

# Building the Santa Fe Artificial Stock Market

Blake LeBaron \*  
Brandeis University

June 2002

## Abstract

This short summary presents an insider's look at the construction of the Santa Fe artificial stock market. The perspective considers the many design questions that went into building the market from the perspective of a decade of experience with agent-based financial markets. The market is assessed based on its overall strengths and weaknesses.

---

\* Graduate School of International Economics and Finance, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, [www.brandeis.edu/~blebaron](http://www.brandeis.edu/~blebaron). The author is a research associate at the National Bureau of Economic Research, and an external faculty member at the Santa Fe Institute.

# 1 Introduction

Most social systems involve complex interactions among many individuals. Much of economics hopes both to greatly simplify human behavior, and to do it in such a way that aggregate macro features can be easily characterized. The success of this traditional approach to describe human behavior has only had mixed success. One area where questions remain is finance where many empirical features are troubling for existing theories. Several of the foundations of the field are in a state of disarray, and new, radically different theories are appearing.<sup>1</sup>

One of the directions that researchers have been taking is the use of agent-based financial markets. These “bottom-up” models of financial markets start from first principals of agent behavior. Using either computational, or more sophisticated mathematical tools they are able to describe macro features emerging from a soup of individual interacting strategies. Since the early 1990’s I’ve been involved with one of the first agent-based financial market platforms, The Santa Fe Artificial Stock Market. Now, with nearly a decade of experience in looking at financial markets from an agent-based perspective, I would like to turn my attention back to the SFI market. I will explore some of the market’s early history, and the debates and design questions that went into its development. From a modern perspective there is still much to like about what the SFI market was trying to do. However, there are also things that I believe make it a less than perfect tool for modern research on financial markets.

Much of my market design philosophy stems from a desire to understand the impact of agent interactions and group learning dynamics in a financial setting. While agent-based markets have many goals, I see their first scientific use as a tool for understanding the dynamics in relatively traditional economic models. It is these models for which economists often invoke the heroic assumption of convergence to rational expectations equilibrium where agents’ beliefs and behavior have converged to a self-consistent world view. Obviously, this would be a nice place to get to, but the dynamics of this journey are rarely spelled out. Given that financial markets appear to thrive on diverse opinions and behavior, a first level test of rational expectations from a heterogeneous learning perspective was always needed.<sup>2</sup>

For this reason I have often used traditional economic models which have well defined homogeneous equilibria as the core for model building. This certainly isn’t necessary for an agent-based approach, but I think at this stage a deeper understanding of dynamics around well studied equilibria is essential. The existence of an equilibrium benchmark also provides an important test case for many of the computational learning tools in use. If for some cases they can actually learn convergence to an equilibrium then their abilities to perform purposeful search through the sets of rules are give

---

<sup>1</sup>See Heaton & Korajczyk (2002) along with articles in that volume for an introduction to behavioral finance issues. A very nice summary of the field is Shefrin (2000). Much of this work is still subject to controversy as shown by Rubenstein (2001).

<sup>2</sup>See for Sargent (1993) for further thoughts on the importance of learning rational expectations. Also, Board (1994) provides some important incites into the theoretical computability of equilibria.

greater credibility.

Beyond learning about the dynamics of certain types of markets subject to multi-agent learning, agent-based approaches have also replicated many empirical features. Although calibration exercises are still in an early stage of development, they are a promising and important route for validation.

This paper describes the Santa Fe Artificial Stock market from the perspective of one of the builders. Section two covers some of the early history of agent-based financial markets. Section three describes an early version of the SFI market. Section four describes the SFI market setup, and provides some perspective on many of the design questions. Section five briefly summarizes the results. Section six describes a new market model structure, and section seven summarizes and gives some ideas for the future of agent-based financial modeling.

## 2 Early Markets

Although one of the most sophisticated of early agent-based markets, the SFI market was not the first. There were several early simulations that tried to look the impact of random behaving agents on various market structures as in Cohen, Maier, Schwartz & Whitcomb (1983). Another early market looked the interactions of specific trading strategies (Kim & Markowitz 1989). Some of the more interesting early markets were concerned with the dynamics of foreign exchange as in Frankel & Froot (1988) and De Grauwe, Dewachter & Embrechts (1993). These follow much of the later literature in being concerned with the interaction of potentially destabilizing trend following strategies and their interactions with others. The SFI artificial market very much follows in the path of these early markets.

There were also several contemporaries to the SFI market which rightfully continue to draw much attention. Among these are the interactions models of Lux (1997), Kirman (1991), and Chiarella (1992) which provide relatively simple tractable dynamic frameworks for agent interactions. Rieck (1994) is also a early market with coevolving trading strategies as in the SFI market framework. Levy, Levy & Solomon (1994) is another early market example which introduced constant relative risk aversion preferences along with varying agent memory lengths.<sup>3</sup> Finally, the market of Beltratti & Margarita (1992) introduced trading strategies based on neural networks, and differing levels of agent sophistication. For further references to other agent-based financial markets see (LeBaron 2001a, LeBaron 2000a).

