of a forecasting rule. If the market was in the linear REE, then as long as the linear forecasting parameters were correct all rules should have the same forecast accuracy. The bit cost gives a gentle nudge in this situation to the all $\#$ rule. It also biases the results slightly toward the REE, but as we will see this bias is not enough in some cases.

Learning occurs by removing old rules, and generating new ones to replace them. Specifically, the 20 worst rules are eliminated out of 100, and 20 new rules are generated to replace them. New rules are generated through one of two procedures, crossover and mutation.¹⁵ In crossover new rules are generated by first choosing two fit parents from the population of existing rules. This is done using tournament selection. For each rule needed, two are chosen, and the strongest of these two will be used. After two parents have been selected, they are “crossed”. Crossing tries to combine useful parts of each parent. A method called uniform crossover is used on the bitstring component of the parent. The new child’s bit string is built one bit at a time, choosing a bit from each parent in the corresponding position with equal probability. Crossing the real part of each rule, (a, b) , is a little more difficult, and there is very little experience in doing this in

¹⁴In many of these cases the i subscript for the agent is left out to keep the notation simpler.

¹⁵Crossover occurs with probability 0.1, and mutation with probability 0.9.

the GA community. We have decided to use one of three different methods chosen at random. First, all the real values are chosen from one parent chosen at random. Second, a real value is chosen from each parent with equal probability. Third, the new values are created from a weighted average of the two parents values, using $\frac{1}{\sigma_j^2}$ as the weight for each each rule. The weights are normalized to sum to 1. In mutation, one parent is chosen using tournament selection. Then bits are flipped at random.¹⁶ The real valued components (a, b) have random numbers added to them.¹⁷ Mutation is an important part of the exploration process in a GA, and helps to maintain a diverse population. Rules inherit the average forecast accuracy, σ_j^2 , of their two parents unless the parents haven't been matched. In this case they receive the median forecast error over all rules.

2.6 Other details

There are several other details which are important to the functioning of the market, but which are not in the basic agent design structure. Each agent maintains a default rule that will be matched in any situation.¹⁸ The real valued parameters, (a, b) , are set to a weighted average of the values for each of the other rules. The weighting is determined by $\frac{1}{\sigma_j^2}$ for each rule.

The GA rule generation procedure is capable of generating rules that will never be matched. For example, a rule could appear in which it is required that the dividend price ratio be greater than 1, and less than 1/2. Obviously, such a rule is useless, and will never be used. Rules that have not been matched for 4000 periods have 1/4 of their bits (1's or 0's) set to #, and their strength is reset to the median value. This procedure is known as generalization, and is an important part of keeping stagnant rules out of the rule book.

When multiple rules are matched, the tie is broken using the active variance estimate, $\nu_{t,i,j}^2$. This is done for several reasons. First, it keeps bad rules from ever being used. Second, it delinks the tie breaking mechanism from the denominator estimate in the share demand equation. Otherwise, this would generate a lot of bias toward larger share holdings because of the selection of lower variance rules. There is already

¹⁶With probability 0.03 each bit in the string undergoes the following changes. $0 \rightarrow \#$ with probability 2/3. $1 \rightarrow \#$ with probability 2/3. $\# \rightarrow 0$ with probability 1/3. Other changes are as expected, ie. $0 \rightarrow 1$ with probability 1/3. On average this preserves the "specificity", or fraction of #'s, of a rule.

¹⁷With probability 0.2 they are changed to a value chosen randomly from their allowable ranges, given in table 1, using a uniform distribution. With probability 0.2 they are changed a small amount chosen uniformly over the range $+ - 0.05\%$ times the allowable range.

¹⁸The bitstring is constrained to be all #.

some of this present since the variance values are closely linked.

At startup the rules are initialized randomly. The bitstrings are set to $\#$ with probability 0.9, and 1 or 0, with probabilities of 0.05 each. The real values, (a_j, b_j) are set to random values in the allowable range (table 1), and the variance estimates σ_j^2 are set to the variance estimate in the rational expectations equilibrium, $\sigma_{p+d}^2 = 4$. Each agent is then allocated 1 share of stock, and 20,000 units of cash.

3 Experiments

3.1 Experimental design

The experiments performed will concentrate on the time series behavior of the simulated price series in several different situations. The objective is to show that for some parameter values the price series are indistinguishable from what should be produced in a homogeneous rational expectations equilibrium. For other parameters convergence does not occur, and the market generates interesting features which appear to replicate some of those found in real financial time series. Also, a few experiments will be performed to check if the agents could learn the rational expectations equilibrium if that was the data they were actually seeing. In these cases the theoretical price series is given to the agents without any trade taking place. This is somewhat like learning from a teacher, and will be referred to as the REE example case. It performs a quick reality check to make sure the agents could learn the REE if they happened to find themselves in it.

The crucial parameter for controlling the behavior of the market is the frequency of learning. All other parameters will remain fixed. We will experiment with two frequencies, $k = 250$, where the agents learn on average once every 250 time periods, and $k = 1000$, where they learn on average once every 1000 time periods. These will be referred to as fast learning and slow learning, respectively.

3.2 Time series features

To generate time series experiments, the market was run for 250,000 time periods to allow sufficient learning, and for early transients to die out. Then time series were recorded for the next 10,000 time periods. This was done for a cross section of 25 separate runs each started at different random starting values. Figure

1 displays a snapshot of one of these runs for the fast learning, $k = 250$, case. The actual price series is compared with the theoretical value from equation 5.

It is interesting to think about what these price series would look like if it were converging close to the REE. Since there is stochastic learning going on, the price would always have an extra amount of variability relative to the REE benchmark, and it is difficult to estimate how much this would be. From equation 7 it is clear that increased price variability alone will just lower the price by a constant amount. To compare the actual results with this, the difference between the actual and theoretical prices is plotted. This series shows a pattern of periods in which it doesn't look too far off a constant difference, and periods in which it fluctuates wildly. This phenomenon appears as a common feature for many of the runs, with the price looking close to the REE for a while, and then breaking away, and doing something else for a short period of time. In other words, the market is not inefficient all the time, but it goes through some phases where the theoretical asset valuation rules are not followed very closely.

These pictures are only representative of what is going on. The next tables present some statistical tests on the time series from the market. Under the null hypothesis of a linear homogeneous REE, price and dividend should be a linear function of their lags. This is clear from the dividend process, equation 1, and the linear dividend to price mapping in equation 5. To check this, the market price and dividend is regressed on a lag and a constant,

$$p_{t+1} + d_{t+1} = a + b(p_t + d_t) + \epsilon_t, \quad (18)$$

and the estimated residual series $\hat{\epsilon}_t$ is analyzed for structure. In the homogeneous REE this series should be independent and identically distributed, $N(0, 4)$. This is done for both the fast learning, $k = 250$, and the slow learning series, $k = 1000$.¹⁹ Results are given in table 4.

