Financial Markets can be at Sub-Optimal Equilibria

Shareen Joshi[†‡] (shareen@santafe.edu), Jeffrey Parker[†]
(jeffrey.parker@reed.edu) and Mark A. Bedau[†‡] (mab@reed.edu)
[†] *Reed College, 3203 SE Woodstock Blvd., Portland OR 97202*
[‡] *Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe NM 87501*

Abstract. We use game theory and the Santa Fe Artificial Stock Market, an agent-based model of an evolving stock market, to study the properties of strategic Nash equilibria in financial markets. We discover two things: there is a unique strategic equilibrium in the market, and this equilibrium is sub-optimal since traders' earnings are not maximized and the market is inefficient. This strategic equilibrium is due to an analogue of the prisoner's dilemma; the optimal global state is unstable because each individual has too much incentive to "defect" and use forecasting rules that pull the market into the sub-optimal equilibrium.

Keywords: finance, efficiency, Nash equilibrium, game theory, agent-based models, prisoner's dilemma

JEL code: G10

1. Introduction

Traditional analyses of financial markets have argued that trading should stabilize asset prices and bring about market efficiency (originally shown by Samuelson (1965); see also Fama (1970) and Malkiel (1992)). But evidence of market inefficiencies, such as bubbles and crashes (see Chancellor (1999) for a historical account and Campbell and Lo (1997) for a review of the empirical evidence), excess volatility (Shiller, 1989), and other non-linearities in time series data (Lo and Mackinlay 1999, Campbell and Lo 1997), has generated interest in sources of market inefficiency. A variety of theoretical, empirical, experimental and computational models are now being used to argue that asset prices may be destabilized by such factors as technical trading or 'trend following' (Brock et al. 1992, Keim and Madhaven 1995, Soros 1994,Werner, deBondt and Thaler 1995), noise-trading (Delong et al. 1990b, Delong et al. 1991), the existence of fashions, fads or sunspots (Shiller 1989, Farmer 1993), and the heterogeneity of beliefs in the market (Marengo and Tordjman1995, Farmer 1998). These factors are thought to cause sustained deviations of asset prices from their fundamental values, increased risk, and altered earnings patterns in the market, thus making markets inefficient. It is an open question, however, whether inefficient markets can be at an equilibrium, or whether they are likely to evolve towards increased efficiency. It is also unclear whether in the long run the average investor profits more in an inefficient market than in an efficient market. We address these questions in this paper. We argue that, although the average trader is better off in an efficient market, efficient markets can be unstable because each individual trader can have a strong incentive to use trading methods that create market inefficiencies. Thus markets can be inefficient and unlikely to evolve towards higher efficiency.

Dropping the assumption of rational expectations and informational efficiency in markets, we study a model of an evolving stock market consisting of "boundedly-rational" traders who learn through their market experience, continually adapting their behavior to changing market conditions (Arthur 1992, Sargent 1993). We develop a simple game-theoretic framework to study the decision-making process of such traders. Central to our analysis is the assumption that traders lack perfect foresight about what other traders will think and do. At the same time, since a trader's profits depend on the behavior of other traders, each trader should make investment decisions on the basis of her best guess about what other traders will be thinking and doing (Arthur 1992, Sargent 1993). In Keynes' words, she must anticipate "what average opinion expects average opinion to be" (Keynes, 1936).

A market operating under these conditions is a complex adaptive system consisting of a co-evolving ecology of heterogeneous traders. A central factor governing the behavior of such markets is the *rate* at which traders revise and adapt their market-forecasting methods. This revision rate, analogous to temperature in physical systems, determines the market's behavior since different forecast-revision rates promote the use of different kinds of forecasting methods in the population of traders. Here we study, using a game-theoretic framework, what forecast-revision rates maximize traders' profits. We also compare the market's behavior under various forecast-revision rates to determine whether the optimal forecasting rate leads to the best state globally.

To investigate the co-evolution of the market-forecasts of boundedly-rational and heterogeneous traders, we study an agent-based computer simulation of such a market, specifically, the Santa Fe Artificial Stock Market (Palmer et al. 1994, Arthur et al. 1997, LeBaron 1997). This model's key characteristic is that risk-averse traders choose their market-forecasting rules from an evolving set of rules, depending on which ones have proved to be the most successful predictors of recent stock-price changes. Previous work in this model (Palmer et al. 1994, Arthur et al. 1997, LeBaron, Arthur, and Palmer 1998) has shown that the market's behavior is consistent with the theory of efficient markets if traders revise their forecasts infrequently, but its behavior shows inefficiencies similar to real markets if traders revise their forecasts frequently.

We build on this work in two ways. First, we investigate the optimal rate at which traders should revise their repertoire of market-forecasting rules. We find that a rapid forecast revision rate is the unique symmetrical strategic equilibrium in the game of choosing a forecast revision rate. Since rapid forecast-revision is a unique symmetrical strategic equilibrium, the market will fall into this equilibrium. Our explanation of the mechanisms creating this equilibrium is broadly consistent with various recent work (Delong et al. 1990b, Shiller 1989, Kirman 1991, Youssefmir, Huberman, and Hogg 1998).

Second, we investigate what happens when all traders use the optimal forecast-revision rate and the market is in its strategic equilibrium. We find that in this strategic equilibrium the market is noisy and risky, with a high variance in prices and high levels of technical trading. Thus, the strategic equilibrium defined by the optimal forecast-revision rate is inefficient compared to hypothetical equilibria with other forecast-revision rates, for excessive variance in prices and higher market noise in this equilibrium cause traders to earn significantly less wealth. Because of these diminished earnings, the market equilibrium is sub-optimal, yet the fact that this is a unique stable equilibrium makes the prospects of the market evolving towards higher efficiency dim. Our explanation of this result is that frequent revision of market-forecasting rules brings about high technical trading which then generates positive feedback in price streams, destabilizing prices and thus worsening every trader's predictions of future price movements. Because their forecasts are less accurate, all traders earn less than they would in a more efficient market. This explanation is also consistent with various recent work (Arthur 1988, Delong et al. 1990a, Delong et al. 1991,Youssefmir, Huberman, and Hogg 1998).