---

<sup>3</sup>This market is covered in extensive detail in Levy, Levy & Solomon (2000).

### 3 SFI Market Early History

The SFI market was born in discussions at the Santa Fe Institute in the late 1980's and early 1990's. It originated as a desire of Brian Arthur and John Holland to build a financial market with an ecology of trading strategies. Successful strategies would persist and replicate, and weak strategies would go away, creating potential niches for the entry of new strategies. The market was to be a continually coevolving soup of strategies. At the foundation of this was a desire to not preload much into the system, and to let evolution do most of the work. In its purist rendition the strategies and macro features would all be emergent. The project would also be able to make use of John Holland's important tools for modeling learning, the genetic algorithm, and classifier system. This early team was quickly augmented with Richard Palmer (a physicist) and Paul Tayler (a computer scientist).

There is actually a kind of first generation SFI market which was ready and operating in the late 1980's and early 1990's. This was eventually published in Palmer, Arthur, Holland, LeBaron & Tayler (1994). This was a simpler, but less economically oriented market of trading dynamics than the eventual SFI market. Rules simply mapped states of the world into buy or sell decisions. The market operated in the following fashion. A price was announced by a market maker to all the traders. Agents found a matching rule for the current market conditions, and came to the market with an order to buy or sell 1 share of stock. Most of the time this market was out of equilibrium with either more buyers or sellers. The smaller of these two sets would get satisfied while the other would get rationed. For example if there were 100 buyers ( $B_t$ ) and 50 sellers ( $S_t$ ), the sellers would all be able to sell 1 share, and the buyers would get only 1/2 share. The price would then be adjusted to reflect either the excess demand or supply.

$$p_{t+1} = p_t + \lambda(B_t - S_t) \quad (1)$$

I have commented previously on some of the problems with this pricing mechanism<sup>4</sup>, but it is important to repeat them here. The first problem is that the results were very sensitive to  $\lambda$ . Setting  $\lambda$  too low generated long persistent trends during which the market stayed in excess demand or supply. Setting  $\lambda$  too high yielded markets in which the price overreacted, and the market thrashed back and forth from excess supply to excess demand.<sup>5</sup> The second problem was that rules were reinforcing based on buying and selling behavior that wasn't getting filled. Since every period some set of traders left the market with their orders unfilled this was an important consideration. However, adjusting fitness functions for expected trade amounts would be a difficult task. The early market was still capable of being run,

---

<sup>4</sup>(LeBaron 2001a).

<sup>5</sup> $\lambda$  is essentially a measure of market depth, or the ease with which traders find others to trade with at a given price. It makes sense that this should be an endogenous feature of the market. Recently, LeBaron (2001d) shows that sudden reductions in depth are related to market instabilities in an agent-based market.

and led to many interesting discussions. In some versions of the software one can still see some of the legacy of this early market in terms of available market making structures.

## 4 The SFI Market Structure

In 1993 I was at SFI directing the economics program, and working on a small scale agent-based financial market of my own. My interest was directed at many of the empirical issues I was dealing with. Most of which were related to understanding the joint dynamics of prices and trading volume. In absence of some kind of theoretical framework these time series exercises were getting increasingly difficult. At this time I was added to the team, and along with Arthur and Palmer we began some extensive modifications to the SFI market. This section briefly summarizes these changes which eventually became the SFI market platform. Also, several limitations of the design questions are discussed as well.

### 4.1 Assets

The SFI market structure in terms of tradeable assets was designed to be as simple as possible. There are two assets. 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 the following autoregressive process,

$$d_t = \bar{d} + \rho(d_{t-1} - \bar{d}) + \mu_t, \quad (2)$$

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 a nonstationary dividend processes. The price of a share of stock,  $p_t$ , is determined endogenously in the market.

The system is simple and fits well with the other mechanisms of the model. Its simplicity does open several important questions. First, it is not calibrated to any sort of actual data. No where is  $t$  specified as to what the time period is. This is fine for early qualitative comparisons with stylized facts, but it is a problem for quantitative calibration to actual time series. Another problem is the fixed interest rate. It would seem sensible to have interest rates set endogenously. Also, the market is definitely a partial equilibrium one since there is no market clearing restriction in place for the debt market. Clearing a second market opens rather extensive design questions, but it does remain an important open issue. Another common critique is the fact that there is only one equity issue. Adding further shares would be useful, especially for studying trading volume, but again this opens up another range of parameters,

and significantly increases model complexity. Until more is understood about the dynamics of these markets, a single risky asset world is a good benchmark.

A final tricky timing issue is how frequently dividends are paid. As the only fundamental, dividends play a key role in the information process of the market. In actual markets traders may have many periods in which to observe the price when no actual fundamental information is changing. This may greatly facilitate the chances for successful learning to occur, and it could also hinder them, in that unusual beliefs are not stopped immediately by new fundamental information. In the SFI market the fundamental changes every period which it least makes it a little different from reality. Solving this timing problem in a realistic way is a difficult, but interesting problem.

