Experiments in Evolutionary Finance *

Blake LeBaron
Department of Economics
University of Wisconsin - Madison
1180 Observatory Drive
Madison, WI 53706
(608) 263-2516
blebaron@facstaff.wisc.edu

August 1995
Revised: November 1995

Abstract

Combining ideas from evolution and learning to understand empirical puzzles in financial markets is a growing area of interest in economic research. This paper provides a short survey of some of the ongoing work in this area with special attention paid to computational models relying on artificial intelligence methods. Also, specific experiments will be analyzed using the Santa Fe Artificial Stock Market. The conclusions tie some of the results from these very different modeling approaches together, and suggest paths for future research.

*The author is grateful to the Alfred P. Sloan Foundation and the University of Wisconsin Graduate School for support. Part of this work represents an on going project begun at the Santa Fe Institute along with W. B. Arthur, J. Holland, R. Palmer, and P. Taylor.

1 Introduction

This paper summarizes some of the continuing work on evolution and learning as applied to financial markets. Papers inspired by the connections between selection and observed rational behavior are not uncommon in economic modeling and can trace some of their origins to Knight, Schumpeter, Friedman, and Alchian, along with many others.[1] The application of these principles to financial markets remains a relatively small body of work, but recently it has been experiencing a surge of interesting work.

Financial markets are probably some of the best places to explore evolutionary approaches in economics. Notions of survival and fitness appear closely linked to overall wealth accumulation. Populations of traders using different strategies may appear and disappear in time horizons short enough to be studied in available time series. The constant disagreement and struggle to interpret information and the information of others in financial markets creates a setting probably best described with an environment allowing continual change and adaptation as a opposed to the more traditional static equilibrium approach.

Research styles used to attack these problems range from purely analytic to heavily computational. This paper points out some interesting connections between the two research styles, and their results. Special emphasis will be placed on computational modeling approaches along with a brief set of experiments from the Santa Fe Artificial Stock Market .[2]

This paper is probably a good introduction for someone interested in learning, finance, and computational approaches to this field. It is clearly directed at finance, and although there is a large overlap with learning in economics in general, this is not a survey of this field.[3] Section 2 looks briefly at the more analytic approaches. Section 3 goes into more detail summarizing computational models. Section 4 will present some experiments from the Santa Fe Artificial Market, and finally section 5 will close and try to identify common themes coming from both lines of research.

2 Analytic Approaches

The earliest approaches to evolution in financial markets were probably arguments about logarithmic preferences and the long run growth of wealth. These appeared in Kelley (1956) and Breiman (1961). They presented more of a normative argument for why agents should strive for fit strategies, as opposed to a positive argument for why only the fittest strategies would survive. Their work suggested a special place for agents maximizing the expected growth rates of their portfolios. Also, Friedman (1953), should probably also be included in the early evolutionary framework. While not a formal model of trader behavior, he was one of the first to make the point that foolish trading strategies should eventually disappear from the trading environment. We will come back to Friedman's arguments later in this paper.

Many loose ends in this literature are tied together in Blume & Easley (1990). They analyze the positive side of evolution in financial markets in great detail, showing that rules maximizing growth rates in wealth are the "most fit" in the evolutionary sense, and traders using these will dominate all others. One of the main results in Blume and Easley is the formal delinking of rationality from

[1] See for example, Knight (1921),Schumpeter (1934),Alchian (1950), and Friedman (1953).

[2] This is part of an ongoing project with other researchers at the Santa Fe Institute, W. B. Arthur, J. Holland, R. Palmer, and P. Taylor, Arthur *et al.* (Forthcoming).

[3] Readers should look to Sargent (1993) for applications to more general economic problems.

wealth fitness. They show in several different cases that rationality is not selected for in their evolutionary process.[4] Although operating in a slightly different framework, Figlewski (1978), was also an early contribution to this area. Figlewski specifically considers wealth dynamics in the convergence of a market to efficiency, and how the differing price impact of more and less wealthy traders can slow this dynamic significantly.

Although not exactly evolutionary, models of general learning in financial markets are closely related. The work of Bray (1982) is probably the most important early piece which gives a mechanism for simple learning to converge under specific conditions. Bray also points out the important problem that in many learning contexts the learners have misspecified models during the entire learning process. Whether we believe such a dichotomy between our agents and data can or should survive remains an open question. Recently, Bossaerts (1994) has gone further by finding that the learning process of different agents, if sufficiently slow, can make financial data sufficiently nonstationary to distort inferences made in time series studies. He relates this to differences coming from time series and cross sectional studies which find deviations from efficiency in the first case, and good agreement with efficiency in the second. It is possible that nonstationarities coming from the learning process may be responsible for this apparently inconsistent results.

