

Technical Trading Creates a Prisoner's Dilemma: Results from an Agent-Based Model

Shareen Joshi
Jeffrey Parker
Mark A. Bedau

SFI WORKING PAPER: 1998-12-115

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu



SANTA FE INSTITUTE

Technical Trading Creates a Prisoner's Dilemma: Results from an Agent-Based Model

Shareen Joshi^{†‡}, Jeffrey Parker[†], and Mark A. Bedau[†]

[†]Reed College, 3203 SE Woodstock Blvd., Portland OR 97202

[‡]Santa Fe Institute, 1399 Hyde Park Rd., Santa Fe NM 87501

December 2, 1998

Abstract

The widespread use and proven profitability of technical trading rules in financial markets has long been a puzzle in academic finance. In this paper we show, using an agent-based model of an evolving stock market, that widespread technical trading can arise due to a multi-person prisoners' dilemma in which the inclusion of technical trading rules to a single agent's repertoire of rules is a dominant strategy. The use of this dominant strategy by all traders in the market creates a symmetric Nash equilibrium in which wealth earned is lower and the volatility of prices is higher than in the hypothetical case in which all agents rely only on fundamental rules. Our explanation of this lower wealth and higher volatility is that the use of technical trading rules worsens the accuracy of the predictions of all agents' market forecasts by contributing to the reinforcement of price trends, augmenting volatility, and increasing the amount of noise in the market.

1 Introduction

Technical and fundamental forecasting rules are widely used by traders in financial markets. While fundamental rules are based on the assumption that prices should stay close to their true worth (the discounted worth of future returns), technical rules are based on the assumption that prices move in predictable historical patterns [30, 31].

The usefulness of fundamental trading rules is adequately explained by the standard theory of efficient markets [8, 17, 12, 14]. In informationally efficient markets consisting of homogeneous agents with rational expectations, the theory predicts that prices should closely track fundamental values, and so it should be possible for an agent to make higher-than-normal profits only in the case

that she successfully forecasts changes in stock fundamentals. Thus fundamental rules might provide accurate forecasts of stock values. Technical trading rules on the other hand, can not be useful forecasters in such markets. Since historical patterns yield no useful information about future stock prices beyond their implications for the stream of dividends, attempting to forecast future prices based on historical patterns should not be profitable. Thus, traditional theoretical models have a difficult time explaining the proven profitability of technical trading rules in financial markets [32, 16, 7, 33].

A wide variety of theoretical and empirical models have been developed to explain why technical trading is widespread in financial markets [18, 19, 29, 32, 13, 11, 23, 26]. This paper uses evidence from an agent-based artificial model of a stock market to explore an explanation of this phenomenon. The key characteristic of this model is that agents' expectations do not follow a fixed rule such as a rational expectations rule. Instead agents choose among an evolving set of expectation rules depending on which ones have proved to be the most successful predictors of recent stock-price changes.

Using this framework, we show first that in a market in which all other traders follow strictly fundamental rules of the kind that would characterize an efficient-market equilibrium, an individual agent might gain from adding technical trading rules to her repertoire of forecasting techniques. Second, using a game theoretic analysis, we show that while the use of technical trading rules (in addition to fundamental rules) is the optimal strategy of a single agent, the use of this strategies by all agents in the market drives the market to a symmetric Nash equilibrium at which wealth is lower for all agents than in a hypothetical equilibrium where all agents use only fundamental rules.

Our explanation of this phenomenon is that the adoption of technical-trading strategies by all agents in the market adds to the noise in the market and thus make it more difficult for everyone to predict future stock-price movements than in a regime where only fundamental rules are used. Because their predictions are less accurate, the presence of technical trading makes agents in the market worse off. Thus, our results suggest that technical trading leads to a prisoners' dilemma in which individual decisions lead to an inefficient social outcome.

Section 2 below describes the Santa Fe Artificial Stock Market model that we use in our argument, section 3 explains our experimental framework, sections 4 and 5 present and explain the results of our experiment, and section 6 concludes by explaining the relevance of these results to the real world.