The first row shows the standard deviation of the residuals from the two different types of runs. The parameters were chosen so that in both cases the theoretical value of these was 2. Both cases show a higher amount of variability than should be there, and there is little difference between the cases. The next row shows the excess kurtosis for the two cases, which should be zero under a gaussian distribution. Both show a significant amount of excess kurtosis, but it is more pronounced for the fast learning case. Qualitatively, this

¹⁹Results for the REE example case are also performed, but in these cases the price is set to the REE value, so the estimates are as expected.

lines up well with the fact that real asset returns are leptokurtotic, but it should be noted that quantitatively, the numbers here are still smaller than those for daily asset returns. The third row shows the autocorrelation in the residuals. This is an important number because it shows that there is little extra linear structure remaining. The small magnitude of this autocorrelation is comparable to the very low autocorrelations observed in real markets. It is an important feature that any artificial market needs to replicate. The fourth row performs a simple test for ARCH dependence in the residuals.²⁰ The numbers reported are the means of the test statistics. The numbers in brackets are the fraction of runs that rejected “no ARCH” at the 95% confidence level. In the fast learning case all the runs rejected “no ARCH”, while in the slow learning case this was reduced to only 44%. A second test for volatility persistence is performed in the next row where the first order autocorrelation of the squared residuals is estimated. The value of 0.064 in the fast learning case, is clearly larger than the slow learning case value of 0.017. The next row presents the BDS test for nonlinear dependence. This test is sensitive to any kind of dependence in a time series, including ARCH, and generally rejects independence for most financial time series.²¹ It shows a large fraction, 84%, rejecting the null hypothesis of independence in the fast learning case, and a much smaller fraction, 24%, rejecting in the slow learning case. The BDS test is sensitive to many departures from independence including ARCH, so the last group of tests may be detecting similar features. The next row shows the mean excess returns for the two runs, which is estimated as $\frac{p_{t+1} + d_{t+1} - p_t}{p_t} - r_f$. The estimate is higher in the fast learning case with a value of 3.06% as compared with 2.89% in the slow learning case.²² It is also much larger than the value from the REE example using the linear benchmark. In this case the mean excess return was estimated at 2.5% with a standard error of 0.002. This gives some evidence for an increase in the equity premium in the fast learning case over both the slow learning, and REE example. The final row presents trading volume in shares per period averaged over the final 10,000 periods of the market. There is a large increase in volume going from the slow to fast learning case with almost a doubling in the amount of trading activity. This change is very strongly significant given the estimated standard errors.

In summary, the results show the fast learning case to be slightly more leptokurtotic than the slow

²⁰This is the test for conditional heteroskedasticity proposed by Engle (1982).

²¹Brock, Dechert, Scheinkman & LeBaron (1996) introduce and derive the BDS test, and Brock, Hsieh & LeBaron (1991) give further descriptions and examples of its use with financial variables.

²²This difference in means is significant at the 95% confidence level with a simple t-test value of 3.0.

learning case. Also, the fast learning cases show much stronger evidence for the presence of ARCH effects, and possibly other nonlinearities. Both cases show little indication of autocorrelation, which is similar to actual markets.

Table 5 looks at several aspects of predictability remaining in the time series beyond that which should be there in the REE equilibrium. This is done by adding extra explanatory variables to a simple linear regression of price and dividend on lags,

$$p_{t+1} + d_{t+1} = a + b(p_t + d_t) + cI_{t,MA(500)} + \epsilon_{t+1}. \quad (19)$$

Here, an indicator variable showing whether the price is above or below a 500 period moving average is added to the regression on lagged $p + d$.

Row 1 of table 5 presents the results for a regression using equation 19 with a 5 period MA for both the slow and fast learning cases. The number reported is the estimated coefficient, c , and the number in parenthesis is the standard error estimated over the cross section of runs. The results show no significant extra predictability coming from the 5 period MA. However, moving down to the next entry in the table shows a different story. The same regression is performed here, but this time the 500 period moving average is used as the additional information. In this case, the coefficient is positive and significant for the fast learning case, but insignificantly different from zero for the slow learning case. This shows a clear distinction between the two cases in terms of technical predictability, with the fast learning case showing evidence for some kind technical trading predictability, and the slow learning case still looking close to the REE result of no predictability. The last row gives a similar finding for one of the dividend price ratio indicators. This more fundamental indicator should also be of little use in the REE equilibrium since all the relevant information is contained in the lagged price. The estimated coefficient is significant and negative for the fast learning case, indicating that when prices are high relative to dividends they are expected to come down. The estimated coefficient for the slow learning case is insignificantly different from zero. The two cases from the fast learning case, line up with certain empirical features of real markets. The usefulness of the price dividend ratios has been shown in many situations to be a useful predictor. Also, technical trading rules

have been found to have some predictive value as well.²³

In the next two figures the dynamic properties of the artificial market are addressed. Specifically, correlations and cross correlations of trading volume. In actual stock markets, trading volume series are highly persistent.²⁴ Figure 2 shows the mean autocorrelations for trading volume in the two artificial market cases, fast and slow learning, along with a comparison series from IBM daily volume from 1962 to 1994. The IBM volume is detrended using a 100 day moving average of past volume. It is clear that all three generate very persistent trading volume series. However, the IBM series appears to have a much longer range dependence than the artificial series. Both the fast and slow learning markets generate similar patterns for volume persistence.

Figure 3 shows the relation between volume and volatility. The positive contemporaneous correlation between them is well known Karpov (1987). The figure shows the cross correlation pattern between volume at time t , and squared returns at time $t + j$.²⁵ The dashed lines are the 95% confidence bounds on the estimated correlations using the standard error from the cross section of the artificial market runs. The dotted line is the corresponding value for IBM daily returns. The figure shows a very close correspondence between the two patterns including a slight tendency for the volatility to lead movements in trading volume.

3.3 Agent forecast properties

Markets with computerized traders offer several advantages over experimental markets with real people. One of these is that it is possible to ask the agents just what they are doing when the experiment is over, and they will have to reveal truthfully the kinds of strategies they are using. In this section several pieces of information from the individual agents will be used to show how they have reacted to the previously reported time series averages.

Figure 4, shows the fraction of bits set for the two different cases, slow learning, and fast learning. Fraction of bits set refers to the fraction of bits that are not don't care symbols, #, averaged over all the rules and agents. In other words, this reports the fraction of rules that are conditioned on different pieces of

²³See Campbell & Shiller (1988a) and Campbell & Shiller (1988b) for evidence on price/dividend ratios, and see Brock et al. (1992) for some evidence on technical trading.

²⁴Since volume series have an upward trend this does depend on the detrending method used, but the general persistence of volume remains relatively robust to various techniques.

²⁵Returns are the residuals from regression 18, $\hat{\epsilon}_t$, for the artificial market data.

information. The dashed lines represent the 25/75 quantiles for the 25 runs, and the solid line is the median. In both cases the bits appear to be settling down to levels below 5 percent. However, there is some weak indication of a continuing increase in the fast learning case.

Figure 5 presents a similar estimate of the fraction of bits set, but now the average is constrained to the 4 moving average technical trading bits. It shows a much larger fraction of bits set for the fast learning case, and what appears to be an upward trend. This shows that there are almost 4 times as many rules conditioning on the technical trading bits in the fast learning case as there are in the slow case. Since there is an upward trend it would be interesting to know what happens to the technical bits further out in the future. It is too costly in terms of computer time to run a lot of long length runs, but one was run out to 1,000,000 time steps. Figure 6 shows the fraction of technical bits that are set in this one long run. It is clear from this picture that the bits eventually stop their increase, but there also appear to be large swings over time in magnitude of technical trading behavior. There are swings from as low as 10 percent, to highs of over 25 percent. This suggests a world in which technical trading is not a consistently useful strategy for agents, but one in which it makes large swings through the population over time.

The next figure, figure 7, repeats the experiment in figure 5 for one of the fundamental bits, $p > 3/4$. It again shows a strong persistence in the use of fundamental information in the fast learning case, and not in the slow learning case. This is the same bit that was used in the regressions in table 5 which showed its usefulness in forecasting future price movements in the fast learning case alone. In the homogeneous REE none of these pieces of information should be of any use to the individual agents.