In Timmermann (1995) the feedback of the price process on traders trying to solve the hard problem of forecasting long run dividend growth induces a pattern of volatility persistence which is similar to that in actual financial data (Bollerslev *et al.* (1990)). Several other papers that generate activity bursts similar to those found in actual financial markets are Brock & LeBaron (1993), Youssefmir & Huberman (1995), Grannan & Swindle (1994). Papers in this style might be called "GARCH generators" after the GARCH and ARCH processes (Engle (1982), Bollerslev (1986)) use to model changing volatility. They all have a common theoretical mission of generating persistent volatility from the input of independent shocks. These will be discussed further after the computational approaches are presented.

3 Computational Models

There is a broad range of approaches to computational modeling in evolutionary finance. To keep these models sorted out it is useful to break them up according to several criteria including the economic environment, learning representations, and trading mechanisms. The economic environment refers to the actual economic trading situation, the types of assets, and the objective functions of the various agents. It can also refer to the way in which economic objectives and payoffs are converted into evolutionary pressures. Several of the examples cited here use familiar structures such as overlapping generations, and standard heterogeneous information frameworks. Other examples use different setups which are harder to map into standard models for financial markets.

The way information is processed and stored by agents is a crucial part of any learning environment. This is just as true for financial models as it is for other applications in artificial intelligence. Modern approaches from computer science such as neural networks, genetic algorithms, and classifiers have been applied to financial problems, along with more traditional methods such as least squares learning. In all cases it is important to have a precise knowledge of what domain the agents knowledge lies in, what types of equilibria lie in that domain, and how the agents move in this this domain by updating beliefs.

[4]This result is a suggested in a special case by DeLong *et al.* (1992).

Finally, in financial markets the actual trading mechanism is crucial. Clearing markets with large numbers of interacting agents is not a trivial task, and the mechanism for doing this may greatly affect results. This might be a desirable feature if it is this clearing mechanism itself that is under study, but it may not be useful if the experimenter is after broader properties of learning.

Several papers operate within a well defined simple economic modeling framework using learning to explore both evolutionary and stability properties of well defined economic environments. In Lettau (1993) an optimal portfolio decision problem is analyzed. Agents trade in a market where they can choose how much of a risky asset to hold each period. The current price does not depend on their demands. Agents' strategies take a linear form in the price, $s = a + bp$. Learning is represented by encoding parameters from the strategy into a bitstring, and evolving bitstrings using a genetic algorithm (GA). Selection is made based on utility payoffs from a sequence of draws of the random asset. In this benchmark model Lettau shows there are some interesting biases in the strategies. The evolving agents tend to take riskier strategies than optimal. He also shows that the continual updating of portfolio weights can replicate certain data on mutual funds where agents are continually changing portfolio composition.

Arifovic (Forthcoming) looks at an overlapping generations model for exchange rates taken from Kareken & Wallace (1981). Agents' policy decisions about consumption and savings are again encoded into a bitstring for GA manipulation. A population of agents is evolved through time using observed utility as the selection criteria. Arifovic finds that the exchange rate does not settle down to any of the known equilibria in the model, but continues to bounce around. She shows that this result is similar to other results coming from laboratory experiments.

A third paper injecting learning into a well known economic environment is Routledge (1994). He looks at the Grossman & Stiglitz (1980) framework for heterogeneously informed agents trading an asset. In this model a noisy signal is available which gives agents information about future dividends. This signal is costly, and agents endogenously decide whether to purchase it. Uninformed agents must infer the value of the signal from the current price. They essentially build a linear expectation for the future dividend payoff as a function of current price. This expectation function is encoded as a bitstring along with the decision of whether to purchase the signal (one bit). Simulations of this model show some very interesting dynamics. First, increasing the noise in the system (specifically the supply of the risky asset) can cause the system to be more likely to settle to the equilibrium with some proportion of informed and uninformed agents. This is the rational expectations equilibrium studied by Grossman & Stiglitz (1980). However, when noise decreases the learning uniformed agents are likely to jump to being informed before getting good estimates of uninformed forecast parameters. This further exacerbates the problem of uninformed agent learning by reducing the population of uninformed agents experimenting with the GA. This often converges to a situation where almost all agents decide to become informed. Further experiments show that holding the proportion of uninformed agents constant for a certain number of initial periods can allow the system to converge. These experiments show the interesting interactions between learning and behavior when both are intertwined in complicated fashion. In this model it appears that different speeds of learning between informed and uniformed agents drive much of what is going on. Learning faster, the informed agents draw in more converts from the uniformed population, which further slows the learning there.