## 4.2 Preferences and Trading

In response to potential criticisms about the ad hoc price setting mechanism used in Palmer et al. (1994) the modern versions of the SFI market have used a simple constant relative risk aversion preference format for equity demands.<sup>6</sup> Individuals in the market form their demands for stock as

$$x_t^i = \frac{\hat{E}_t^i(p_{t+1} + d_{t+1}) - (1 + r_f)p_t}{\gamma \hat{\sigma}_{p+d,i}^2}, \quad (3)$$

where the subscript  $i$  represents the fact that beliefs may differ across agents.  $r_f$  is the risk free rate, and  $\gamma$  is the coefficient of absolute risk aversion. Agents use a classifier system to make predictions on the first and second moments for stock returns. This is described in the next section. A rule tells each agent how to forecast future returns, and what the conditional variance of this forecast is. The forecast is assumed to be linear in the current price and dividend.

$$\hat{E}_t^i(p_{t+1} + d_{t+1}) = a_j(p_t + d_t) + b_j \quad (4)$$

The  $j$  subscript refers to the rule chosen by agent  $i$ . This restricted forecasting rule along with the demand for shares gives a demand function which is linear in  $p_t$ . Setting the total number of shares to a given fixed value then allows for the solution of equation 3 for a temporary equilibrium price. After the price is set, agents update their portfolios, and trading volume is recorded.

This method for handling market clearing is clean and simple, and doesn't involve troubling issues with price adjustment which are part of many other mechanisms. The temporary equilibrium aspect also makes it appealing to those who perceive financial markets operating well over short time horizons. There is also a well defined heterogeneous

---

<sup>6</sup>This is very standard in finance. See for example Grossman & Stiglitz (1980).

rational expectations equilibrium for this model which is most clearly defined in LeBaron, Arthur & Palmer (1999) as parameters for an equilibrium pricing function mapping the dividend into a current price. This benchmark forms a key test for the market. When and how does it converge to this equilibrium? If it does for some parameters we can feel a little more comfortable with the market design and setup. Furthermore, if we can understand why it isn't converging in some instances, then much can be discerned about the out of equilibrium structure, and learning dynamics of the market.

There are several weaknesses of this approach. First, it is obvious that the price forecast itself is linear in  $p_t$ . It is not clear what impact this restriction has on the market dynamics. Another problem is the share demand function itself. Even though wealth is shifting among traders there is no place for wealth to enter into the share demand equation. In actual markets it is clear that wealthier traders will have a greater impact on prices. In this market they do not. A change to constant relative risk aversion would alleviate this problem.<sup>7</sup> Another important modification is to move to full intertemporal preferences with agents evaluating consumption each period. In general this would involve some potentially complicated issues in terms of the interactions between portfolio and consumption choices.<sup>8</sup>

A more fundamental critique would be that real financial markets are never truly in equilibrium. They are institutions which are designed to handle buying and selling in real time, and therefore may be far from the ideal world of Walrasian market clearing. This market mechanism has simply avoided the entire microstructure aspect of the market. It is assuming that supply and demand are somehow balancing in some unspecified market institution. This seems like a good first assumption to start with, and might be fine for thinking about longer horizon price patterns. However, realistic modeling of market institutions and trading is an important extension that needs to be considered.<sup>9</sup>

### 4.3 Classifiers and Forecasting

The linear agent forecast is only a local forecast used by agents. Agents are equipped with multiple linear models which can be used. Appropriate forecasting models are selected according to specific market conditions. The setup used is based on the classifier systems developed in Holland, Holyoak, Nisbett & Thagard (1986). Classifier rules were originally a mapping from states of the world, represented as bit-strings, into actions. This was the case with the early SFI market too, but the later version changed this to a mapping from the conditions into forecast parameters.

This is sometimes been referred to as a condition/forecast classifier.<sup>10</sup>

<sup>7</sup>Levy et al. (1994) is an example of CRRA preferences in an agent-based market.

<sup>8</sup>LeBaron (2001c) and LeBaron (2001b) present examples of this in a simplified setting using log utility.

<sup>9</sup>See Daniels, Farmer, Iori & Smith (2002) for a recent analysis of the dynamics of trading.

<sup>10</sup>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 (1999), 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.

A classifier rule is given by a bit-string and a parameter vector. An example would be

$$(\#, 0, 1, \#; a_j, b_j, \sigma_j^2). \tag{5}$$

The first part of the rule matches current conditions in the market. A 1 would match a true condition, and a 0 a false condition. The # sign is the don't care symbol, and it matches either a true or false condition. This allows rules to be of varying specificity. Some might condition on many events and some might condition on very few. Table 1 gives some examples of the actual states that a rule matches.

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 1: Matching Examples

In the construction of classifier systems the market designer is given a large amount of design power by setting the actual conditions which will be used. The SFI market is no exception. The conditions used in the SFI market runs are given in table 2. MA refers to a moving average of past prices. Agents can build rules conditioned on this information, but it would be difficult to move beyond this set. Agents certainly have the power to ignore any piece of information, and it is often the case that basic questions are about why certain pieces of information get used in actual market runs.