Section 2 below describes the Santa Fe Artificial Stock Market and explains how we use it to study market equilibria, section 3 identifies and characterizes the equilibria we find in the market, and sections 4 and 5 explain our findings and discuss their relevance to real financial markets.

2. The Model and Theoretical Framework

We focus on two main aspects of the decision making of traders: the formation of expectations of future prices, and the revision of those expectations over time. In the following two subsections, we first describe the agent-based model which carries out the formation and actual evolution of traders' expectations. Then we explain the game-theoretic framework we use to find the optimal rate of revising market expectations.

2.1. The Santa Fe Artificial Stock Market

This section briefly describes the Santa Fe Artificial Stock Market, developed by Brian Arthur, John Holland, Blake LeBaron, Richard Palmer, and Paul Taylor. More detailed descriptions are available elsewhere (Palmer et al. 1994, Arthur et al. 1997, LeBaron 1997). When mentioning model parameters below, we indicate the specific parameter values used in the present work with typewriter font inside brackets [`like this`].

The Santa Fe Artificial Stock Market is an agent-based model in which traders continually explore and develop rules to forecast future prices in a financial markets, buy and sell assets based on the predictions of their most accurate rules, and revise or discard these rules based on their accuracy. Each trader acts independently, but the returns to each trader depend on the decisions made simultaneously by all the other traders in the market.

The market contains a fixed number N [`25`] of traders each of whom is endowed with an initial sum [`10000`] of money (in arbitrary units). Time is discrete. Traders may invest their money either in a risk-free asset or a risky stock. The risk-free asset is perfectly elastic in supply and pays a constant interest rate r [`10%`]. The risky stock, of which there are a total of N shares, pays a stochastic dividend d_t that varies over time according to a stationary first-order autoregressive process with a fixed coefficient [`0.95`]. The past- and current-period realization of the dividend is known to the traders at the time they make their investment decisions.

At each time step each trader must decide to allocate her wealth between the risky stock and the risk-free asset. She does this by forecasting the price of the stock in the next time period with a certain forecasting rule (described below). Forecasts are used to make an investment decision through a standard risk aversion calculation. Each trader possesses a constant absolute risk-aversion (CARA) utility function of the form

$$U(W_{i,t+1}) = -\exp(-\lambda W_{i,t+1}) \tag{1}$$

where $W_{i,t+1}$ is the wealth of trader i at time $t+1$ and λ [`0.5`] is the trader's degree of risk aversion. In order to determine i's optimal stock holding $x_{i,t}$ at time t, this utility function is maximized subject to the following constraint:

$$W_{i,t+1} = x_{i,t}(p_{t+1} + d_{t+1}) + (1+r)(W_{i,t} - p_t x_{i,t}) \tag{2}$$

where $x_{i,t}$ is trader i's demand for the stock at time t. If we assume that trader i's predictions at time t of the next period's price and dividend are normally distributed with (conditional) mean and variance, $E[p_{t+1} + d_{t+1}]$ and $\sigma^2_{i,t,p+d}$, and if we assume that the distribution of forecasts is normal, then, as Arthur et al. (1997) explain, trader i's demand for the stock at time t should be:

$$x_{i,t} = \frac{E_{i,t}(p_{t+1} + d_{t+1}) - p_t(1+r)}{\lambda \sigma^2_{i,t,p+d}} \tag{3}$$

where p_t is the price at time period t. The bids and offers submitted by traders need not be integers; the stock is perfectly divisible.

The aggregate demand for the stock must equal the number of shares in the market. Traders submit their decisions to the market *specialist*—an extra agent in the market who functions as a market maker. The specialist collects bids and offers from traders. Since the total demand is not to exceed the number of shares available, the specialist determines the market clearing price by solving the following equation together with equation 3:

$$\sum_{i=1}^{N} x_{i,t} = N.$$

Note that this process assumes that the specialist has access to the demand functions of the agents in the market. We believe it is justified since market makers actually have a keen sense of the demand functions of traders in real world markets, and avoid using inventory to balance day-to-day demand.

Each trader has a set of [`100`] forecasting rules. Each forecasting rule in the set has the following form:

IF (the market meets state D_i) THEN $(a = a_j, b = b_l)$

where D_i is a description of the state of the market and a_j and b_l are the values of the forecasting variables a and b. The values of a and b are used to make a linear forecast of the next period's price and dividend using the equation:

$$E(p_{t+1} + d_{t+1}) = a(p_t + d_t) + b \tag{4}$$

The values of the variables a and b in a trader's initial set of forecasting rules are selected randomly from a uniform distribution of values centered on the values that would create a homogeneous rational-expectations equilibrium in the market (for details on this process, refer to Arthur et al., 1997). As time progresses, the traders discard ineffective forecasting rules and try out new forecasting rules, so the values of a and b in a trader's set of rules evolve, as described below.

A market descriptor D_i matches a state of the market by an analysis of price and dividend history. A market state consists of a set of market conditions, and a market descriptor is a boolean function of those market conditions. There are fourteen different market conditions that are used to define market states, so forecasting rules can distinguish 2^{14} different market states. A market descriptor is represented as an array of fourteen bits, corresponding to the fourteen market conditions, with `1` signaling that the condition in question obtains, `0` indicating that the condition fails, and `#` indicating that the condition is to be ignored.

The breadth and generality of a market descriptor depends positively on the number of `#` symbols in its market descriptor; descriptors with many `0`s and `1`s match more narrow and specific market states. As the traders' sets of forecasting rules evolve, the number of `0`s and `1`s in the rules can change, making the rules sensitive to either more specific or more general market states. An appropriate reflection of the complexity of the population of forecasting rules possessed by the traders is the number of market states that their rules can distinguish, which is related to the number of bits that are set to `0` or `1` in the rules' market descriptors.

The market conditions defining market states fall into two main categories: *technical* conditions and *fundamental* conditions. Technical market conditions pertain to the recent history of the stock price, and the bits reflecting technical conditions are called *technical* bits. Technical market conditions concern issues taking one of these two forms:

"Is the price greater than an n period moving-average of past prices?" where $n \in \{5, 20, 100, 500\}$.