Other computational papers operate with a more open structure both for the economic model and the information representation. Beltratti & Margarita (1992) use a neural network for their agents' representation of a trading world in which agents trade bilaterally. They bump into each other at random when a price is proposed, and decide if their respective valuations are amenable to trade. They have agents of differing levels of sophistication, which is determined by the complexity of their networks. Costs can be paid to improve network complexity if desired. Even in this very different setup they observe fluctuations between the fraction of smart and dumb agents in their model. When the price settles down it no longer pays to be smart, and dumb agents dominate, but this leads to price instabilities bringing back smart agents.

In Rieck (1994) the strategies of agents are explicitly modeled as functions of current and past market information. The representative strategies are inspired by actual trader strategies such as moving average trading rules. The agents follow well defined parameterized demand functions which yield limit orders in price and quantity. Evolution is allowed over both different types of demand functions, and the parameters in the functions based on wealth levels. The market is cleared simply by calculating the aggregate supply and demand functions from the limit orders as is done in many experimental markets. Rieck finds that in his simulations it is difficult for technically oriented strategies to be evolutionarily removed from the population.

Marengo & Tordjman (1995) use a classifier based system to model open ended belief formation in a market. In this framework, classifier systems map features of the market information into actions of buying and selling. Agents follow the behavior of their matched classifier rule. Rules of this type can be used to map many features into actions, and they will be detailed further below. The market is cleared using a price adjustment mechanism which allows excess demand, and adjusts the price to try and reduce this excess demand. In their experiments they observe, as in the other cases, a market which does not really settle down, and often moves through apparent changing regimes.

4 Experiments From the Santa Fe Stock Market

4.1 The Environment

The Santa Fe Stock Market, which is an ongoing project of a team of Santa Fe Institute researchers, is outlined in detail in Arthur *et al.* (Forthcoming). Also, the general process of learning and induction in economics and finance is described in more detail in Arthur (1994) and Arthur (1995). This model tries to combine both a well defined economic structure in the market trading mechanisms, along with inductive learning using a classifier based system. This section gives a brief outline of the market structure along with a few experiments.

The market setup is simple and borrows much from existing work such as Bray (1982), and Grossman & Stiglitz (1980). In this framework, one period, myopic, constant absolute risk aversion utility,[5] CARA, agents must decide on their desired asset composition between a risk free bond, and a risky stock paying a stochastic dividend. The bond is in infinite supply and pays a constant interest rate, r. The dividend process is a well defined stochastic process. For these experiments the dividend will follow an AR(1) process,

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

where ϵ_t is gaussian, independent, and identically distributed, and $\rho = 0.99$ for all experiments. It is well known that under CARA utility, and gaussian distributions for dividends and prices, the

[5] This utility function is given by $u(x) = \frac{-1}{\gamma} e^{-\gamma x}$ where γ is the coefficient of absolute risk aversion.

demand for holding shares of the risky asset by agent i, is given by,

$$s_{t,i} = \frac{E_{t,i}(p_{t+1} + d_{t+1}) - p(1+r)}{\lambda \sigma^2_{t,i,p+d}}$$

where p_t is the price of the risky asset at t. $\sigma^2_{t,i,p+d}$ is the conditional variance of $p + d$ at time t, for agent i. λ is the coefficient of absolute risk aversion, and $E_{t,i}$ is the expectation for agent i at time t. [6] [7] Assuming a fixed number of agents, N, and a number of shares equal to the number of agents gives,

$$N = \sum_{i=1}^{N} s_i$$

which closes the model.

In this market there is a well defined linear homogeneous rational expectations equilibrium (REE) in which all traders agree on the model for forecasting future dividends, and the relation between prices and the dividend fundamental. An example of this would be

$$p_t = b + a d_t.$$

The parameters a and b can be easily derived from the underlying parameters of the model by simply substituting the pricing function back into the demand function, and setting it equal to 1, which is an identity and must hold for all d_t.

At this point, this is still a very simple economic framework with nothing particularly new or interesting. Where this breaks from tradition is in the formation of expectations. Agents' individual expectations are formed using some of the learning algorithms that have already been mentioned. Specifically, agents use a classifier system to try to determine the relevant state of the economy, and this in turn leads to a price and dividend forecast which will go into the demand function.

The classifier is a modification of Holland's condition-action classifier, Holland (1992), which is called a condition-forecast classifier. It maps current state information into a conditional forecast of future price and dividend. Current market information is summarized by a bitstring, and each agent possesses a set of classifier rules which are made up strings of the symbols, 1, 0, and #. 1 and 0 must match up with a corresponding 1 or 0 in the current state vector, and # represents a wild card which matches anything. For example, the rule 00#11 would match either the string 00111, or 00011. An all # rule would match anything. In standard classifier systems there is a determination made on which is the strongest rule depending on past performance, and the rule then recommends an action. Here, the strongest rule maps into a real vector of forecast parameters, a, b which the agent uses to build a conditional linear forecast as follows,

$$E_{t,i,j}(p_{t+1} + d_{t+1}) = a_{i,j}(p_t + d_t) + b_{i,j}.$$

[6] $E_{t,i}$ is not the true expectation of agent i at time t. This would depend on bringing to bear all appropriate conditioning information in the market which would include beliefs and holdings of all other agents. Here, it will refer to a simplified price and dividend forecasting process used by the agents. In some benchmark cases this representation will include the true conditional expectation in the set of possible predictors.