2 The Santa Fe Artificial Stock Market

The Santa Fe Artificial Stock Market described in this paper was developed by Brian Arthur, John Holland, Blake LeBaron, Richard Palmer, and Paul Taylor [27, 3]. It is an agent-based artificial model in which agents continually explore and develop forecasting models, buy and sell assets based on the predictions of those models that perform best, and confirm or discard these models based on their performance over time. At each time period in the market, each agent

acts independently, following her currently best models, but the returns to each agent depend on the decisions made simultaneously by all the other agents in the market.

The following sections provide a brief introduction to the Santa Fe Artificial Stock Market model. More detailed descriptions are available elsewhere [27, 3, 24, 20]. When mentioning model parameters below, we indicate the specific parameter values we used in the work reported here with typewriter font inside brackets [`like this`].

2.1 The Market

The market contains a fixed number N [25] of agents each of whom is endowed with an initial sum of money (in arbitrary units) [10000]. Time is discrete. At a given time period each agent decides how much of her money to invest in a risky stock and how much to invest in a risk-free asset. 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 agents at the time they make their investment decisions.

2.2 Agents and Market Forecasting Rules

Agents must make decisions to allocate their wealth between the risky stock and the risk-free asset. They do this by forecasting the price of the stock in the next time period. This forecast, in turn, is based on one of a set of [100] rules, each of which has the following form:

IF (the market meets condition D_i) THEN ($a = k_j, b = k_l$)

where D_i is a description of the state of the market, k_j and k_l are constants, and a and b are forecasting parameters. The values of the variables a and b are used to make a linear forecast of next period's price using the equation:

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

The forecasting parameters a and b are initially selected randomly from a uniform distribution of values centered on the values that would create a homogeneous rational-expectations equilibrium in the market [3].

Market descriptors $\{D_i\}$ match certain states of the market by an analysis of the price and dividend history. The descriptors are boolean functions that relate to a number [14] of "market conditions." The response of each rule to the set of market conditions is represented as an array of bits in which 1 signals that the rule is to be used only if that condition is true, 0 indicates that the rule is to be used only if the condition is false, and # indicates that that condition is irrelevant for the application of the rule.¹

¹Since there are 14 boolean market descriptors, it is possible to distinguish 2^{14} different market states.

The breadth and generality of the market states to which a specific rule applies depends positively on the number of # symbols in its market descriptor; rules with descriptors with many 0s and 1s recognize more narrow and specific market states. As the market evolves, these strings are modified periodically by a genetic algorithm, so the number of 0s and 1s might go up or down, allowing them to respond to more specific or general market conditions. An appropriate reflection of the complexity of the population of forecasting rules possessed by the agents is the number of specific market states that their rules can distinguish. This is measured by calculating the number of bits that are set to 0 or 1 in the rules' market descriptors.

There are two main categories of market conditions to which descriptors are attached. One pertains to the recent history of the stock price; the descriptors associated with these conditions are called *technical* trading bits. The other main kind of conditions pertains to the relationship between the stock's price and its fundamental value; the descriptors of these conditions are called *fundamental* trading bits. (Two additional condition bits were set at "always on" and "always off" to reflect the extent to which agents act on useless information.) While trading rules based solely on fundamental conditions and descriptors detect immediate over- or under-valuation of a stock, technical trading rules 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).

The market conditions corresponding to the descriptors in the technical forecasting rules (i.e., rules with some fundamental trading bits set) take 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 $n \in \{5, 20\}$.

The conditions in the fundamental rules (i.e., rules with only fundamental bits set) all take this form:

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

In the present context, then, fundamental rules are sensitive to only current prices and dividends; they ignore any trends in those quantities. Only technical rules can detect any patterns in the market.