The convergence plots presented so far have been concerned with the bitstring components, or conditioning information, part of the forecasting rules. They have ignored the actual linear forecast parameters. The next figure looks at some of the convergence properties of the linear forecast parameters. Figure 8 displays the convergence of the forecast parameter, a , which is the coefficient on lagged $(p + d)$ in the prediction rules. In the homogeneous REE equilibrium this has a value of 0.95 which is marked by a thick black line. The dashed lines are again the 25/75 percent quantiles. The top panel displays the fast learning case, and the second panel, the slow learning case. Both show a slight positive bias, and are different from the REE value. The third panel is the first result presented from the REE example case mentioned in section 3.1.

Here agents are shown the true REE price for each dividend. They don't trade. They just watch what the price should be, and update their forecasts according to the usual GA updating system. This value shows a slight downward bias that should be expected since the least squares estimate of an AR(1) parameter would be downward biased. The bias is not disappearing asymptotically because the squared errors are based on a down weighted rolling horizon,²⁶ and do not extend out over all periods. This feature shows that the rules would be capable of getting close to what they should be in the REE, and something very different is going on in the other cases.

Figure 9 shows the average value of the variance estimate, $\sigma_{i,j}^2$, for all rules. This is a general volatility estimate for all agents. In the homogeneous REE equilibrium this value should be equal to 4 as marked by the dark line in each figure. Once again, both simulations show a slightly upward bias in this value indicating an estimated price variability larger than the theoretical value. The REE example case is right on target in terms of its estimated variance. This once again shows that given the correct answer the agents are able to learn it.

Few of the estimated forecast parameter figures have shown much difference between the fast and slow learning cases. This may seem unusual given the time series evidence presented earlier. One reason for the similarity is that the dispersion of these parameters has not yet been analyzed. They may have a similar central tendency, but differing amounts of dispersion. Figure 10 shows the mean absolute deviation of the variance estimate. This is estimated for each agent, and then averaged across agents, giving a sense of the dispersion in variance estimates across different rules. In the REE these dispersions should be zero since the conditional variance is constant. This figure shows some interesting differences between the fast and slow learning cases represented by the solid, and dashed lines respectively. Initially, the slow learning case shows more dispersion, and falls relatively slowly. However, in the long run it falls to a point where there is much less dispersion over the rules. It ends up with less than half the mean absolute deviation than in the fast learning case. The slow convergence at first is probably due to the fact the rules are getting updated more slowly. In the long run this clearly shows that there is something different going on in the fast learning cases.

The final figure, figure 11, demonstrates the distinction between individual agent versus aggregate phenomena. An interesting fact about financial markets is that volume and volatility are positively correlated.

²⁶Equation 15.

This is a little puzzling given the theoretical setup of agent demands used in this paper. Given that volatility enters into the denominator, it seems likely that when volatility is higher actual trading volume will be lower. This is not directly obvious, but looking at equation 4 it is clear that increases in estimated volatility, all things equal, will reduce the variability in share demand from shocks to future expected prices given current prices. Therefore, the desired reshuffling of shares will be greater when volatility is lower. This is very close to an argument that traders get scared when volatility is high, and don't want to change their positions much in response to new information.

This conjecture is confirmed in figure 11 which plots the cross correlation of trading volume with the average estimated volatility, the average value of $\sigma_{i,j}^2$, over all rules and agents. This shows a strong negative contemporaneous correlation. Since the volatility estimate is highly persistent by design, the crosscorrelation between lagged volatility and volume is also strongly negative. The value begins to move positive since there should be no direct connection from trading volume to future volatility estimates. This pattern appears to confirm what one might think should be the relation between volume and volatility at the micro level. This does not need to hold at the macro level since this figure completely ignores the possible dispersions in volatility estimates across agents. Volatile periods might also be ones in which the dispersion in volatility estimates over agents is quite large, driving large amounts of trading volume. This figure shows how the micro/macro pictures in a market can be quite different, and thinking in terms of a representative agent can fool one in terms of what dynamics are possible.

3.4 Result Summary

The results show two distinct outcomes from the simulated market. One of these, the slow learning case, looks very close to what should be going on in the linear REE equilibrium. In these runs, lagged prices contain all the information needed for forecasting, expected returns, kurtosis, and volatility persistence are also closer to what they should be in the benchmark case. Also, agents clearly are ignoring extraneous information that is offered to them as they should. This is a good indication that this market structure is capable of learning the REE under the right conditions.

In contrast, changing only one parameter alters the results dramatically. Increasing the frequency of rule

updating causes the market to change. The market appears to not settle down to any recognizable equilibrium. In the time series we observe many features common to real markets, including weak forecastability, volatility persistence, and higher expected returns. Correlations between volume and volatility appear to line up with some features of individual stock returns. Finally, the behavior of the agents themselves is consistent with the time series, in that they endogenously form rules that build on the information that the time series tests find useful.

It would not be all that surprising to see such a dramatic change in outcomes if many different parameters were changed. What is surprising is that a single parameter causes such a large change. Moreover, this is a crucial parameter in terms of the economics behind the agent and market behavior. If agents update their rules at longer horizons, the relative value of rules will be based on longer horizon features of the time series. In the shorter horizon updating, it seems more likely that agents might be able to update rules that perform well against the current pool of other agent strategies at the current time. Given that these strategies themselves are changing over time, at longer horizons the idiosyncratic parts of individual behavior may wash out, leaving only the fundamental economic structures which pull the agents into the REE. This logic stresses the importance of how stationary the environment is, and how much the agents believe in a stationary versus changing world view.²⁷

This sort of short term aspect to decision making has some connections to the growing area of behavioral finance.²⁸ Many parts of this market such as rules of thumb, and decision making heuristics have a definite behavioral aspect, but there are some key differences. One of the most important is that traders are behaving optimally subject to their information representations. In many behavioral models there are obvious biases and mistakes that could be corrected to improve economic welfare. Here, there are no such immediate biases, and in this continually changing world it is not obvious that the agents might be following the best strategies possible, given their bounded models. This does ask two key questions. What would happen if the agents were allowed to build ever more complicated models, and what if the speed of learning itself were put under some kind of evolutionary dynamic? Both of these are difficult questions, which cannot be answered in the current market structure.

²⁷A recent paper looking at this issue using a different framework is Benink & Bossaerts (1996).

²⁸See Thaler (1991) and Thaler (1992) for an overview of this large area.

4 Conclusions and Future Research

These results show that the artificial stock market is able to replicate certain time series features from real markets. Among these are predictability, volatility, and volume relations both static and dynamic. Also, agent beliefs are consistent with the time series world that they have constructed. In other words, they concentrate on rules using those pieces of forecasting information which appear significant in the time series tests, but should not be important inside a homogeneous linear REE. The artificial market is also capable of generating what appears to be near equilibrium REE behavior under certain circumstances. These important benchmarks showed that when the speed of learning was slowed down, forcing agents to look at longer horizon features, the market approached the REE in terms of time series, and agent rule features. This shows not just the importance of this one parameter, but the importance of forecast horizons in general, and events that trigger learning and rule updating.

The results presented are purely qualitative. There is no intention to try and calibrate magnitudes to real data. Time horizons have been purposely left vague. In future work it will be necessary both to take strong stands on what the horizon lengths are, and quantitatively line up with actual data. Another aspect of the empirical work, is that the market shows little ability to address questions about excess volatility. One possible reason for this is that the rule reinforcement system uses continuous updating. Forecasted dividends are revealed in the next period and rule accuracy is updated. There are no long periods in which the agents only see price information along with small hints about actual fundamentals as there are in real markets.