[7] This demand function is valid if the shocks around the above expectations are Gaussian. This is true in the rational expectations equilibrium, but it may not hold in many situations. We are assuming that disturbances are not far enough from Gaussian to alter this demand. If they are, it does not invalid the analysis, but it does break the link between this demand function and 1 period CARA utility.

The subscripts i, j indicate that this is for agent i, using rule j out of its rule set. This is not the only way to build forecasts, and agents could be constructed using many other parametric classes of rules and forecasts. However, it does offer some useful features. First, the REE is embedded in the forecasts since the optimal forecast is a linear one. Second, if we force agents to decide on rules, using all information except for p_t in deciding which rule to use, the forecast gives a linear demand function in p_t above, which can be used for easy market clearing. The following is a list of the bits used:

1-7 Price*interest/dividend $> 1/2, 3/4, 7/8, 1, 9/8, 5/4, 3/2$

8 Price $>$ 5-period MA

9 Price $>$ 10-period MA

10 Price $>$ 100-period MA

11 Price $>$ 500-period MA

12 always on

13 always off

14 random (on off)

All matched rules are evaluated according their accuracy in predicting price and dividends. Each rule keeps a record of its squared error according to,

$$e^2_{t,i,j} = \beta e^2_{t-1,i,j} + (1-\beta)((p_{t+1} + d_{t+1}) - E_{t,i,j}(p_{t+1} + d_{t+1}))^2$$

This accuracy is used in two places. First, if multiple rules are matched, then only the most accurate is used. Second, this is an entry into the fitness measure which guides the evolution of new rules over time. Overall, rule strength is defined as,

$$f_{t,i,j} = \frac{1}{e^2_{t,i,j}(1 + cs)}$$

Specificity, s, refers to the number of bits that are set in a rule (not #), and c is a cost for setting bits. What this does is put some evolutionary pressure toward simpler rules. This is most important in judging REE equilibrium. In the true REE, any bit configuration is equivalent to any other as long as the linear parameters are correct. This cost introduces a weak drift toward the all #) bit configuration if we were really in the REE.

Agents perform a very simple estimation procedure to build their estimates of conditional variances,

$$\sigma^2_{t,i,p+d} = \alpha\sigma^2_{t-1,i,p+d} + (1-\alpha)((p_{t+1} + d_{t+1}) - E_{t,i}(p_{t+1} + d_{t+1}))^2,$$

where $E_{t,i}$ is give by the rule that was actually used by agent i, at time t.

The final important part involves the evolution of new rules. Agents are chosen at random according to a poison process with average time lags of 500 periods to update their current forecasting rule sets. The worst performing 15 percent of the rules are dropped out of an agent's rule set, and are replaced by new rules. New rules are generated using a genetic algorithm with uniform crossover and mutation. For the bitstring part of the rules, crossover chooses two fit rules as parents, and takes bits from each parent's rule string at random.[8] Mutation involves changing the individual bits at random. The GA has become a very common search algorithm for solving high dimensional search problems, but many of properties are still under investigation. We use it as a tool to produce novel hypothesis which will enter into the general competition with other rules. Crossing over the real components of the rules is not a commonly performed procedure, and it is done using three different methods chosen at random. First, both parameters, a and b, are taken from one parent. Second, they are each chosen randomly to come from one of the parents. Third, a weighted average is chosen based on strength.

4.2 Experiments

This section presents a few simple experiments from the artificial market. These simulations give only a taste of some of the results from Arthur *et al.* (Forthcoming). Some of the tests and parameters here are changed to demonstrate what is doable in this framework.

The primary goal here is to better understand the stability of the REE equilibrium under this learning dynamic, and to explore how and when agents might be enticed to look at other pieces of information in a learning environment. We will pay particular attention to the technical trading bits that the agents are offered.

In the first run, 50 agents are started up with random initial conditions and bits. The bits are set to 1 or 0 with probability 0.1. Figure 1 shows the fraction of certain key bits set over time. The always on bit is a control bit which should be a useless piece of information. Remember, that with positive bit costs it is not costless to have this bit set in a rule. Its values are similar to many other bits which are quickly shut off in the trading process. The technical trading bits show a clearly different pattern, persisting in relatively high proportions, fluctuating in strength, and even reversing positions over time. It is clear they are not settling to any limit over this time horizon. Since all the information bits are superfluous to forecasting in the REE the significant setting of any of them is an important deviation.