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

Several important things should become apparent when one looks at table 2. First, it contains two useless bits in the form of an always on bit, and an always off. These are intended to be zero information bits which can then be

compared with the others to assess if the number of rules using the real information are larger than those using the useless information. Another issue which appears is that rules can easily be created that would never be used. For example, requiring that bit 1 be false, and bit 2 to be true would require the price dividend ratio to be less than  $1/4$  and greater than  $1/2$ . This will end up leaving many rules which will never get used. The classifier system monitors rules, and when they haven't been used for a long time they are modified by changing a 0 or 1 bit-string value into a #. This is known as "generalization" in the classifier literature. This helps to accommodate the problem of unused rules, but it is not a very elegant solution.

Agent rule sets are an evolving set of 100 of these rules. They are required to hold a default rule which is given as all #. There is no rule sharing at all between agents. In this sense there is no form of social learning in this market. Other early agent markets such as Arifovic (1996) involve extensive social learning. This complete lack of communication is an extreme, and is a feature of the SFI market.<sup>11</sup>

Rules are evaluated based on forecast accuracy which is updated using,

$$\sigma_{t,i,j}^2 = \left(1 - \frac{1}{\tau}\right)\sigma_{t-1,i,j}^2 + \frac{1}{\tau}((p_t + d_t) - (a_{i,j}(p_{t-1} + d_{t-1}) + b_{i,j}))^2, \quad (6)$$

for only the matched rules each period. The value of  $\tau$  is fixed and set to 75, but it is an interesting open parameter. This exponential weighted forecast error is reminiscent of many learning algorithms, and it tries to capture possible nonstationarities in the series while maintaining an accurate estimate of forecast performance. This value remained fixed through all the SFI market experiments in the early papers. Experiments with different values, and changing this value across agents would be very interesting. Agents proceed to match rules each period. They may match more than one rule, in which case they use the rule which has had greater forecast accuracy over the recent past.

As a market designer I can say both good and bad things about the use of classifiers in the SFI market. In a general learning context they are very appealing for their simplicity, and their ability to tackle high dimensional state spaces. This was their original intention, and I think they continue to provide a useful metaphor for learning. They also give the researcher a fairly easy way to see what types of information agents are actually using. However, implementing them involves some complicated issues that require several non-obvious assumptions. Also, they are not particularly well designed for the real valued learning problems which are common in financial and economic applications. In the SFI market this required hard wiring key breakpoints for the information variables as in table 2.<sup>12</sup> It would be good to see research on classifier systems continue both in economic and noneconomic contexts, but future market designers should realize that they come with some extensive costs.

---

<sup>11</sup>Vriend (2000) provides a nice comparison of individual and social learning in agent-based settings.

<sup>12</sup>In a recent variant on this theme, Tay & Linn (2001) replace the hard breakpoints with fuzzy-logic constructions.

There is one last issue involved with implementing behavioral rules. In the SFI market rules generate forecasts, and forecasts are used to map into eventual behavior. In some agent-based markets the forecasting process is completely eliminated.<sup>13</sup> Which procedure is best? This has never been completely studied. Avoiding the forecasting step provides a simpler more direct method for studying agent learning. It also avoids some potentially difficult issues related to statistical decision theory. Forecasting agents still have some appeal in that this is often what people are trying to do, and the experimenter can often get a better outside objective feel for rule performance when a forecasting objective is involved.

#### 4.4 The Genetic Algorithm

The SFI market could be run as an experiment in the dynamics of a set of rule based trading agents. However, it was always designed to study the emergence of trading patterns as agents learned over time. Therefore, some learning mechanism is needed to modify the trading rules. Like many other models of learning the SFI market uses a genetic algorithm (GA).<sup>14</sup> The GA is implemented with a certain probability by agents in an asynchronous fashion. The probability of GA activation by each agent is an important parameter which determines the “speed of learning” of the agents.

When agents update their rules with the GA they estimate the strength of each rule based on forecast accuracy and a small penalty for specificity which measures the number of non # elements in each rule. This yields an evolutionary pressure towards parsimony. Also, in the benchmark equilibrium none of the information variables are useful and this fitness pushes the market to rules which use none of the bit-string information. The GA initially removes the 20 worst rules out of the set of 100. New rules are generated by choosing parents, and implementing crossover and mutation operators on them. The details of these operations are covered in LeBaron et al. (1999). Several design issues are important and will be highlighted here. First, the GA is modified to handle the real valued elements of the classifier rules. This differentiates the SFI market from several other early markets which often used binary strings for the GA, and then mapped them into real values for the final problem. It has always seemed more intuitive to modify the GA operators to handle real valued parameters in a sensible fashion. Another interesting feature of the SFI GA is that it uses a very simple selection method, tournament selection. This is a very unselective version of selection. For example, when a parent is needed, two candidates are chosen from the surviving 80 rules, and the more fit of these is used as parent 1. This process is repeated for parent 2. This is much less selective than other forms of selection

---

<sup>13</sup>See Arifovic (1996) or Lettau (1997) as early examples. LeBaron (2001c) also follows this approach.