"Is the price higher than it was n periods ago?" where $n \in \{5, 20\}$.

Fundamental market conditions pertain to the relationship between the stock's price and its fundamental value; the bits reflecting them are called *fundamental* bits. Fundamental market conditions all concern issues of this form:

"Is the price greater than n times its fundamental value?" where $n \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}, 1, \frac{9}{8}\}$.

(A third minor category with two market conditions have their corresponding bits set either always on and always off, reflecting the extent to which traders act on useless information.)

Forecasting rules with descriptors that use technical bits (i.e., with technical bits set to `0` or `1`) are called *technical* rules, and rules with no such bits set are called *fundamental* rules. Fundamental trading rules detect immediate over- or under-valuation of a stock; they are sensitive to the current price and dividend but ignore any trends in those quantities. Technical rules can detect recent patterns of increase or decrease in stock prices and might predict a continuation or reversal of the trend (depending on the associated values of a and b).

A simplified example might help clarify the structure of market forecasting rules. Suppose that there is a three-bit market descriptor.[1] The first bit corresponds to the fundamental market condition in which the price is 75% higher than its fundamental value, the second bit corresponds to the technical condition in which the price is greater than the 20-period moving average of past prices, and the third bit corresponds to the technical condition in which the price has gone up over the last fifty periods. Then the descriptor `#10` matches all those market states in which the price exceeds its 20-period moving average of past prices but it has not risen over the last 50 periods. Note that the `#` symbol makes this descriptor insensitive to whether the price is 75% greater than its fundamental value. Putting this together, the full decision rule

IF `#10` THEN $(a = 0.96, b = 0)$

has the following meaning: If the stock's price exceeds its 20-period moving average but has not risen over the past 50 periods, then the (price + dividend) forecast for the next period is 96% of the current period's price. Since this rule's market descriptor uses some technical bits, this is considered to be a technical trading rule.

If the market state in a given period matches the descriptor of a forecasting rule, the rule is said to be *activated*. Many of a trader's forecasting rules may be activated at a single time, thus giving the trader many possible forecasts to choose among. The forecasting rule that trader actually uses in chosen at random from among the currently activated rules with a probability proportional to the rule's accuracy. Once the trader has chosen a specific rule to use, the rule's a and b values determine the trader's forecast, which then determines her investment decision at that time. A forecasting rule's *accuracy* is defined as the moving-average of the variance of the error (the difference between the forecast price and the true price):

$$\text{accuracy} = (A \times (\text{old } v)) + ((1 - A) \times \text{current squared error}) \qquad (5)$$

with $A = \exp(-1/\tau(v))$ for a healing-time parameter $\tau(v)$. Accuracy is a weighted average of previous forecasting errors defined over a window of previous performance, and the healing-time parameter $\tau(v)$ specifies the size of this window.

[1] Recall that the forecasting rules in the model we study actually use fourteen-bit descriptors.

A genetic algorithm (GA) allows each trader's population of forecasting rules to evolve and improve. Whenever the GA is invoked, it substitutes new forecasting rules for a fraction [`12%`] of the weakest forecasting rules in the trader's pool of rules. A rule's strength indicates how well it has been performing, and it is similar to its accuracy except that the current estimated variance v of each rule's error (the difference between true prices and the rule's forecasted prices) is transformed using the relation:

$$\text{strength} = C - (v \times \text{bitcost} \times \text{specificity}) \tag{6}$$

where bitcost [`0.01`] is the cost of using a certain bit (in the real world, this would be the cost of acquiring new information), specificity is the number of conditions in the rule set either on or off, and C is a constant chosen to make the strength positive.

New rules are created from the more successful rules in the trader's pool of rules. Pairs of successful rules (chosen randomly with probability proportional to their strength) play the role of "parents" and the new rules that they spawn are modified versions of their parent rules. The forecasting parameters a and b of the offspring rules are a linear combination of the forecasting parameters of the parent rules. The offsprings' market descriptors are just like their parents, except for the effects of two genetic operations: with some probability [`0.1`] one-point crossover combines subparts of the two parent descriptors, and with another probability [`0.03`] mutation randomly and uniformly changes some bits in the market descriptors. New rules are all assigned an initial accuracy value equal to the average accuracy in the existing population of rules.

The operation of the GA may be compared to a real-world consultant. The GA is designed so that, over time, poorly performing rules are replaced by rules that are likely to perform better, much as a client following the advice of a consultant replaces poorly performing trading methods with those that are likely to be more profitable. The frequency with which the GA is invoked, then, corresponds to the rate at which trading methods are revised.

It is important to note that traders in this model learn in two ways: First, as each rule's accuracy varies from time period to time period, each trader preferentially uses the more accurate of the rules available to her; and, second, on an evolutionary time scale, the pool of rules as a whole improves through the action of the genetic algorithm.

2.2. A Method for Studying Market Equilibria

In this paper we focus on one particular aspect of a trader's strategy for trading in the market: the rate of revision of the market-forecasting rules used to predict the market's behavior and thereby determine the investment decision. This is an important decision faced by real traders. Since revising forecasting rules takes time and effort, using the right forecast-revision rate can increase earnings. But choosing the right forecast-revision rate is complicated. Since a trader's earnings is affected by the behavior of all other traders, her optimal revision-rate would depend on that of all other traders. Trading, however, is accompanied with a great deal of secrecy, and making assumptions of what all other traders are doing is difficult and risky.

To study how rational traders would choose their forecast-revision rate in the face of this uncertainty, we adopt a game-theoretic framework and ask if a trader has an optimal revision rate *regardless* of the revision rate of other traders. For the purposes of this investigation, a trader's *strategy* is the adoption of a specific market-forecast revision rate. Since the forecast-revision rate in the Santa Fe Artificial Stock Market is controlled by the interval at which the genetic algorithm (GA) is invoked for each trader, a trader's strategy in this model amounts to adopting a given GA-invocation interval, which we will call K. We investigate a trader's optimal GA-invocation interval K by evaluating different possible choices for K on the assumption that other traders all have adopted some given but unknown baseline GA-invocation interval, which we will call $\overline{K}$.[2] If we assume that there are S relevantly different possible GA-invocation intervals, then our trader confronts a classic $S \times S$ decision problem, with columns indexed by possible situations characterized by $\overline{K}$ and rows indexed by possible choices characterized by K. To make a rational decision in this context, a trader needs to know the relative value or payoff of all possible choices in all possible situations. Our criterion for social and individual welfare is a trader's final accumulated wealth, which includes wealth from all sources (interest payments from risk-free assets, returns from stocks, and cash holdings). So, to determine a trader's rational strategy in this decision matrix we need to observe the trader's final wealth in the $S \times S$ possible scenarios.