One major restriction in this market is the homogeneity of the agents. Apart from divergences in their rule sets they are identical. This can be viewed as a strength in terms of most of the results obtained. It is important that the time series results were derived without having to resort to using any type of ad hoc differences across our agents. The results are made much stronger since they were subject to this constraint. However, in terms of future work, these are restrictions that may be relaxed.

This computer experiment should be viewed as part of a growing area of prototypes for research into the dynamics of markets. Being early versions, these prototypes are probably still far from more generally accepted analytic work. As they proceed over time this may change. One of the most important issues for the future is simplification. It is tempting in a computer-modeling world to add everything, and to allow

model complexity to increase. The analytic world provides bounds keeping this expansion in check because complex models quickly become intractable. Computer simulated markets with adaptive learning agents still contain a huge number of parameters whose values are not well understood either in terms of human behavior, or in terms of their impact on the computer runs.²⁹ However, more traditional models are also being forced to get more complex to keep up with the growing list of empirical financial market puzzles. There may come a time when we are forced to choose between a very complex representative agent situated in a world subject to a complicated stochastic structure, versus relatively simple learning agents subject to independent shocks.

²⁹This parameter proliferation has caused many economists to be a little weary of the potentials of these models, (Sargent 1993). One potential solution to this may be the added restriction of using experimental data as guidance for the learning mechanisms used.

References

- Admati, A. R. (1991), 'The informational role of prices', *Journal of Monetary Economics* **28**, 347–361.
- Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.
- Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. & Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, in W. B. Arthur, S. Durlauf & D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15–44.
- Bäck, T. (1996), *Evolutionary Algorithms in Theory and Practice*, Oxford University Press, Oxford, UK.
- Benink, H. & Bossaerts, P. (1996), An exploration of neo-austrian theory applied to financial markets, Technical report, California Institute of Technology, Pasadena, CA.
- Bollerslev, T., Chou, R. Y., Jayaraman, N. & Kroner, K. F. (1990), 'ARCH modeling in finance: A review of the theory and empirical evidence', *Journal of Econometrics* **52**(1), 5–60.
- Bray, M. (1982), 'Learning, estimation, and the stability of rational expectations', *Journal of Economic Theory* **26**, 318–339.
- Brock, W. A., Dechert, W. D., Scheinkman, J. A. & LeBaron, B. (1996), 'A test for independence based on the correlation dimension', *Econometric Reviews* **15**(3), 197–235.
- Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.
- Brock, W. A., Hsieh, D. & LeBaron, B. (1991), *Nonlinear Dynamics, Chaos, and Instability: Statistical Theory and Economic Evidence*, MIT Press, Cambridge, MA.
- Brock, W. A., Lakonishok, J. & LeBaron, B. (1992), 'Simple technical trading rules and the stochastic properties of stock returns', *Journal of Finance* **47**, 1731–1764.

- Brock, W. A. & LeBaron, B. (1996), 'A dynamic structural model for stock return volatility and trading volume', *Review of Economics and Statistics* **78**, 94–110.
- Campbell, J. & Mankiw, G. (1989), Consumption, income, and interest rates: Reinterpreting the time series evidence, in O. Blanchard & S. Fischer, eds, 'NBER Macroeconomics Annual 1989', MIT Press, Cambridge.
- Campbell, J. & Shiller, R. (1988a), 'Stock prices, earnings, and expected dividends', *Journal of Finance* **63**, 661–676.
- Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- Campbell, J. Y. & Shiller, R. (1988b), 'The dividend-price ratio and expectations of future dividends and discount factors', *Review of Financial Studies* **1**, 195–227.
- Chiarella, C. (1992), 'The dynamics of speculative behaviour', *Annals of Operations Research* **37**, 101–123.
- Cochrane, J. (1989), 'The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives', *American Economic Review* **79**, 319–337.
- Conlisk, J. (1996), 'Why bounded rationality?', *Journal of Economic Literature* **34**, 669–700.
- Day, R. H. & Huang, W. H. (1990), 'Bulls, bears, and market sheep', *Journal of Economic Behavior and Organization* **14**, 299–330.
- De Grauwe, P., Dewachter, H. & Embrechts, M. (1993), *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*, Blackwell, Oxford.
- de la Maza, M. & Yuret, D. (1995), A model of stock market participants, in J. Biethahn & V. Nissen, eds, 'Evolutionary Algorithms in Management Applications', Springer Verlag, Heidelberg, pp. 290–304.
- De Long, J. B., Shleifer, A. & Summers, L. H. (1990), 'Noise trader risk in financial markets', *Journal of Political Economy* **98**, 703–738.

- Engle, R. F. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.
- Evans, G. W. (1989), 'The fragility of sunspots and bubbles', *Journal of Monetary Economics* **23**, 297–317.
- Fama, E. F. (1991), 'Efficient capital markets: II', *Journal of Finance* **45**, 1575–1617.
- Fogel, D. B. (1995), *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*, IEEE Press, Piscataway, NJ.
- Frankel, J. A. & Froot, K. A. (1988), Explaining the demand for dollars: International rates of return and the expectations of chartists and fundamentalists, in R. Chambers & P. Paarlberg, eds, 'Agriculture, Macroeconomics, and the Exchange Rate', Westview Press, Boulder, CO.
- Genotte, G. & Leland, H. (1990), 'Market liquidity, hedging, and crashes', *American Economic Review* **80**, 999–1021.
- Goldberg, D. E. (1989), *Genetic Algorithms in search, optimization and machine learning*, Addison Wesley, Reading, MA.
- Grossman, S. (1989), *The Informational Role of Prices*, The MIT Press, Cambridge, MA.
- Holland, J. H. (1992), *Adaptation in Natural and Artificial Systems*, MIT Press, Cambridge, MA.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E. & Thagard, P. R. (1986), *Induction*, MIT Press, Cambridge, MA.
- Karpov, J. M. (1987), 'The relation between price changes and trading volume: A survey', *Journal of Financial and Quantitative Analysis* **22**, 109–126.
- Kocherlakota, N. (1996), 'The equity premium: It's still a puzzle', *Journal of Economic Literature* **34**(1), 42–71.
- Kurz, M. (1994), 'On rational belief equilibria', *Economic Theory* **4**, 859–876.

- Kurz, M. & Beltratti, A. (1995), The equity premium is no puzzle, Technical report, Stanford University, Stanford, CA.
- LeBaron, B. (1998*b*), Technical trading rules and regime shifts in foreign exchange, *in* E. Acar & S. Satchell, eds, ‘Advanced Trading Rules’, Butterworth-Heinemann, pp. 5–40.
- LeBaron, B. (forthcoming 1998*a*), ‘Agent based computational finance: Suggested readings and early research’, *Journal of Economic Dynamics and Control*.
- Lettau, M. (1997), ‘Explaining the facts with adaptive agents: The case of mutual fund flows’, *Journal of Economic Dynamics and Control* **21**, 1117–1148.
- Lettau, M. & Uhlig, H. (1997), Volatility bounds and preferences: A analytical approach, Technical report, CentER for economic research, Tilburg University, Tilburg.
- Levich, R. M. & Thomas, L. R. (1993), ‘The significance of technical trading-rule profits in the foreign exchange market: A bootstrap approach’, *Journal of International Money and Finance* **12**, 451–474.
- Lux, T. (1998), ‘The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions’, *Journal of Economic Behavior and Organization* **33**, 143–165.
- Marengo, L. & Tordjman, H. (1995), Speculation, heterogeneity, and learning: A model of exchange rate dynamics, Technical Report WP-95-17, International Institute for Applied Systems Analysis, Vienna, Austria.
- Margarita, S. & Beltratti, A. (1993), Stock prices and volume in an artificial adaptive stock market, *in* ‘New trends in Neural Computation: International Workshop on Artificial Networks’, Springer-Verlag, Berlin, pp. 714–719.
- Marimon, R., McGrattan, E. & Sargent, T. J. (1990), ‘Money as a medium of exchange in an economy with artificially intelligent agents’, *Journal of Economic Dynamics and Control* **14**, 329–373.
- Mitchell, M. (1996), *An Introduction to Genetic Algorithms*, MIT Press, Cambridge, MA.