<sup>14</sup>The 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).

which might weight probabilities of rule choices based on fitness (fitness proportional selection). Tournament selection serves two useful purposes. It keeps the population of rules from converging too fast. Also, it means that fitness values are only ordinal in that any monotonic transformation of these would give similar selection properties. This is not true for many other selection methods. Given that the fitness values may not be a perfect measure of the usefulness of each rule, the ordinal aspect of tournament selection is very appealing in the SFI market.

One final difficult feature of the GA is how to evaluate new rules. In some market settings the rules can be immediately evaluated and then compared with the existing rule sets. One can restrict the addition of new rules to only those that are improvements over existing rules. This is the election operator developed in Arifovic (1996). Unfortunately, in the SFI market, the forecast performance of a new rule is difficult to assess. The market conditions may be changing, and taking each rule back through some fixed window would be difficult. In the SFI market rules inherit the average forecast accuracy of their two parents. It is not clear that this would be a sensible estimate for the accuracy of a new rule. Setting the fitness values for new rules is another interesting design issue.

## 4.5 Software Systems

I have concentrated so far on abstract model design issues, but have not addressed basic computational problems such as modeling languages and platforms. These issues are often ignored in many papers on agent-based systems since we'd like to hope that they are second order problems which are not important to the basic results. Since this volume is concerned with mapping several agent based systems into a common software platform I felt some commentary would be useful.

The SFI market was originally programmed in the c programming language on UNIX based workstations. It was later modified to objective-c and developed extensively on NeXT computers.<sup>15</sup> The objective extensions in objective-c were very useful for the software design process, and do help to compartmentalize data and algorithms in the agent-based world. The software used on the NeXT also had a very sophisticated user interface and real time graphics showing the evolution of the market as it was running. These features are also part of the Swarm toolbox as well. Eventually, the market was ported into Swarm, and the most reliable and up to date version is a Swarm version.<sup>16</sup> The choice of computing platform is an important one. One would like to have ease of use for designers who are probably in continual state of design changes. There also is a need for extensive portability. Anyone building these types of models has to have in mind a useful distribution that others can eventually use and build on. Finally, one needs access to common tools such as statistical, matrix, and graphic routines. It is not the purpose of this paper to be a proponent

---

<sup>15</sup>Three out of the five (Arthur, LeBaron, Palmer) coauthors were NeXT users.

<sup>16</sup>The market is maintained by Paul Johnson at <http://ArtStkMkt.sourceforge.net>. He has made many changes to make the code more readable and reliable. Johnson (2002) is a very nice description of many of the software coding issues he encountered.

for any one package, but it is safe to say that there is no single software platform available that accomplishes all these goals. Hopefully, there will be further convergence so that agent based modelers will be more likely to share models and software tools.

## 5 Results

Most of the key early empirical results from the SFI market are presented in Arthur, Holland, LeBaron, Palmer & Tayler (1997) and LeBaron et al. (1999). Similar to many other agent-based markets it generates several features similar to actual financial data. Return series show excess kurtosis, very little linear autocorrelation, and persistent volatility. Trading volume is also strongly persistent and correlated with price volatility. Most of the series generated also display small amounts of predictability in ways that are similar to actual returns series. Certain technical trading rules yield a small amount of predictability, as do price/dividend ratios. One important feature of the empirical results presented in LeBaron et al. (1999) is that they use a cross section of 25 different market runs. If the market process were known to be ergodic, then running from different initial conditions would be unnecessary. However, there is no way to prove ergodicity for these computational models, so reporting the summaries over different starting values is important.

The second key result is that these features are very sensitive to the learning speeds of agents, or the frequency with which they run the genetic algorithm. When they update their rules frequently the market is more likely to generate the patterns common to actual financial time series. When they update their rules more slowly the market is very close to what would be predicted in the homogeneous rational expectations equilibrium. This shows both that the learning agents are capable of finding the equilibrium, but it is not clear in actual settings how people might coordinate to slow down their learning.

The SFI market results are only a first step in a calibration exercise. First, the model and the dividend process, are not aligned with actual macro finance series as in studies such as Mehra & Prescott (1985). This makes success or failure in the aggregate time series dimension difficult to assess. Second, and maybe more importantly, the quantitative values in the SFI market don't really line up that well to actual data. One example is return variances themselves. LeBaron et al. (1999) report in their table 4 that the standard deviation of returns changes from 2.135 to 2.147 when going from the slow learning to the fast learning cases. Although this is an increase, the SFI market is not a mechanism for greatly magnifying fundamental volatility. Also, the excess kurtosis levels in the fast learning case are only about 0.32 which is much lower than actual daily or weekly returns series. In many ways it is doing the right thing qualitatively, but under shooting quantitatively.

The paper also reports some results on the behavior of the agents. These often get ignored in summaries of the SFI market. Observed behavior in the fast learning case is significantly different from that in the slow learning case. Agents reveal much more heterogeneity in their forecast parameters which should be expected. They also are more likely to be using rules which are conditioned on trend following technical trading rules. This is the important aspect of strategy reinforcement that the market was developed to analyze. Through no coordination other than through movements in the price itself agents appear to be using similar trading strategies. The contribution of this behavior to the dynamics of prices in the SFI market remains an interesting open question.