[2] For simplicity, we restrict our attention here to the case in which a trader can assume that all other traders are following the *same* strategy, defined by $\overline{K}$. We are in the process of examining the more general situation in which a trader assumes that the population of other traders might be following some *mixed* strategy defined by an arbitrary collection of different GA-invocation rates, and the preliminary results of this work corroborate our conclusions here.

The process of each trader attempting to find her optimal strategy creates a symmetric simultaneous-move N-person game, and we ask whether a trader in this game has a dominant strategy (i.e., a strategy that outperforms all other possible strategies no matter what strategies other traders follow). In a symmetric N-person game, a single trader's dominant strategy is also the dominant strategy of other traders. The existence of such a dominant strategy creates a symmetric Nash equilibrium in the market, i.e., an equilibrium in which no trader can gain more by unilaterally switching strategies, for all other traders are in a symmetrical situation; a dominant strategy for one is a dominant strategy for all.

To find the dominant strategy in the Santa Fe Artificial Stock with N traders, we fix the background GA-invocation rate $\overline{K}$ of $N-1$ traders and systematically vary the GA-invocation rate K of a single trader, observing how the single trader's wealth varies as a function of K given $\overline{K}$. This allows us to fill out the payoffs in the $\overline{K}$ column in the $S \times S$ decision matrix. We systematically do this for the columns corresponding to each possible baseline GA-invocation rate $\overline{K}$ in the population of $N-1$ traders. Other than varying the GA-invocation rates $\overline{K}$ and K, we hold all model parameters constant in all the market simulations.

To summarize the information in the $S \times S$ decision matrix in a reaction function for the single trader, we determine a trader's optimal GA-invocation interval K' for each baseline GA-invocation interval $\overline{K}$. The reaction function assigns to every value of $\overline{K}$ that GA-invocation interval K' at which a trader maximizes her wealth. In effect, we collapse the rows in the $S \times S$ table and record the row index (K value) of the highest value in each $\overline{K}$ column. This yields for each $\overline{K}$ that rate K' at which a trader's earnings are maximized.

There is a symmetric Nash equilibrium in the market wherever the population's baseline GA-invocation interval $\overline{K}$ equals the profit-maximizing GA-invocation interval K' of the single trader, i.e., at those values of $\overline{K}$ for which $\overline{K} = K'$. A symmetric Nash equilibrium is stable because each trader uses the same strategy to maximize her utility, *regardless* of what other traders in the market are doing. No matter which trader we consider, it is not possible for her to increase her utility by pursuing any other strategy, and so there is no incentive for her to leave the equilibrium.

The values of $\overline{K}$ that we studied were $\overline{K} = 5$, 10, 50, 100, 250 and 5000. For each value of $\overline{K}$, we examined about a dozen different GA-invocation rates K for the singular trader, and for each choice of $\overline{K}$ and K we ran about a half-dozen simulations of the market, with different random number seeds, observing the singular trader's earnings. The singular trader's optimal GA-invocation rate K' for each population baseline rate $\overline{K}$ was identified by finding her maximum earnings as a function of K.

Finally, we studied the properties of the strategic equilibria we found by comparing them with hypothetical equilibria at other GA-invocation intervals. Since a strategic equilibrium in our framework exists only when all traders revise their market forecasts at the *same* rate, we gave all traders the same GA-invocation interval and studied the market's behavior at various GA-invocation intervals. Specifically, at different GA-invocation intervals we observed the variance of the price stream, the complexity of the evolved forecasting rules, and the wealth earned by traders. The specific GA-invocation intervals at which we explored hypothetical equilibria were $K = 10$, 25, 50, 100, 250, 500, 750, 1000, 1500, 2500, 5000, and 10000.[3] At each value of K we simulated the market 30 different times, each time for 300,000 time steps.[4]

3. Results

3.1. A Unique Market Equilibrium Exists

Figure 3.1 presents the reaction function for an trader in the market, i.e., the trader's optimal GA-invocation interval K' given that all other traders are using some fixed background GA-invocation interval $\overline{K}$. This reaction function has three notable aspects.

First, when other traders in the market are using a relatively small rule-updating interval ($\overline{K} < 100$), i.e., they are revising their forecasts frequently, the single trader profits from using a larger interval, thus updating her forecasts *less* frequently. We believe the explanation for this is that most traders in this regime are updating their rules too frequently to be able to exploit the longer-term patterns in the market. By invoking the GA at a slower rate, the single trader is able to exploit these longer-term patterns and, since she does not have enough market-impact to alter those patterns that she detects, she profits significantly.

Second, when other traders in the market are using a large rule-updating interval, i.e., they are updating their forecasting rules relatively slowly ($\overline{K} > 100$), it is profitable for a single trader to revise here forecasts more frequently. Our explanation of this is similar to the previous case. The single trader detects the short-term patterns in the market while those traders who are updating their forecasts infrequently are incapable of exploiting these short-term patterns. By

[3] At significantly higher or lower values of K the market behavior is determined by special boundary conditions (Joshi and Bedau, 1998).

[4] Each simulation of the market took on average about one hour on a DEC Alpha.