- Rieck, C. (1994), Evolutionary simulation of asset trading strategies, in E. Hillebrand & J. Stender, eds, 'Many-Agent Simulation and Artificial Life', IOS Press.
- Routledge, B. R. (1994), Artificial selection: Genetic algorithms and learning in a rational expectations model, Technical report, GSIA, Carnegie Mellon, Pittsburgh, Penn.
- Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.
- Simon, H. (1969), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.
- Steiglitz, K., Honig, M. L. & Cohen, L. M. (1996), A computational market model based on individual action, in S. H. Clearwater, ed., 'Market-Based Control: A paradigm for distributed resource allocation', World Scientific, New Jersey, pp. 1–27.
- Sweeney, R. J. (1986), 'Beating the foreign exchange market', *Journal of Finance* **41**, 163–182.
- Sweeney, R. J. (1988), 'Some new filter rule tests: Methods and results', *Journal of Financial and Quantitative Analysis* **23**, 285–300.
- Taylor, M. & Allen, H. (1992), 'The use of technical analysis in the foreign exchange market', *Journal of International Money and Finance* **11**(3), 304–14.
- Taylor, S. J. (1992), 'Rewards available to currency futures speculators: Compensation for risk or evidence of inefficient pricing?', *Economic Record* **68**, 105–116.
- Thaler, R. H. (1991), *Quasi Rational Economics*, Russell Sage Foundation, New York, NY.
- Thaler, R. H. (1992), *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, The Free Press, New York, NY.
- Woodford, M. (1990), 'Learning to believe in sunspots', *Econometrica* **52**(2), 277–307.
- Youssefmir, M. & Huberman, B. A. (1997), 'Clustered volatility in multiagent dynamics', *Journal of Economic Behavior and Organization* **32**, 101–118.

Zeeman, E. (1974), 'On the unstable behavior of stock exchanges', *Journal of Mathematical Economics* **1**, 39–49.

Parameter	Simulation Value
γ	0.5
d	10
r_f	0.10
ρ	0.95
σ_μ^2	0.07429
σ_{p+d}^2	4.00
f	6.3333
e	16.6880
$(1 - \rho)((1 + f)d + e)$	4.5011
a range	[0.7, 1.2]
b range	[-10.0, 19.0]

Table 1: Parameter Values

Bit	Condition
1	Price*interest/dividend $> 1/4$
2	Price*interest/dividend $> 1/2$
3	Price*interest/dividend $> 3/4$
4	Price*interest/dividend $> 7/8$
5	Price*interest/dividend > 1
6	Price*interest/dividend $> 9/8$
7	Price $>$ 5-period MA
8	Price $>$ 10-period MA
9	Price $>$ 100-period MA
10	Price $>$ 500-period MA
11	On: 1
12	Off: 0

Table 2: Condition Bits

Rule	Matches
$(1, \#, 1, \#, 0)$	$(1, 0, 1, 0, 0)$
	$(1, 1, 1, 0, 0)$
	$(1, 0, 1, 1, 0)$
	$(1, 1, 1, 1, 0)$

Table 3: Matching Examples

Description	Fast Learning	Slow Learning
Std.	2.147 (0.017)	2.135 (0.008)
Excess Kurtosis	0.320 (0.020)	0.072 (0.012)
ρ_1	0.007 (0.004)	0.036 (0.002)
ARCH(1)	36.98 [1.00]	3.159 [0.44]
ρ_1^2	0.064 (0.004)	0.017 (0.002)
BDS	3.11 [0.84]	1.28 [0.24]
Excess Return	3.062 (0.050)	2.891 (0.028)
Trading Volume	0.706 (0.047)	0.355 (0.021)

Table 4: Residual Summary Statistics

Means over 25 runs. Numbers in parenthesis are standard errors estimated using the 25 runs. Numbers in brackets are the fraction of tests rejecting the no ARCH, or independent identically distributed null hypothesis for the ARCH and BDS tests respectively at the 95% confidence level.

Description	Fast Learning	Slow Learning
$MA(5)$	0.009 (0.013)	-0.008 (0.007)
$MA(500)$	0.074 (0.014)	-0.025 (0.015)
$rP/D > 3/4$	-0.443 (0.104)	0.050 (0.093)

Table 5: Forecasting Regressions

Means over 25 runs. Numbers in parenthesis are standard errors estimated using the 25 runs.

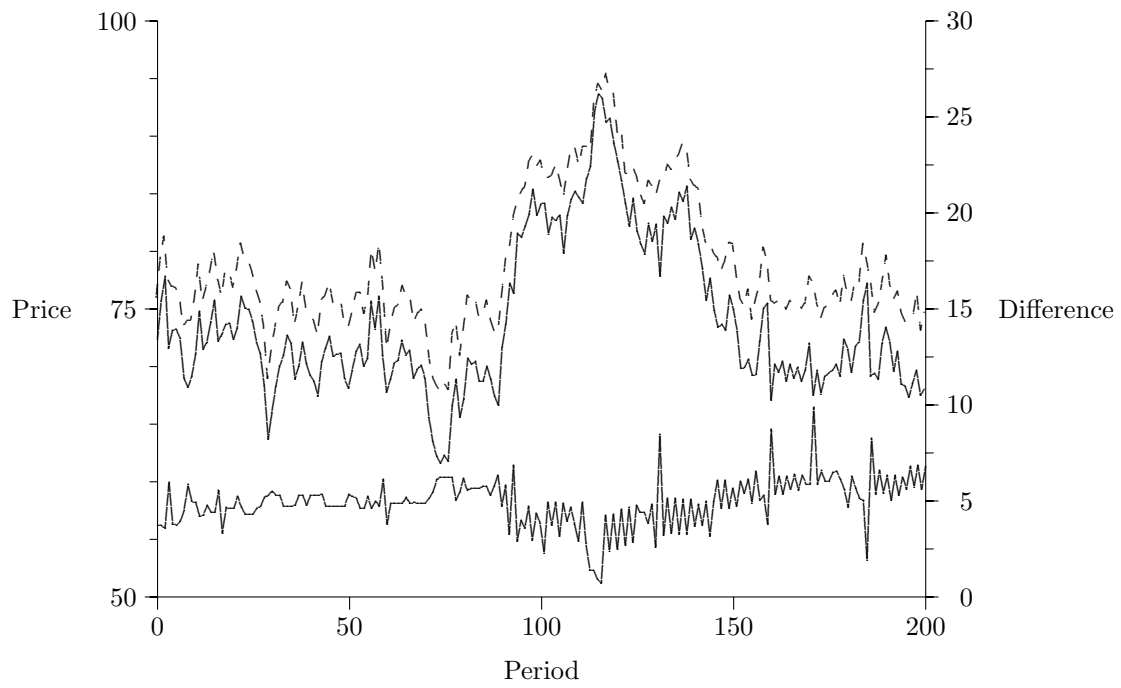


Figure 1: Price Series

Starting period = 250,000, solid line = actual price, dotted line = theoretical, bottom line = difference.