To demonstrate there is something really interesting going on in figure 1, as opposed to spurious bit setting, a second experiment is run. While we do not really have any knowledge about what the properties of a large group of learning agents should be, it is clear that for a single agent playing against a group of REE agents, the optimal policy should be that of the other REE agents. In this experiment a group of agents locked into the REE parameters will trade with a small number of randomly initialized learning agents, to see if the learners will learn the REE in this case. We set the small number of traders to 5, and there are 45 REE traders. The bit distributions for this experiment are given in figure 2. The numbers are percentages of just the learning agents, so they are comparable with figure 1. It is clear in comparing figure 2 with figure 1 that few of the fits are set in the case with lots of REE traders around. There are a few brief spikes when there are bursts of bit setting, but few of these last for very long.

We turn now to the time series generated by the two different experiments. A snap shot is taken well into the experiment, from t=1,000,000 to 1,020,000, for a total of 20,000 time points. This is taken for both experiments. In the REE equilibrium, where price is a linear function of dividends, prices should follow an autoregressive process of order 1. Residuals of an AR(1) are analyzed in

[8] Selection is by tournament selection, which means that for every rule that is needed two are picked at random, and the strongest is taken.

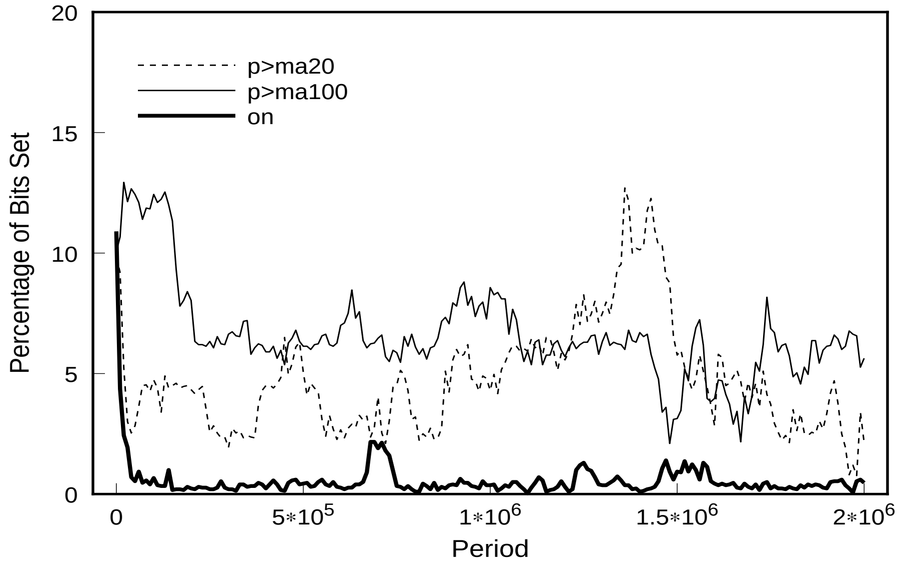

Figure 1: *Percentage of bits set for random startup*

table 1.[9] In general, they are remarkably close to their theoretical properties for both series. They appear close to normal, uncorrelated, with no skewness, or excess kurtosis. The series also show no evidence for nonlinearities, but these tests are not presented.

	Random Start	Fixed Agents
Mean*100	0.000	0.000
$\bar{p}$	78.8	80.2
Std.*100	1.764	1.791
Skew	0.003	-0.002
Kurtosis	2.940	3.008
ρ	0.989	0.991
ACF(1)	0.006	-0.008
ACF(2)	0.007	-0.001
ACF(3)	0.002	0.007
ACF(4)	-0.011	0.000
ACF(5)	0.001	0.012
Bartlett	0.007	0.007

Table 1:
Summary statistics for price residual series. Price residuals from an AR(1). ACF(n) is the autocorrelation at lag n. Bartlett is the asymptotic standard error for the ACF's.

These similarities end in table 2. In the REE equilibrium the current price should be sufficient information for forecasting the future price. No other piece of time t information should help in forecasting the future. However, in the first experiment agents are very interested in using rules conditioning on the technical trading bits. In table 3, p_t is regressed on lagged price, plus the 2

[9] Residuals could also be obtained by using the theoretical AR(1) coefficient (0.99). Results using this method are similar to those presented here.