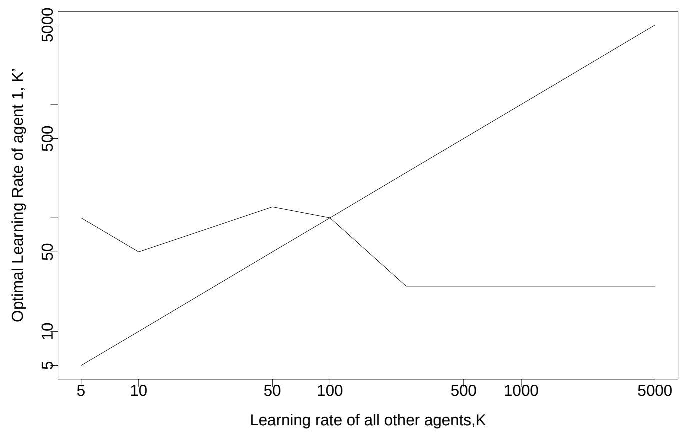

Figure 1. An trader's reaction function, indicating her optimal GA-invocation interval, K', given the baseline GA-invocation interval, $\overline{K}$, used by all other traders. The symmetric Nash equilibrium is where $\overline{K} = K' = 100$.

frequently updating her rules, the single trader however, can exploit these patterns. Since she is the only one who is exploiting them, she does not alter them, and thus earns significant profits.

The third and most important conclusion from the graph is that a trader's optimal GA-invocation interval K' matches the background GA-invocation interval $\overline{K}$ in the population at one and only one value: $K' = \overline{K} = 100$.[5] At this background GA-invocation interval, a single trader cannot gain by either raising or lowering her forecast-revision rate; her best choice is to keep updating her rules at the same rate as everyone else. But since all traders are in symmetric situations, the single trader's optimal strategy is also the optimal strategy of all other traders. The use of this strategy by all the traders in the market leads the market to a unique symmetric Nash equilibrium.

It is easy to see why the equilibrium described above is the market's one and only stable state. In all other states there is an incentive (usually, a strong incentive) for a trader to deviate from the baseline GA-invocation interval in the market. Since *all* traders feel this same incentive, and since all traders are self-interested, the traders' strategies will be in a stable and steady state *only* when they converge on the GA-invocation rate that maximizes earnings regardless of the strategies followed by other traders. For example, if one trader assumed

[5] We have not tried to determine what sets the time-scale at 100, nor how this time-scale depends on the parameters of the stock market.

that other traders were slow at revising their forecasts, it would be in her interest to start revising her forecasts faster. But since this is the choice that confronts everyone in the market, everyone would start revising their forecasts faster. A similar situation would occur when a single trader assumed that everyone else is revising their forecasts very fast. She would slow down her revision rate, but for the same reasons, so would everyone else in the market. The only strategic equilibrium occurs when no trader has any incentive left to change her strategy from that adopted by the rest of the population. This is the equilibrium we observe at $K = 100$.

It is important to note that the fact that the market is at a strategic equilibrium does *not* imply that prices, trading volumes or any other aspect of the market besides those strategies are stable, nor does it imply that the market is in any sense efficient. In fact, in the following section we show that many important aspects of the market's behavior in this equilibrium are unstable and this makes the market inefficient.

3.2. The Market Equilibrium Is Sub-Optimal

Figures 3.2 and 3.2, show how the variance of the price stream and the complexity of the trader's evolved market-forecasting rules vary as a function of the rate at which forecasting rules are evolving, i.e., the GA-invocation rate, K. The general pattern in these data is consistent with previous work (Palmer et al. 1994, Arthur et al. 1997, Joshi and Bedau 1998, LeBaron, Arthur, and Palmer 1998).

The most significant earlier finding was that this market exhibits two quite different kinds of behavior, corresponding to different rates at which market-forecasting rules are being revised by the genetic algorithm (GA). Although there is a gray area of intermediate GA-invocation rates at which the market's behavior is somewhat ambiguous, this gray area separates two clearly different kinds of markets. When the GA-invocation interval is large ($1000 \leq K \leq 10,000$) so forecasting rules are evolving relatively slowly, prices are more stable (Fig. 3.2), evolved forecasting rules are less complex (Fig. 3.2), and the levels of technical trading are significantly lower (Fig. 3.2). In addition, earlier work has shown that trading volumes are low and there is little evidence of non-linearity, lepto-kurtosis, volatility persistence or volume-volatility correlation when the GA-invocation interval is high (Palmer et al. 1994, Arthur et al. 1997, LeBaron, Arthur, and Palmer 1998). Since this kind of behavior resembles the predictions of the theory of efficient markets, this regime of the Santa Fe Artificial Stock Market has been termed “Rational Expectations Regime”.

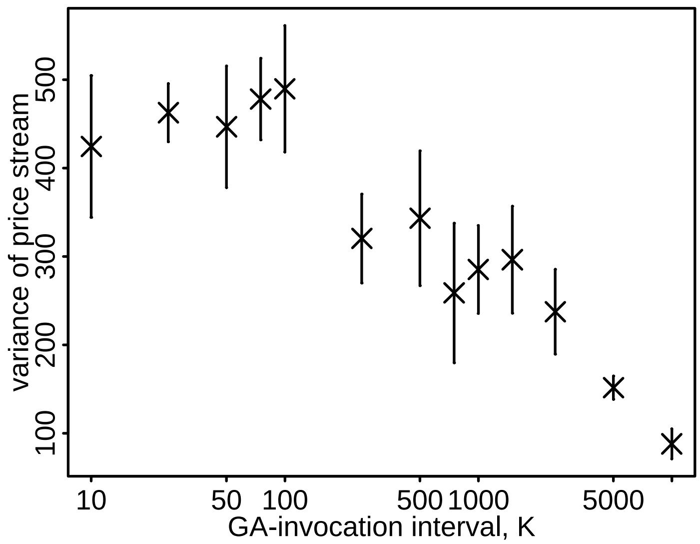

Figure 2. Variance of the stock price time series as a function of GA-invocation interval, K. The data shown averages values from the last 50,000 time steps in 30 simulations each lasting 300,000 time steps; error bars indicate one standard deviation. Note that the price stream variance is highest when the GA is being invoked frequently, i.e., in the Complex regime, roughly $10 \leq K \leq 100$, and that the variance is significantly lower when the GA is being invoked infrequently, i.e., in the RE regime, roughly $1000 \leq K \leq 10,000$.