In an equilibrium corresponding to the predictions of the efficient markets theory, agents would use only an optimal fundamental rule (based on the actual parameters of the time-series process driving dividends) which would outperform all rules based on technical conditions. Our model differs from this in that agents do not know the parameters of the dividend process, and thus they must experiment with alternative forecasting rules based on fundamental (and perhaps technical) conditions in seeking to improve their forecasts.

An example might help clarify the structure of market forecasting rules. Suppose that there is a two-bit market descriptor, the first bit of which corresponds to the market condition in which the price has gone up over the last fifty periods, and the second bit of which corresponds to the market condition in which the price was 75% higher than its fundamental value. Then the descriptor 10 matches any market state in which the stock price has gone up for the past fifty periods and the stock price is not 75% higher than its fundamental value. The full decision rule

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

can be interpreted as “If the stock’s price has risen for the past fifty periods and is now not 75% higher than its fundamental value, then the (price + dividend) forecast for the next period is 96% of the current period’s price.”²

If the market state in a given period matches the set of descriptors of a forecasting rule, the rule is said to be *activated*. A number of an agent’s forecasting rules may be activated at a given time, thus giving the agent many possible forecasts among which to choose. The agent decides which of the active forecasts to use by measuring each rule’s accuracy and then choosing at random from among the active forecasts with a probability proportional to accuracy. Once the agent has chosen a specific rule to use, the rule’s a and b values provide a forecast of next period’s price.

Forecasts are used to make an investment decision through a standard risk aversion calculation. Each agent possesses a constant absolute risk-aversion (CARA) utility function of the form

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

where $W_{i,t+1}$ is the wealth of agent i at time $t + 1$, and $0 < \lambda[0.5] \leq 1000$. This utility function is maximized subject to the following constraint:

$$W_{i,t+1} = x_{i,t}(p_{t+1} + d_{t+1}) + (1 + r_f)(W_{i,t} - p_t x_{i,t})$$

where $x_{i,t}$ is agent i ’s the demand for the stock at time period t . Under the assumption that agent 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 $\sigma_{i,t,p+d}^2$, and the distribution of forecasts is normal, agent i ’s demand for the stock at time t , should be [3]:

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

The bids and offers submitted by agents need not be integers; the stock is perfectly divisible. The aggregate demand for the stock must equal number of shares in the market.

²Recall that the actual descriptors used in our modeling exercises contain 14 descriptors rather than just two.

Agents 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 agents, announces a ‘trial price’, and if the market does not clear or if his inventory does not stay within acceptable bounds, repeats this process. When the market clears, the ‘trial price’ becomes the current period’s market price.

2.3 Evolution of Market Forecasting Rules

A genetic algorithm (GA) provides for the evolution of the population of forecasting rules over time. Whenever the GA is invoked, it substitutes new forecasting rules for a fraction [12%] of the least fit forecasting rules in each agent’s pool of rules. A rule’s success, or “fitness” is determined by both how accurately it has forecasted prices and by how complex it is (the GA has a bias against complex rules). New rules are created by first applying the genetic operators of mutation, crossover, and inversion to the bit strings of the more successful rules in the agent’s rule pool. The forecasting parameters a and b of the offspring are a linear combination of the forecasting parameters of the parent rules. New rules are assigned an initial accuracy rating by averaging the accuracy of their parent rules.

The GA may be compared to a real-world consultant. It replaces current poorly performing rules with rules that are likely to perform better much the same way as a consultant urges her client to replace poorly performing trading strategies with those that are likely to be more profitable.

It is important to note that agents in this model learn in two ways: First, as each rule’s accuracy varies from time period to time period, each agent 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.

3 Experimental Methods

In this paper we study one particular aspect of an agent’s general strategy for trading in the market: whether technical rules should be included in her collection of forecasting rules. So, in this framework an agent’s *strategy* is either to include technical trading rules in her repertoire of trading rules, or to exclude them entirely and instead use only fundamental rules. We restrict our attention to just these two strategies to make our argument simple but realistic. In particular, we exclude the strategy of using *only* technical rules as unrealistic; no matter how much faith people have in technical trading rules, they generally seem to take economic fundamentals into consideration as well.