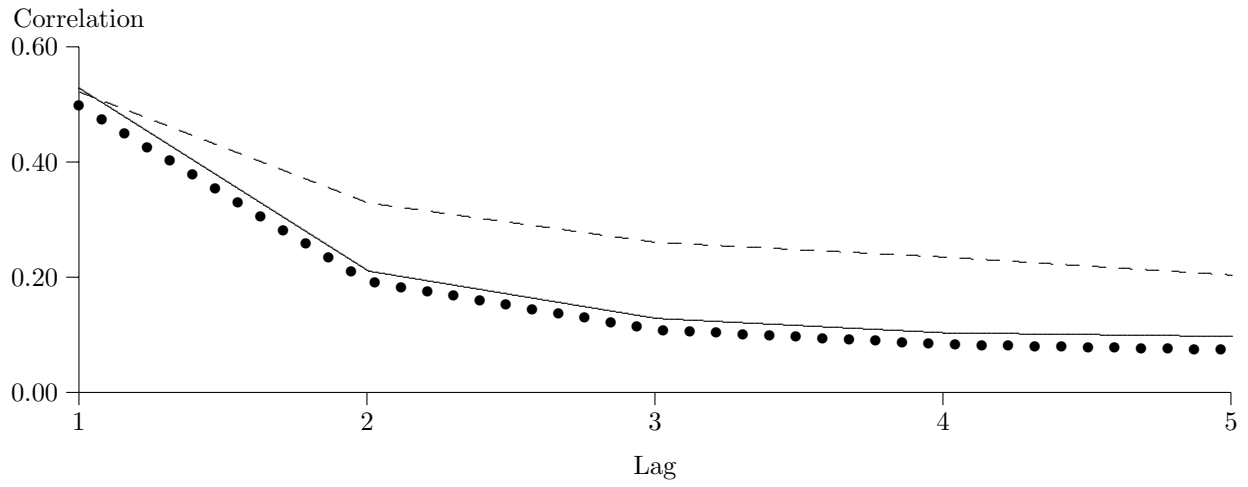


Figure 2: Volume Autocorrelations

Fast learning mean = solid, slow learning mean = dots, IBM daily = dashes.

Crosscorrelation

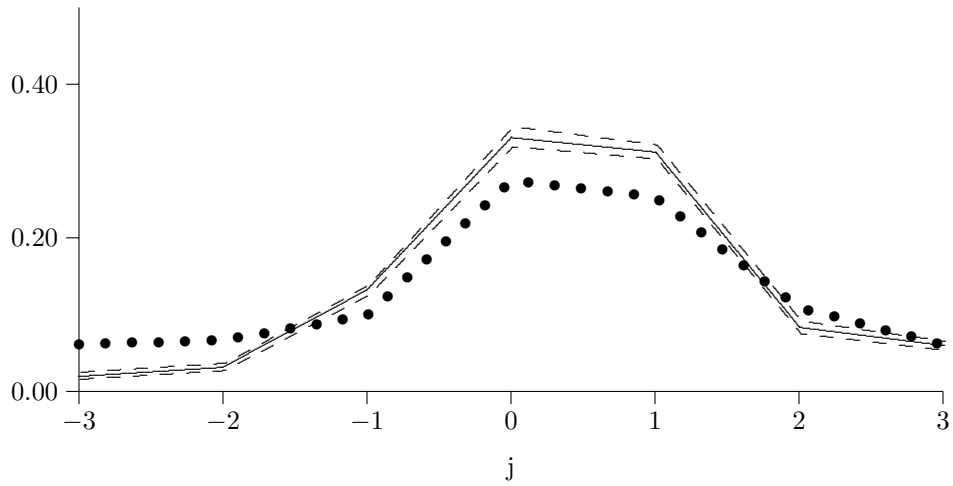


Figure 3: Crosscorrelations: Squared Returns and Trading Volume

Correlation of squared returns at $(t + j)$ with volume at t . Solid line = mean fast learning, dashes = 95% confidence intervals using cross section, dots = IBM Daily.

Fraction of Bits Set

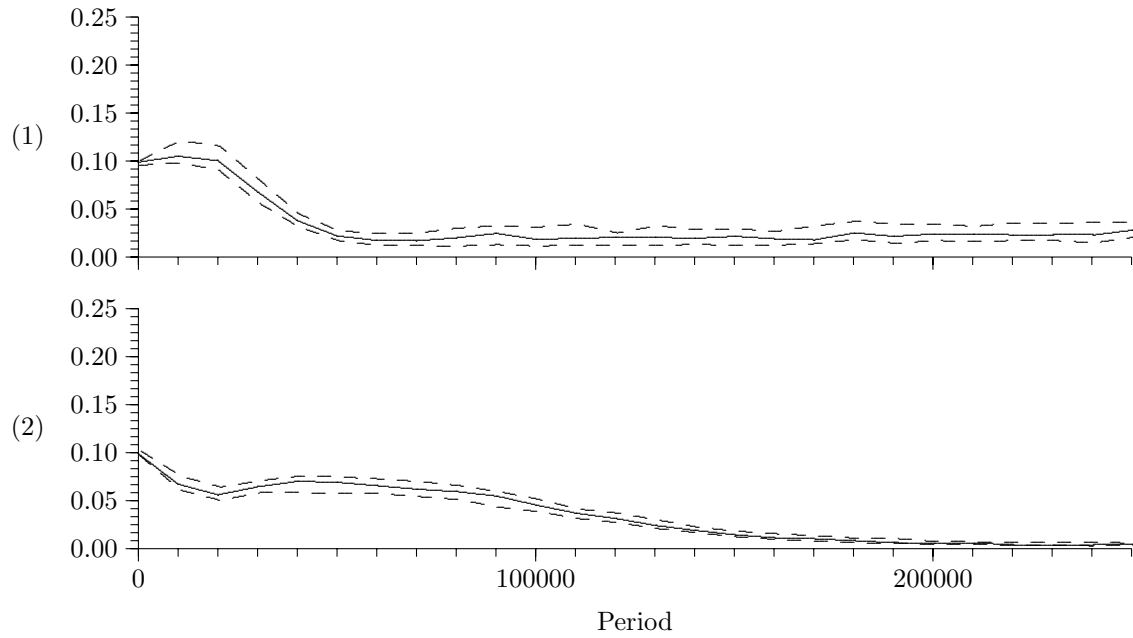


Figure 4: Average Bits Set

(1) Fast Learning, (2) Slow Learning. Median and 25/75 Quantiles: 25 runs.