On the other hand, when the GA-invocation interval is small ($10 \leq K \leq 100$) so forecasting rules are evolving relatively quickly, the variance of the price time series is relatively high (Fig. 3.2). In addition, the complexity of the forecasting rules produced by the course of evolution is relatively high (Fig. 3.2), as is the level of technical trading (Fig. 3.2). Earlier work has also shown that trading volumes are higher and there is evidence of non-linearity, lepto-kurtosis, persistence of volatility, and significant volume-volatility correlation when the GA-invocation interval is small and market forecasts are being revised quickly (Palmer et al. 1994, Arthur et al. 1997, LeBaron, Arthur, and Palmer 1998), so this regime of the Santa Fe Artificial Stock Market has been termed "Complex". For convenience and consistency with this earlier work, we will adopt this terminology of "Complex" and "RE" markets here.

The unique strategic equilibrium in the rate of forecast evolution at $K = 100$ clearly falls within the Complex regime of the market ($10 \leq K \leq 100$). What is important to note here is that this Complex equilibrium is characterized by much lower earned wealth than in the RE regime (Fig. 3.2). Thus, in a straightforward sense, this market equi-

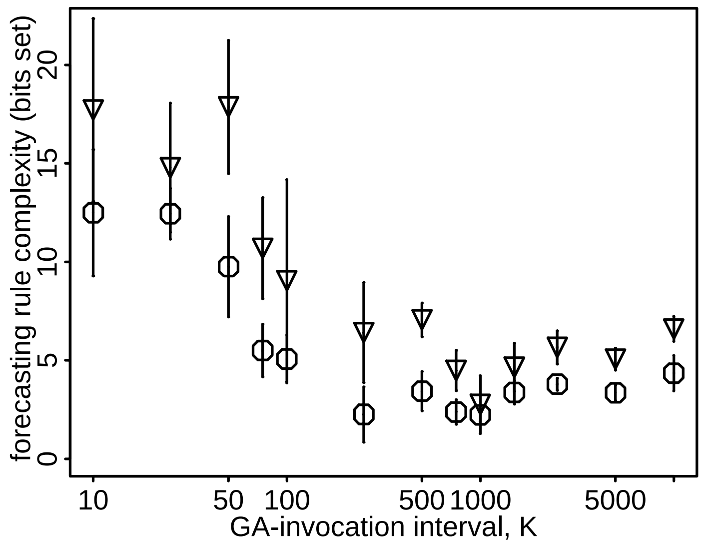

Figure 3. Complexity of evolved market-forecasting rules as a function of GA-invocation interval, K. The data shown averages values from the last 50,000 time steps in 30 simulations each lasting 300,000 time steps; error bars, which overlap for some values of K, indicate one standard deviation. Forecasting-rule complexity is measured by the number of technical trading bits (triangles) and fundamental trading bits (circles) set in the market descriptors of the forecasting rules. The number of bits set is normalized (i.e., divided) by the total number of bits available. At the start of a simulation before the forecasting rules start evolving, the average number of technical and trading bits set is 10 (slightly below the midline in this plot). Note that evolution raises the forecasting rule complexity when the GA is being invoked frequently, i.e., in the Complex regime, roughly $10 \leq K \leq 100$, and that evolution lowers the forecasting rule complexity when the GA is being invoked infrequently, i.e., in the RE regime, roughly $1000 \leq K \leq 10,000$.

librium is sub-optimal. One can justify the claim that this equilibrium is sub-optimal in some other ways as well. For one thing, the relatively high variance of the price stream makes all investments relatively risky. Also, since the complexity of the evolved forecasting methods is significantly higher, it is reasonable to assume that the cost of formulating and testing these forecasting methods would also be higher. But the bottom line is that everyone in this symmetric Nash equilibrium is earning less money. Clearly, the equilibrium is sub-optimal.

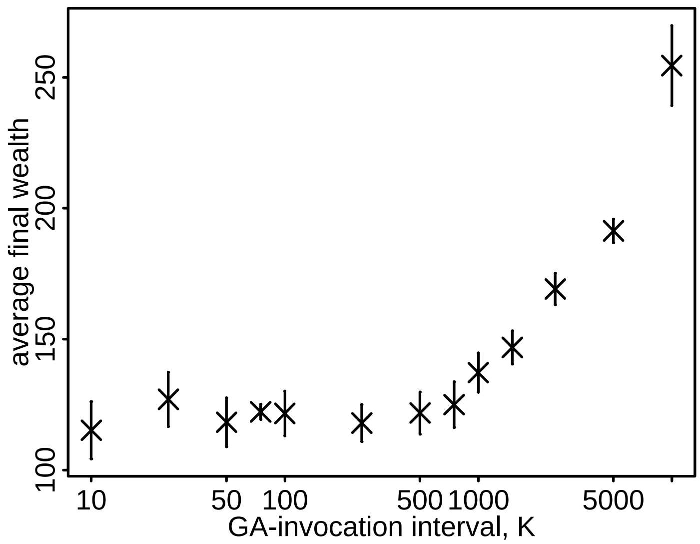

Figure 4. Accumulated wealth of traders in the market as a function of GA-invocation interval, K. The data shown averages values from the last 50,000 time steps in 30 simulations each lasting 300,000 time steps; error bars indicate one standard deviation. (Wealth values shown are in 10,000s of the arbitrary monetary units used in this model.) Note that wealth is lowest when the GA is being invoked frequently, i.e., in the Complex regime, roughly $10 \leq K \leq 100$, and that wealth is significantly higher when the GA is being invoked infrequently, i.e., in the RE regime, roughly $1000 \leq K \leq 10,000$. Comparison with Fig. 3.2 shows that wealth is roughly inversely proportional to price variance.

4. Discussion

If the unique strategic equilibrium in this market is inefficient and suboptimal, as we have argued, then two important questions arise: First, exactly what causes the reduced wealth at the Nash equilibrium? Second, is there any way that the market can be moved to a more efficient outcome?