## 6 A New Agent-based Market

My experiences with the SFI market were extremely valuable, and I think it has provided many useful lessons. However, I shut down my research on this platform in 1997, and began the design of a new market. As I've mentioned earlier there were several design issues that really needed a radical change from what was going on in the SFI market. These included:

1. Moving to intertemporal constant relative risk aversion preferences where wealth distributions determine the amount of price impact different traders and strategies have.
2. Calibrating fundamental movements to actual financial time series.
3. Coordinating social learning between agents.
4. Replacing the classifier systems with a more streamlined mechanism for mapping past information into trading strategies using a neural network inspired nonlinear functions.

In several recent papers, (LeBaron 2001*c*, LeBaron 2001*b*, LeBaron 2002*b*) I have used a new framework which I think of as a second generation artificial market. It is a platform designed to make serious attempts at empirical calibration, and it uses an economic structure which is more directly comparable to the literature in macroeconomics and finance. Beyond this it tries to simplify much of the structure in the SFI market by reducing parameters and zeroing in on some of the key dimensions of the problem.

While the SFI market emphasized learning speeds, this new market emphasizes memory length as a key heterogeneity dimension, and determining factor for convergence to homogeneous equilibrium pricing. Short memory traders using small amounts of past time series appear to contribute to a significant magnification of fundamental volatility and are also difficult to evolutionarily remove from the market. Speed of learning changes do have similar

features to the SFI market and they can be checked directly as in LeBaron (2002*b*). However, the short memory traders alone seem to be an interesting feature which replicates a kind of behavior which may be common in many economic and financial situations.<sup>17</sup> Calibration results in (LeBaron 2001*b*, LeBaron 2002*a*) are not perfect, but they are more interesting than those from the SFI market. It seems clear that future versions of this market will be able to faithfully replicate most of the desired features, and to do this with reasonable quantitative accuracy.

Some basic lessons and design philosophies are maintained from the SFI platform. Trading strategies are not hard wired, but are allowed to emerge as agents learn and adapt over time. Many key parameters are not set, but are set in an evolutionary dynamic. Finally, trading is again a simple market clearing mechanism which, in this case, needs to be done numerically. A modified genetic algorithm is used for agent learning which borrows much from what was learned in designing a real valued GA for the SFI market. This market represents a natural second generation evolutionary step for the SFI market type of structure. It is important to realize that while there is a great need to explore given market structures, there also needs to be progress and evolution on how to build simple, sensible agent-based models for financial markets. This evolution will clearly continue.

## 7 The Future of Agent-based Finance

What are the basic lessons from early markets like the SFI market and others, and where are they leading us? I think the key lesson is to remind us of the the fragility of the the equilibrium models that we consider in finance and macroeconomics. The hope that the interactions of rational traders might lead to a relatively simple aggregate outcome is only possible under highly restrictive assumptions, and the computational world shows us that an even wider array of simple learning models share this feature. Dynamics of both states, and expectational perceptions need to play a far bigger role in modeling than it has in the past. Furthermore, we need to think about heterogeneity in more sophisticated endogenous forms. The ability to analyse a world of endogenous heterogeneity is one of the cornerstones of the agent-based modeling approach. The simple world of a “rational” and “noise” trader must be replaced with a rich array of beliefs whose rationality and fitness is very dependent on the world they live in, and simultaneously the worlds they create.<sup>18</sup> The SFI market provided an initial window into this world.

At a practical level current research will probably continue in several directions. First, efforts to quantify and calibrate markets to actual data will be a key goal. This area has moved from a more qualitative style as in the SFI market to more direct quantitative experiments. As in (Mehra & Prescott 1985) markets will be constructed to roughly

---

<sup>17</sup>See Rabin (forthcoming 2002) for further examples and theoretical justifications for overweighting the recent past.

<sup>18</sup>Many agent-based markets remind us that cross sectional heterogeneity is an important part of the evolutionary process, and its dynamics should not be ignored.

align with actual data in a way that makes reasonable sense as a representation of the economy as a whole. Early successes in terms of volatility, leptokurtosis, and trading volume give a strong indication that this will be a fruitful search. Second, it seems clear that the theoretical work paralleling the computational side will continue. Here, generic features from multi-agent systems may emerge answering the criticism that much of what is going on is specifically crafted into any single model. Finally, the artificial markets will continue to provide useful thought experiments on how real markets function. This amounts to “turning the dials” in an intelligent fashion on computer simulation platforms. The basic questions about convergence to a rational expectations equilibrium and the learning speed in the SFI market is an early example of this.<sup>19</sup> These experiments help is to understand what it is that is special about financial markets that may lead to less than optimal dynamics. The computer serves as a tool to point out which behavioral issues may be critical, and may also give hints as to where to look in the data to see traces of these elements.