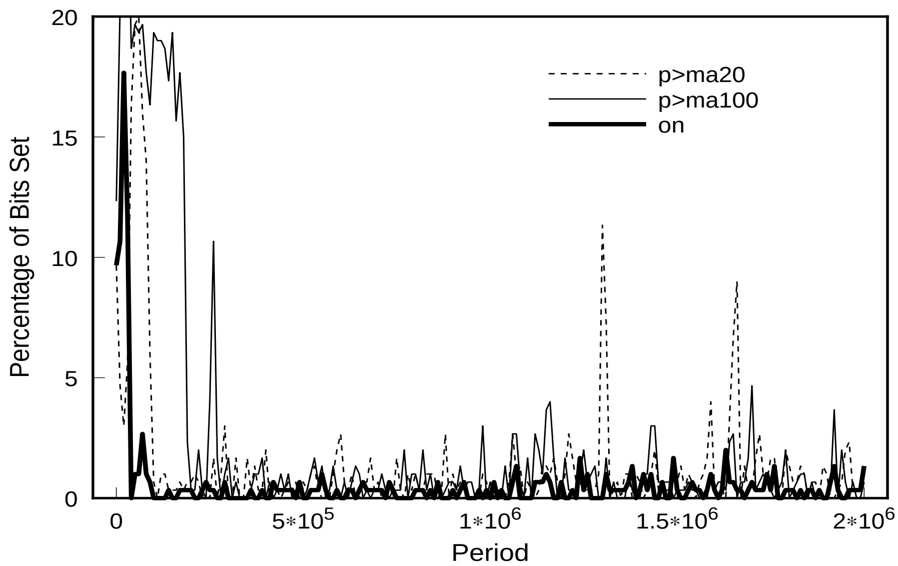

Figure 2: *Percentage of bits set for REE agents and learning agents*

strongest technical trading bits used by the agents in period 1,200,000. The table shows a significant impact of the 100 period moving average technical bit on the p_t forecast. This information should not be useful in the REE equilibrium, and it is not for the fixed agent case where none of the technical bits are significant. One deviation from some of the results seen in actual data is the sign of the trading rule coefficient in table 3. When the price is above the moving average the conditional time series expectation of future prices is lower. This may reflect a trading range type of strategy which emphasizes long range reversals when prices are far out of line. This differs from results in Brock *et al.* (1992) which find evidence for trend following sorts of predictability.

The results from these tables and figures suggest something very different is going on in the two simulations. The second one, with the fixed agents, is behaving according to the predictions of the REE, but the first simulations show some important differences. It looks like technical trading

$$p_t = \beta_0 + \beta_1 p_{t-1} + \beta_2 I_{ma5} + \beta_3 I_{ma100}$$

	Random Start	Fixed Agents
β_0	0.678	0.824
(std)	(0.132)	(0.123)
β_1	0.991	0.991
(std)	(0.001)	(0.001)
β_2	0.032	0.008
(std)	(0.026)	(0.026)
β_3	-0.078	0.038
(std)	(0.031)	(0.031)

Table 2:
Information from technical indicators. I_{ma5} is 1 when the price is above a 5 period moving average, and zero when it is below. I_{ma100} is defined in a similar fashion for a 100 period moving average.

behavior has appeared from the set of predictors in the first case, but not in the second. The presence of a large number of already learned agents has kept it away.

In the final experiment a test is run to see if there is another way to keep technical trading from occurring. In figure 1 it appears that the technical trading activity is far from stationary as the strengths of the different rules vary quite dramatically over time. It is therefore possible that there is an interesting coevolutionary dynamic going on between agents using the different rules. To test this, the final experiment tries to stop this kind of interaction by restricting the bits available to only one price/dividend bit($pr/d > 3/4$), a single technical bit (price>100 day moving average), and an always on bit. Results of this experiment are shown in figure 3. We see little activity of any kind starting up in the various bits. This single experiment clearly needs to be backed up by further explorations, but it is suggestive of a strong interaction between the various types of rules in the trading environment. A simple evolutionary race over a fairly small set of information may not result in a rich emergent dynamics. It is probably true that this is the result of some kind of interaction between the rules using information bits of various types.