Fraction of Bits Set

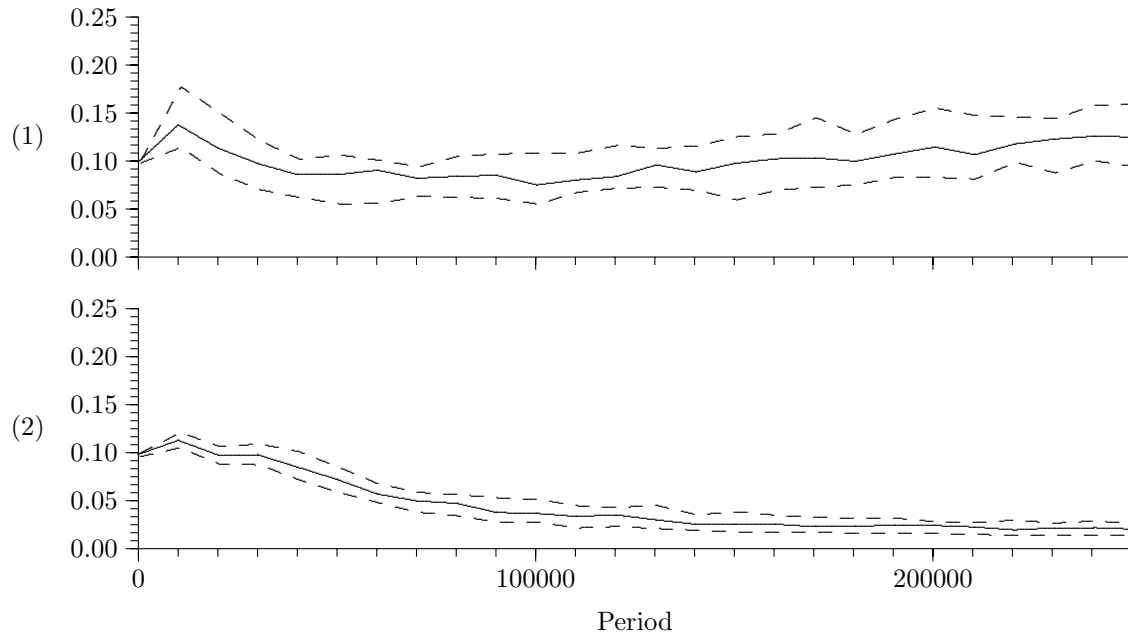


Figure 5: Technical Trading Bits
(1) Fast Learning, (2) Slow Learning. Median and 25/75 Quantiles: 25 runs

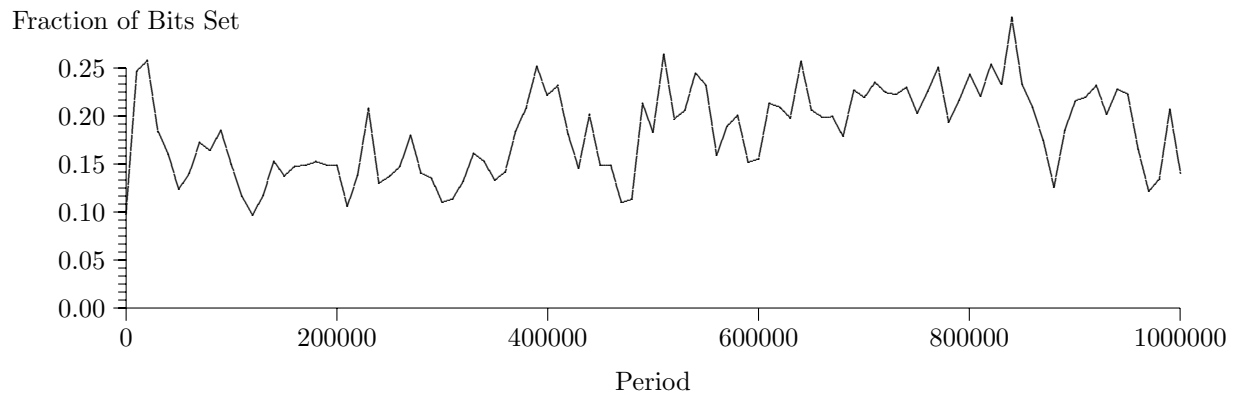


Figure 6: Long Run experiments: Technical trading bits

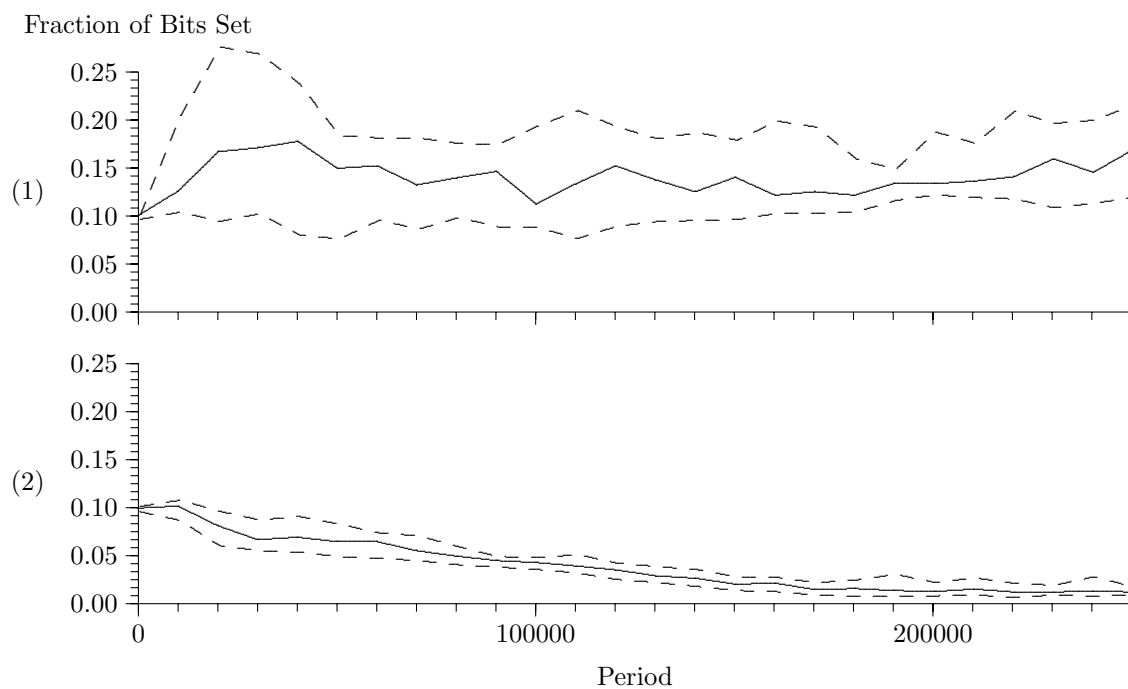


Figure 7: Fundamental Trading Bit ($\frac{prf}{d} > \frac{3}{4}$)

(1) Fast Learning, (2) Slow Learning. Median and 25/75 Quantiles: 25 runs.