The SFI market and several of its contemporary platforms also brought an entire new set of computer tools into the economics realm. The usefulness and applicability of these tools still remains a open question. Even the style in which they are implemented remains subject to some debate in the agent-based modeling community. This summary has touched on a few of the components of the SFI market, concentrating on the classifier system, and the genetic algorithm. In both cases it is important to realize that the standard computer science tools need to be modified to better suit the problem at hand. The classifier remains controversial as a modeling tool in economics. While it seems an intuitive way to approach simple rule-based behavior it contains many practical complications that make it difficult to implement. It will be good to see research on the classifier continue in the future. The genetic algorithm is a kind of foundation for many agent-based approaches, and the SFI market was one of the early applications. It generally continues as a kind of robust method for modeling learning without too much additional cognitive baggage.<sup>20</sup> In general most results have been robust to changes in basic algorithm specifications, but this result has not be broadly documented within or across model platforms. In summary the computer science tools side of the SFI market opened up many questions which still remain unanswered.

I’m often asked to turn the table, and try to explain the major potential weaknesses of this line of research. I think it does have some potential problems on the horizon. The first relates to the sensitivity of results in the SFI market to learning speeds. This can be taken as an interesting result, but it could also be a weakness too. If a single parameter for which we know little about in reality can change the outcome so dramatically then we may always be in a state of uncertainty concerning potential model predictions. It seems possible that many aggregate dynamics may rest on how fast agents are responding to each other. Tune this too fast, and the evolutionary process concentrates on

---

<sup>19</sup>See (LeBaron 2002b) for other examples of convergence testing and issues.

<sup>20</sup>For support on this point from a cognitive science perspective see Clark (1997).

a continuing process of adapting to what the other person did last period which leads to a dynamic with only weak tendencies to converge to an equilibrium. If the parameter is tuned to a slower learning rate, then the agents can adapt to the underlying economics of the situation, and they can actually learn how to behave in a rational expectations equilibrium. This sensitivity may be a kind of Achilles heal to the entire evolutionary finance agenda. Researchers should probably keep this in mind, but should still continue with what they are doing. A second problem might be that some forms of evolution in actual markets are happening at time scales that are large relative to the length of data we have.<sup>21</sup> The conjecture is that people have been steadily learning about risk over the last century of financial data. As this learning has proceeded they have emphasized equity investments more heavily and have driven down the equity premium. This debate is interesting in its own right, but is particularly troubling to those trying to calibrate agent-based markets since it would be difficult to estimate or calibrate to a type of evolution and learning happening across the entire available data set.

Despite its faults the Santa Fe artificial market still has extensive appeal to researchers from many different disciplines. I think this is due to its dual personas with both radical agent dynamics, and conservative economic structure. Its radical side comes from the fact that it leaves the agent behavior open ended. They try to figure out what they can from the time series that go by with very little help preloaded into them. Their continuing struggle to find market inefficiencies and capitalize on them is very appealing to those of us believing that certain parts of efficient market arguments can't be ignored, and obvious patterns will not last long in the data.<sup>22</sup> The conservative side of the market comes from its very standard setup which yields simple benchmarks, and results which can be readily interpreted. As this paper documents there are many more experiments that can be done within the basic framework. I hope that further distribution of this software will help researchers to take on some of these difficult problems.

---

<sup>21</sup>An example of this is the current debate about the level of the equity premium, (Arnott & Bernstein 2002) and (Fama & French 2002). Also, certain types of trading rules have been shown to generate different results over recent periods, (Sullivan, Timmerman & White 1999) and (LeBaron 2000b).

<sup>22</sup>The recent market of Chen & Yeh (2001) takes open ended learning to a higher level by modeling agent behavior with genetic programming, which evolves forecasting equations themselves.