Our answer to the first question is presented in more detail elsewhere (Joshi, Parker and Bedau, 1998), but the following is a summary. When traders revise their forecasting rules very frequently, levels of technical trading are high (Palmer et al. 1994, Arthur et al. 1997, Joshi and Bedau 1998, LeBaron, Arthur, and Palmer 1998). As more traders adopt technical rules, the incentives for technical trading can reinforce themselves in a new way. In effect, if enough traders in the market buy into similar enough technical trading rules, positive feedback can make the rules self-fulfilling prophecies. For example, if all traders *believe* that the price of a stock will go up, they will all want to buy the stock, creating an excess demand and driving its price up—thereby making their belief in a price increase true. But now, the self-fulfilling prophecies created by technical trading dramatically increase the variability of prices in the market, causing bubbles and crashes. This increased market noise decreases the accuracy of *all* forecasting rules, and this, in turn, drives down all traders' wealth because less accurate rules are less profitable. The gains from self-reinforcing technical trends are short lived; in the long run, correction toward fundamental value bursts the bubbles. In other words, technical trading creates a negative externality in the market by altering the patterns that market forecasts exploit. It worsens everyone's forecasts by driving prices away from their fundamental value and increasing noise. When all traders engage in significant technical trading, they worsen each other's forecasts, there is a loss of efficiency, and the average earnings in the market are lowered.

The second question is whether the socially optimal outcome (where all traders revise their market-forecasting rules infrequently) can ever prevail in the market. Our answer to this question is that this is highly unlikely given the large number of players that are interacting with each other. Traders are generally self-interested, desiring to maximize their own profits rather than the average profits in the market as a whole. It would be almost impossible to induce all traders in the market to co-operate by slowing down the rate at which they revise their market-forecasting rules, because it is not in their individual interests to do so. The situation is analogous to an N-person prisoner's dilemma: when each trader acts independently and rationally, everyone makes the same choice and a sub-optimal outcome prevails. And although co-operation can spontaneously emerge among self-interested players in iterated two-person prisoner's dilemmas (Axelrod, 1984), in analogous iterated N-person games co-operation is almost impossible to achieve (Lindgren, 1998).

It is important to keep in mind that the problem of determining the optimal forecast-revision rate is not posed or solved by the traders in the artificial market we simulate. The revision rate for each trader in the model is fixed for the duration of each simulation, so the model does not explain the dynamical process by which the traders would converge on the optimal revision rate. Nevertheless, by studying the long-term effects of various combinations of revision rates, we have discovered that frequent forecast revision is a stable and unique strategic equilibrium. Given that this equilibrium exists, it is reasonable to expect that it would be discovered in a decentralized market with boundedly rational traders provided they could try different revision rates and compare their long-term payoffs. Investigating the dynamical process of converging on the equilibrium is a natural and important topic for future work, and it could be studied by modifying the Santa Fe Artificial Stock Market to allow the traders to choose their learning rate according to some learning algorithm.

5. Summary and Conclusion

Our simulations using the Santa Fe Artificial Stock Market suggest that financial markets can end up in situations analogous to prisoner's dilemmas in which frequent revision of forecasting rules and extensive technical trading lead to increased price variability and thus reduced earnings. When traders do not know *a priori* what other traders are doing, their optimal strategy is revise forecasting rules frequently. But when this dominant strategy prevails and market beliefs co-evolve rapidly, the market falls into a symmetric Nash equilibrium with relatively low average earnings for traders. This happens because frequent forecast-revision creates high levels of technical trading, and this creates a negative externality in the market by causing positive-feedback and destabilizing prices, thus decreasing all traders' earnings.

Though the Santa Fe Artificial Stock Market and game-theoretic framework considered in this paper are both simplifications of real-world markets, we believe that they capture some important elements of decision making in those markets. Both the Santa Fe Artificial Stock Market and our game-theoretic analysis mimic the uncertainty, bounded rationality, and agent heterogeneity that underlie learning and decision making in most real markets. The genetic algorithm in the model we study is a mechanical yet sophisticated process for learning better market-forecasting rules. Our analysis shows that there is a strategic equilibrium in the rate of this learning process, one which mirrors some of those very aspects of real markets that violate the predictions of traditional market models.

Much research remains to be done in establishing the robustness of these results to variations both in the model's parameters and in the structural design of the model itself. It is not yet clear to what extent our results depend on the particular model parameters we used, nor is it yet clear to what extent changing the structure of the model would change our results. In related work we have separated the effects of technical trading from the forecast-revision rate and shown that technical trading by itself can create an analogous prisoner's dilemma (Joshi, Parker and Bedau, 1998), and our current work includes investigating the endogenous evolution of the forecast-revision rate, to see if the market finds the strategic equilibrium we identified here. In general, it would be very interesting to use a variety of theoretical and empirical methods to investigate the effects of market-forecast revision rates on market equilibria in other artificial models and in real markets. All of this research acquires a special importance from the potentially quite significant conclusion of the initial explorations in the artificial market reported here: that even though traders would be better off if this could be prevented, frequent revision of market-forecasting rules seems to be inevitable.

Acknowledgements

Thanks to the authors of the Santa Fe Artificial Stock Market, especially Richard Palmer and Blake LeBaron, for making their source code available to us and helping us to use it productively. For helpful discussion or comments on previous drafts, thanks to Doyne Farmer, Blake LeBaron, Michael de la Maza, John Moody, Chris Moore, Norman Packard, Richard Palmer, Stephanie Forrest, Robert Seymour, and two anonymous reviewers for this journal. Thanks also to the Santa Fe Institute for support to SJ and MAB during the final stages of this project.

References

Acar, E. and S. Satchell, 1997, *Advanced Trading Rules* (Butterworth-Heinemann: Woburn).

Arthur, W.B., 1988, Self-reinforcing mechanisms in economics, in: P.W. Anderson, K. J. Arrow and D. Pines, eds., *The Economy as an Evolving, Complex System I: Proceedings of the Evolutionary Paths of the Global Economy Workshop* (Addison-Wesley: Redwood City) 15-44.

Arthur, W.B., 1989, Positive feedbacks in the economy, *Scientific American* (February), 92–99.

Arthur, W.B., 1992, On learning and adaptation in the economy, Working Paper 92-07-038, Santa Fe Institute, Santa Fe NM.