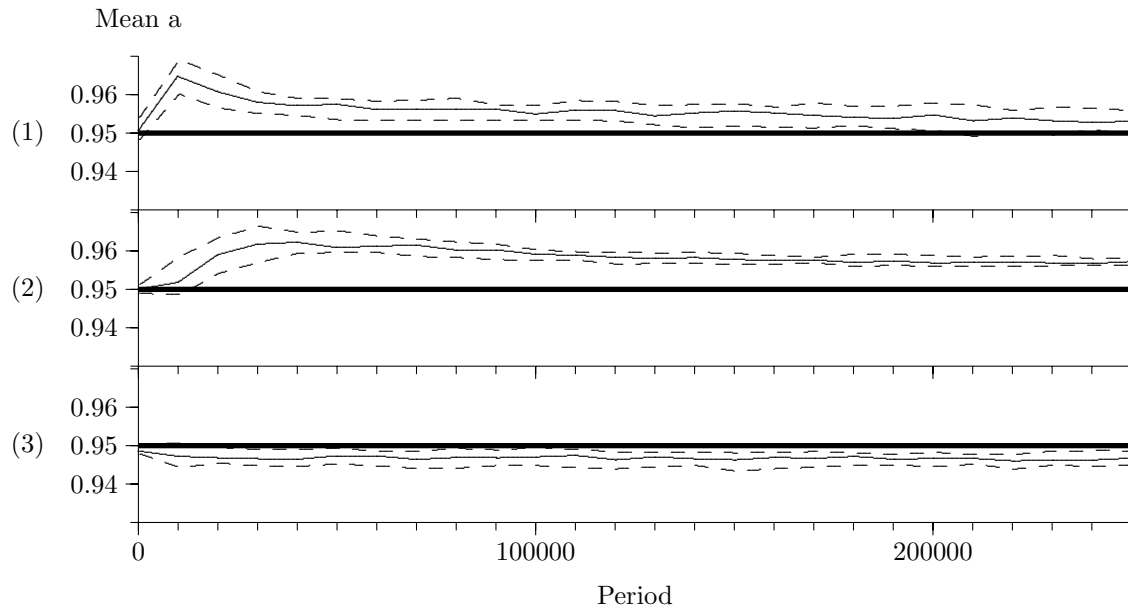


Figure 8: Mean Forecast Parameter a

(1) Fast Learning, (2) Slow Learning, (3) REE Example. Median and 25/75 Quantiles: 25 runs.

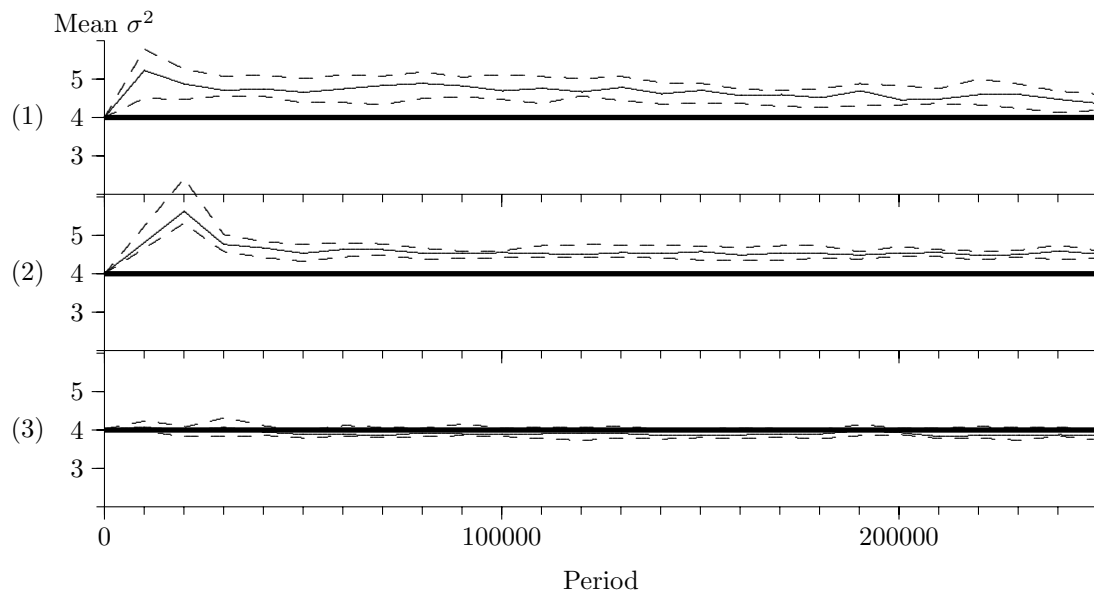


Figure 9: Mean Variance Estimate σ^2

(1) Fast Learning, (2) Slow Learning, (3) REE Example. Median and 25/75 Quantiles: 25 runs.

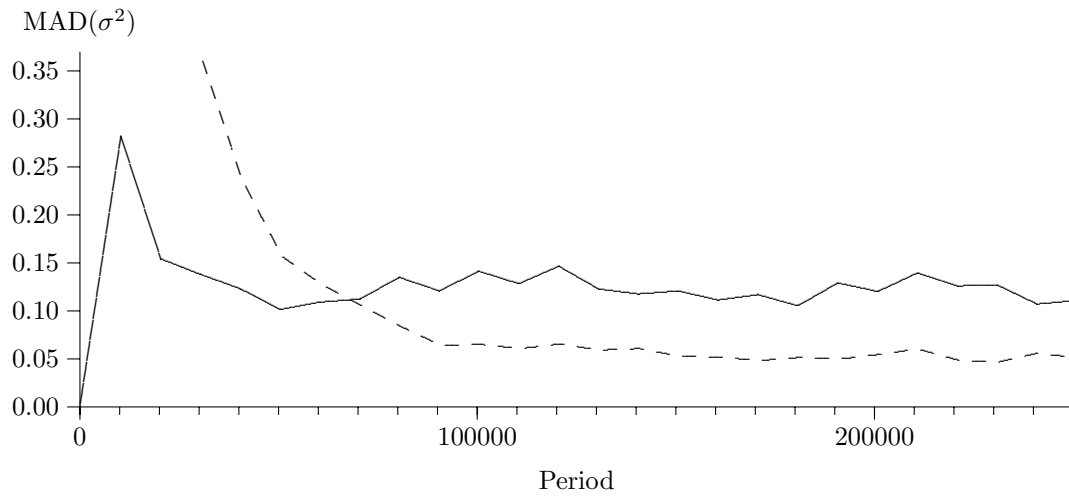


Figure 10: Mean Absolute Deviation for Variance Estimate σ^2
Fast Learning: solid, Slow Learning: dashes

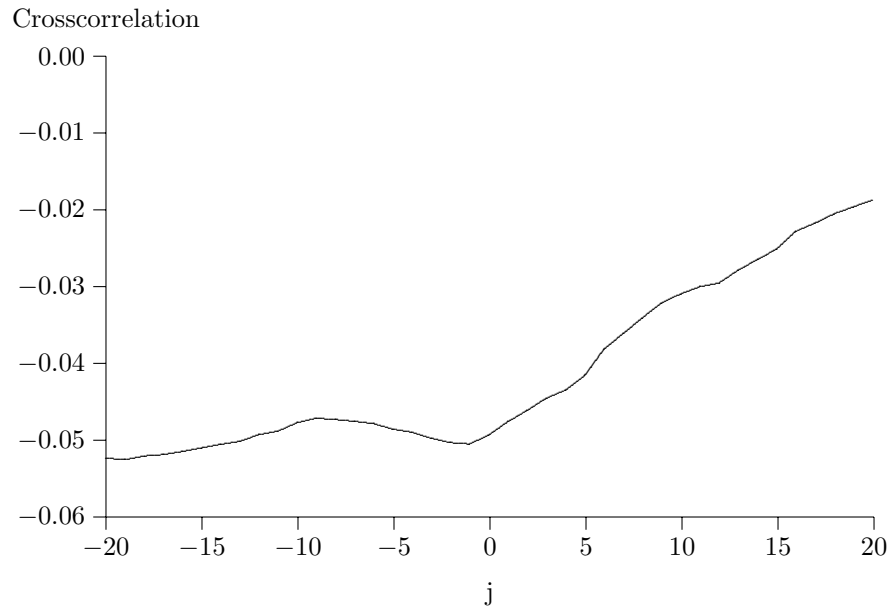


Figure 11: Crosscorrelation: Volume(t) and Individual Volatility(t+j)