## References

- Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.
- Arnott, R. D. & Bernstein, P. L. (2002), 'What risk premium is "normal"?'', *Financial Analysts Journal* **58**(2), 64–85.
- 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.
- Beltratti, A. & Margarita, S. (1992), Evolution of trading strategies among heterogeneous artificial economic agents, in J. A. Meyer, H. L. Roitblat & S. W. Wilson, eds, 'From Animals to Animats 2', MIT Press, Cambridge, MA.
- Board, R. (1994), 'Polynomial bounded rationality', *Journal of Economic Theory* **63**, 246–270.
- Chen, S. H. & Yeh, C. H. (2001), 'Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market', *Journal of Economic Dynamics and Control* **25**, 363–394.
- Chiarella, C. (1992), 'The dynamics of speculative behaviour', *Annals of Operations Research* **37**, 101–123.
- Clark, A. (1997), Economic reason: The interplay of individual learning and external structure, in 'The frontiers of the new institutional economics', Academic Press, San Diego, CA, pp. 269–290.
- Cohen, K. J., Maier, S. F., Schwartz, R. A. & Whitcomb, D. K. (1983), 'A simulation model of stock exchange trading', *Simulation* **41**, 181–191.
- Daniels, M. G., Farmer, J. D., Iori, G. & Smith, E. (2002), How storing supply and demand affects price diffusion, Technical report, Santa Fe Institute, Santa Fe, NM.
- De Grauwe, P., Dewachter, H. & Embrechts, M. (1993), *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*, Blackwell, Oxford.
- Fama, E. F. & French, K. R. (2002), 'The equity premium', *The Journal of Finance* **57**(2), 637–631.
- 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.
- Goldberg, D. E. (1989), *Genetic Algorithms in search, optimization and machine learning*, Addison Wesley, Reading, MA.
- Grossman, S. & Stiglitz, J. (1980), 'On the impossibility of informationally efficient markets', *American Economic Review* **70**, 393–408.
- Heaton, J. & Korajczyk, R. (2002), 'Introduction to *review of financial studies* Conference on market frictions and behavioral finance', *Review of Financial Studies* **15**(2), 353–362.
- 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.
- Johnson, P. (2002), 'What I learned from the artificial stock market', *Social Science Computer Review* **20**(2), 174–196.
- Kim, G. & Markowitz, H. (1989), 'Investment rules, margin, and market volatility', *Journal of Portfolio Management* **16**(1), 45–52.
- Kirman, A. P. (1991), Epidemics of opinion and speculative bubbles in financial markets, in M. Taylor, ed., 'Money and Financial Markets', Macmillan, London.
- LeBaron, B. (2000a), 'Agent based computational finance: Suggested readings and early research', *Journal of Economic Dynamics and Control* **24**(5-7), 679–702.
- LeBaron, B. (2000b), 'The stability of moving average technical trading rules on the Dow Jones index', *Derivatives Use, Trading, and Regulation* **5**(4), 324–338.
- LeBaron, B. (2001a), 'A builder's guide to agent-based financial markets', *Quantitative Finance* **1**(2), 254–261.
- LeBaron, B. (2001b), 'Empirical regularities from interacting long and short memory investors in an agent based stock market', *IEEE Transactions on Evolutionary Computation* **5**, 442–455.
- LeBaron, B. (2001c), 'Evolution and time horizons in an agent based stock market', *Macroeconomic Dynamics* **5**(2), 225–254.

- LeBaron, B. (2001*d*), Financial market efficiency in a coevolutionary environment, *in* 'Proceedings of the workshop on simulation of social agents: Architectures and institutions', Argonne National Laboratories, pp. 33–51.
- LeBaron, B. (2002*a*), Calibrating an agent-based financial market to macroeconomic time series, Technical report, Brandeis University, Waltham, MA.
- LeBaron, B. (2002*b*), 'Short-memory traders and their impact on group learning in financial markets', *Proceedings of the National Academy of Science: Colloquium* **99**(Supplement 3), 7201–7206.
- LeBaron, B., Arthur, W. B. & Palmer, R. (1999), 'Time series properties of an artificial stock market', *Journal of Economic Dynamics and Control* **23**, 1487–1516.
- 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. (1999), 'Rules of thumb versus dynamic programming', *American Economic Review* **89**, 148–174.
- Levy, M., Levy, H. & Solomon, S. (1994), 'A microscopic model of the stock market: cycles, booms, and crashes', *Economics Letters* **45**, 103–111.
- Levy, M., Levy, H. & Solomon, S. (2000), *Microscopic Simulation of Financial Markets*, Academic Press, New York, NY.
- Lux, T. (1997), 'Time variation of second moments from a noise trader/infection model', *Journal of Economic Dynamics and Control* **22**, 1–38.
- 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.
- Mehra, R. & Prescott, E. (1985), 'The equity premium: A puzzle', *Journal of Monetary Economics* **15**, 145–161.
- Mitchell, M. (1996), *An Introduction to Genetic Algorithms*, MIT Press, Cambridge, MA.
- Palmer, R., Arthur, W. B., Holland, J. H., LeBaron, B. & Tayler, P. (1994), 'Artificial economic life: A simple model of a stock market', *Physica D* **75**, 264–274.
- Rabin, M. (forthcoming 2002), 'Inference by believers in the law of small numbers', *Quarterly Journal of Economics*

- Rieck, C. (1994), Evolutionary simulation of asset trading strategies, in E. Hillebrand & J. Stender, eds, 'Many-Agent Simulation and Artificial Life', IOS Press.
- Rubenstein, M. (2001), 'Rational markets: Yes or no? The affirmative case', *Financial Analyst Journal* **17**, 15–29.
- Sargent, T. (1993), *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.
- Shefrin, H. (2000), *Beyond Fear and Greed*, Harvard Business School Press, Cambridge, MA.
- Sullivan, R., Timmerman, A. & White, H. (1999), 'Data-snooping, technical trading rule performance and the bootstrap', *Journal of Finance* **54**, 1647–1691.
- Tay, N. S. P. & Linn, S. C. (2001), 'Fuzzy inductive reasoning, expectation formation and the behavior of security prices', *Journal of Economic Dynamics and Control* **25**, 321–362.
- Vriend, N. (2000), 'An illustration of the essential difference between individual and social learning, and its consequences for computational analysis', *Journal of Economic Dynamics and Control* **24**, 1–19.