Arthur, W.B., J.H. Holland, B. LeBaron, R. Palmer, and P. Tayler, 1997, Asset pricing under endogenous expectations in an artificial stock market, in: W. B. Arthur, D. Lane, and S.N. Durlauf, eds., *The Economy as an Evolving, Complex System II* (Menlo Park: Addison-Wesley).

Axelrod, R., 1984, *The Evolution of Cooperation* (New York: Basic Books).

Bierman, S.H, and L. Fernandez, 1993, *Game Theory with Economic Applications* (Redwood City: Addison-Wesley).

Bray, M., 1981, Futures trading, rational expectations, and efficient markets hypothesis, *Econometrica* 49, 575–96.

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

Campbell, J.Y. and A. Lo., 1997, *The Econometrics of Financial Markets* (Princeton: Princeton University Press).

Chancellor, E., 1999, *Devil Take Hindmost: History of Financial Spceculation* (Farrar Straus and Giroux).

Cootner, P., 1967, *The Random Character of Stock Market Prices* (Cambridge: MIT Press).

Delong, J.B., A. Schleifer, L.H. Summers, and J. Waldmann, 1990a, Positive feedback and destabilizing rational speculation, *Journal of Finance* 45, 379–395.

Delong, J.B., and A. Schleifer, L.H. Summers and J. Waldmann, 1990b, Noise trader risk in financial markets, *Journal of Political Economy* 98, 703–738.

Delong, J.B., A. Schleifer, L.H. Summers, and J. Waldmann, 1991, The survival of noise traders in financial markets, *Journal of Business* 64, 1–18.

Fama, E.F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–420.

Farmer, R.A., 1993, *The Macroeconomics of Self-Fulfilling Prophecies* (Cambridge: MIT Press).

Farmer, J.D., 1998, Market force, ecology and evolution, Working Paper 98-12-116, Santa Fe Institute, Santa Fe NM.

Frankel, J.A., and K.A. Froot, 1990, Chartists, fundamentalists, and trading in the foreign exchange market, *AEA Papers and Proceedings* 80, 181–185.

Joshi, S., and M.A. Bedau, 1998, An explanation of generic behavior in an evolving financial market, in: R. Standish, B. Henry, S. Watt, R. Marks, R. Stocker, D. Green, S. Keen, and T. Bossomaier, eds., *Complex Systems '98, Complexity Between the Ecos: From Ecology to Economics* (Complexity Online: Sydney) 326-332.

Joshi, S., J. Parker, and M.A. Bedau, 1998, Technical trading creates a prisoner's dilemma: Results from an agent-based model, Working Paper No. 98-12-115, Santa Fe Institute, Santa Fe NM, to appear in the proceedings of *Conference on Computational Finance 1999.*

Keim, D.B. and A. Madhaven, 1995, Anatomy of the trading process: Empirical evidence of the behavior of institutional traders. *Journal of Financial Economics* 37, 371-398.

Keynes, J.M., 1936, *The General Theory of Employment, Interest and Money* (New York: Harcourt, Brace). Reprinted in 1973 (New York: W.W. Norton).

Kirman, A., 1991, Epidemics of opinion and speculative bubbles in financial markets, in: M. Taylor, ed., *Money and Financial Markets* (London: Macmillan).

Kurtz, M., 1994, On the structure and diversity of rational beliefs, *Economic Theory*, 4: 877–900.

LeBaron, B., W.B. Arthur, and R. Palmer, 1998, Time series properties of an artificial stock market, *Journal of Economic Dynamics and Control*, forthcoming.

LeBaron, B., 1997, Technical trading rules and regime shifts in foreign exchange markets, in: E. Acar and S. Satchell, eds., *Advanced Trading Rules* (Woburn: Butterworth-Heinemann).

Lindgren, K., 1998, Coevolution of strategies in N-person prisoners dilemma games, presented at the workshop *Evolutionary Dynamics—Exploring the Interplay of Selection, Neutrality, Accident, and Function*, Santa Fe Institute, Santa Fe, NM, October 5-9, 1998.

Lo, A. W., and A. C. Mackinlay, 1999, *A Non-Random Walk Down Wall Street* (Princeton University Press).

Malkiel, B., 1992, Efficient market hypothesis, in: P. Newman, M. Milgate, and J. Eatwell, eds., *New Palgrave Dictionary of Money and Finance* (London: Macmillan Press Limited).

Marengo, L., and H. Tordjman, 1995, Speculation, heterogeneity, and learning: A model of exchange rate dynamics, Working Paper WP-95-17, IIASA.

Palmer, R.G., W.B. Arthur, J.H. Holland, B. LeBaron, and P. Tayler, 1994, Artificial economic life: a simple model of a stock market, *Physica D* 75, 264–274.

Rieck, C., 1994, Evolutionary simulation of asset trading strategies, in: E. Hillenbrand and J. Stender, eds., *Many Agent Simulation and Artificial Life* (Washington DC: IOS Press).

Samuelson, P., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* 6, 41-49.

Sargent, T.J., 1993, *Bounded Rationality and Macroeconomics* (Clarendon Press: Oxford).

Shleifer A. and L.H. Summers, 1990, The noise trader approach to finance, *Journal of Economic Perspectives* 4, 19–33.

Schwager, J.D., 1995, *Technical Analysis* (New York: John Wiley and Sons).

Shiller, R., 1989, *Market Volatility* (Cambridge: MIT Press).

Soros, G., 1987, *The Alchemy of Finance* (New York:Simon and Schuster).

Soros, G., 1994, *The Theory of Reflexivity* (New York: Soros Fund Management).

Summers, L.H., 1986, Does the market rationally reflect fundamental values?, *Journal of Finance* 46: 591–601.

Werner, F.M., deBondt, and R. H. Thaler, 1995, Financial decision making in markers and firms: A behavioral perspective, in: R. Jarrow et al., eds., *Handbooks in Operations Research and Management Science*, Vol. 9 (Elsevier Science).

Youssefmir M., B. Huberman, and T. Hogg, 1998, Bubbles and market crashes, *Computational Economics* 12, 97-114.