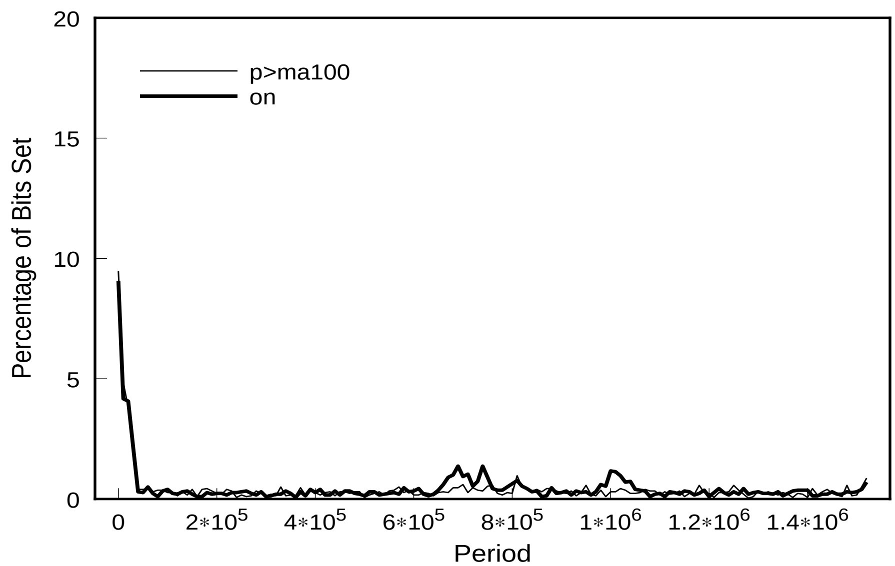

Figure 3: *Percentage of bits set for REE agents and learning agents*

5 Conclusions, Connections, and Future Directions

Even though this field still is quite young, I think it is possible to bring some kind of synthesis to the selection of papers surveyed here, both computational and analytic. There are some interesting common features appearing in many very different approaches.

The first is the issue of instability. Many of the approaches here show some type of instability often leading to periods of quiet followed by periods of activity. While this kind of "volatility clustering" is not common in many traditional theoretical models of asset pricing it almost appears common place in the world of evolution and learning. Few of these papers have been able to exactly target what the underlying causes might be, but some interesting conjectures come from models such as those in Brock & Hommes (1994) and Youssefmir & Huberman (1995). In both these cases agents behave using forecasting rules. In a near REE equilibrium in these models the usefulness of complicated predictors often disappears, shifting the population to simpler rules. However this drive to simplicity, sets the system up for future instabilities, as these rules are less adapted to handle out of equilibrium behavior. There is also a problem in some of these models that near an equilibrium the relative fitness values of different predictors may get very close, scattering agents across many types of predictors, and a certain mass of them away from the REE predictors. In the computational models these instabilities may be the cause for many of the computer simulations finding it hard to reach standard equilibrium solutions.

Another interesting feature is the impact of costs. Models such as Brock & LeBaron (1993), Routledge (1994), Arthur *et al.* (Forthcoming) all have costs playing some role. Both Brock & LeBaron (1993) and Routledge (1994) have costs impacting agents' endogenous decisions to purchase information signals causing a rich set of dynamics with the proportion of informed agents fluctuating over time rather than settling down. In a recent paper, Evans & Ramey (1995) use costs of calculation to show how market fluctuations can be affected. Their agents reduce their forecast horizons when calculation costs are raised. This often increases the chances of speculative bubbles.

Several design questions will probably need to be explored. A general move away from the bitstring/real coding that is often used would be a useful change. These representations do not allow the learning algorithm to really have a clean metric of how far apart rules are in the search space, and search algorithms cannot usefully judge local versus global moves in this space. More relevant coding representations will be a useful change. Also, a move to more detailed market clearing methods will be an important addition. The more ad hoc price setting techniques have the danger of generating rules that detect the price setting methods, and not more fundamental aspects of the markets under study.

One other area that has not been addressed in this paper is the issue of security design. Viewing securities themselves as part of the evolutionary process is also a possibility. There are certainly obvious analogies, that make this sound interesting. It would be impossible to have an option on a future without out having the futures market there in the first place. Securities can only appear when the environment is right for them to make their entrance. Even though the stories are appealing this may be a difficult line of research which will have to combine the literature on optimal security design with the evolutionary processes described here.

A final question that few, if any, of the models have addressed is the speed of learning, and the horizons that agents look at. Many of these models depend critically on some rate of change or mutation in the agent populations. They also depend on how backward looking many of the rules and strength measures are, but in most cases these key parameters are set exogenously by the experimenter. It may be interesting to eventually endogenize these in some way, or at least tune them to reasonable rates of learning. For example, we looked at examples of the artificial market after 1,000,000 time steps had gone by. Is it reasonable to wait this long? Of course, this depends on the amount of time corresponding to a period, but it is a very important question. Also, mutation rates might evolve themselves to some type of optimum.[10] Could agents somehow also adjust their time horizons in some way to adjust to incoming information, or is it possible that they will always will be somewhat slow in their learning and adjustment because of the need to

[10] The work by BenPorath *et al.* (1993) is an approach in this direction.

balance out exploiting a stationary environment while getting ready for a sudden regime shift?

There is clearly much more to be done, and the area is wide open for new entries. Both computational and analytic approaches should be able to complement each other extensively in the search for understanding financial markets as part of an evolutionary process, as opposed to a static equilibrium.

References

Alchian, A. 1950. Uncertainty, evolution, and economic theory. *Journal of Political Economy* **58**, 211–221.

Arifovic, J. Forthcoming. The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy* .

Arthur, W. B. 1994. Inductive behavior and bounded rationality. *American Economic Review* **84**.

Arthur, W. B. 1995. Complexity in economic and financial markets. *Complexity* **1**(1), 20–25.

Arthur, W. B., J. H. Holland, B. LeBaron, R. Palmer and P. Taylor. Forthcoming. Asset pricing under inductive reasoning in an artificial stock market. Technical report, Santa Fe Institute, Santa Fe, NM.

Beltratti, A. and S. Margarita. 1992. Evolution of trading strategies among heterogeneous artificial economic agents. Technical report, Instituto di Ecoomia G. Prato, Universita di Torino, Torino, Italy.

Ben-Porath, E., E. Dekel and A. Rustichini. 1993. On the relationship between mutation rates and growth rates in a changing environment. *Games and Economic Behavior* **5**, 576–603.

Blume, L. and D. Easley. 1990. Evolution and market behavior. *Journal of Economic Theory* **58**, 9–40.

Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **21**, 307–328.

Bollerslev, T., R. Y. Chou, N. Jayaraman and K. F. Kroner. 1990. ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics* **52**(1), 5–60.

Bossaerts, P. 1994. Time series analysis of inefficient markets. Technical report, California Institute of Technology, Pasadena, CA.

Bray, M. 1982. Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory* **26**, 318–339.

Breiman, L. 1961. Optimal gambling systems for favorable games. In *Proceedings of the Fourth Berkeley Symp. of Math Statistics, and Probability*, J. Newyman and E. Scott (eds), vol. 1. University of California Berkely Press, Berkely, CA.

Brock, W. A. and C. H. Hommes. 1994. Rational routes to randomness. Technical report, University of Wisconsin, Madison, WI.

Brock, W. A. and B. LeBaron. 1993. Using structural modelling in building statistical models of volatility and volume of stock market returns. Technical report, University of Wisconsin - Madison, Madison, Wisconsin.

Brock, W. A., J. Lakonishok and B. LeBaron. 1992. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance* **47**, 1731–1764.

DeLong, J. B., A. Schleifer, L. Summers and R. Waldman. 1992. 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. and G. Ramey. 1995. Calculation, adaptation and rational expectations. Technical report, University of California, San Diego, San Diego, CA.

Figlewski, S. 1978. Market "efficiency" in a market with heterogeneous information. *Journal of Political Economy* **86**(4), 581–597.

Friedman, M. 1953. The case for flexible exchange rates. In *Essays in positive economics*. University of Chicago Press, Chicago, IL.

Grannan, E. R. and G. H. Swindle. 1994. Contrarians and volatility clustering. Technical report, Unviversity of California, Irvine, Irvine, CA.

Grossman, S. and J. Stiglitz. 1980. On the impossibility of informationally efficient markets. *American Economic Review* **70**, 393–408.

Holland, J. H. 1992. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA.

Kareken, J. and N. Wallace. 1981. On the indeterminacy of equilibrium exchange rates. *Quarterly Journal of Economics* **96**, 207–222.

Kelley, J. L. 1956. A new interpretation of information rate. *Bell Systems Technical Journal* **35**, 917–926.

Knight, F. 1921. *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston, MA.

Lettau, M. 1993. Risk-taking bias in a financial market with adaptive agents. Technical report, Tilburg University, Tilburg, The Netherlands.

Marengo, L. and H. Tordjman. 1995. Speculation, heterogeneity, and learning: A model of exchange rate dynamics. Technical Report WP-95-17, International Institute for Applied Systems Analysis, Vienna, Austria.

Rieck, C. 1994. Evolutionary simulation of asset trading strategies. In *Many-Agent Simulation and Artificial Life*, E. Hillebrand and J. Stender (eds). IOS Press.

Routledge, B. R. 1994. Artificial selection: Genetic algorithms and learning in a rational expectations model. Technical report, University of British Columbia, Vancouver, British Columbia.

Sargent, T. 1993. *Bounded Rationality in Macroeconomics*. Oxford Uinversity Press, Oxford, UK.

Schumpeter, J. A. 1934. *The Theory of Economic Development*. Harvard University Press, Cambridge, MA.

Timmermann, A. 1995. Volatility clustering and mean reversion of stock returns in an asset pricing model with incomplete learning. Technical report, University of California, San Diego, San Diego, CA.

Youssefmir, M. and B. A. Huberman. 1995. Clustered volatility in multiagent dynamics. Technical report, Santa Fe Institute, Santa Fe, NM.