To investigate whether or not including technical trading rules is advantageous for traders, we contemplate a single agent confronted with a choice between our two strategies. The agent assumes that other traders in the market all follow one or the other of these two strategies—either all include technical trading rules or all exclude them—but the agent does not know which of

these two possibilities occurs. Thus, the agent confronts a classic 2×2 decision problem.

To make a rational decision, the agent needs to know the relative value or payoff of each choice in each situation. Our criterion for social and individual welfare is terminal or final wealth.³ So, to determine the payoffs in the decision matrix, we observed the final wealth of the agent in four different conditions:

- A** The agent *includes* technical rules and all other traders *include* them.
- B** The agent *includes* technical rules and all other traders *exclude* them.
- C** The agent *excludes* technical rules and all other traders *include* them.
- D** The agent *excludes* technical rules and all other traders *exclude* them.

By comparing the agent's payoffs in these four possible situations, we can determine whether there is a dominant strategy for this decision.⁴

Note that, since all agents in the market act independently and simultaneously, each time period in the market can be considered to be a multi-person simultaneous-move game. Furthermore, each agent's decision can be construed in exactly the form of the single agent considered above. So, if the single-agent decision considered above has a dominant strategy, it will be rational for all agents to use it and the simultaneous-move game will reach a symmetric Nash equilibrium [5]. Thus, situations **A** and **B** above are the only potential symmetric Nash equilibria in our context.

Expected payoffs in situations **A–D** were determined by simulating the artificial market 45 times in the four corresponding circumstances. In each simulation, there were 26 agents in the market: one agent following a given strategy and 25 other agents all following another given strategy (possibly the same strategy as the single agent). Each simulation was run for 300,000 time periods to allow the asymptotic properties of the market to emerge and to reduce the dependence of the results on initial conditions. The same 45 random sequences for dividends and initial distributions of rule descriptors among agents were used for all four experiments.

Previous work has shown that the evolutionary learning rate is a crucial parameter controlling the behavior of this model. All our simulations here were carried out at a learning rate of 100, i.e., with the genetic algorithm invoked for each agent once every 100 time periods. We chose this learning rate for two related reasons. First, we wanted to insure that agents had a realistic possibility of *using* technical trading rules. Since previous work [27, 3, 24, 20] has firmly established that high (statistically significant) technical trading actually occurs in the market only at learning rates in this neighborhood, our experimental design requires us to use such a rate. Furthermore, recent work [21] has shown

³The final wealth of an agent in the market includes wealth from all sources: interest payments from the risk free asset, returns from stocks, and cash holdings (money not invested).

⁴A dominant strategy is defined as one that outperforms all other strategies *regardless* of the strategies being used by other agents [5].

		ALL OTHER TRADERS	
		technical rules included	technical rules excluded
THE AGENT	include technical rules	A: 113 ± 6.99	B: 154 ± 6.68
	exclude technical rules	C: 97 ± 6.68	D: 137 ± 5.10

Table 1: The decision table for an agent contemplating whether to include technical trading rules to make her market forecasts, when she is uncertain whether or not the other traders in the market are doing so. The agent’s payoff in each of the four situations **A–D** is her expected final wealth (divided by 10^4 , to make more readable), derived by averaging the results of 45 simulations of each situation. Errors bounds are calculated using standard deviations of the 45 simulations.

that agents will *choose* this learning rate if given the choice, for this learning rate maximizes their wealth. Thus, market behavior at radically different learning rates has dubious relevance to our investigation.

4 Results

Table 1 shows the expected payoffs to the agent in the four situations **A–D**. These payoffs were calculated by averaging the agent’s final wealth in repeated simulations of each of the four situations. This decision matrix supports three conclusions.

First, note that the agent’s dominant strategy is to *include* technical trading rules, since the payoff in **A** exceeds that in **C** and the payoff in **B** exceeds that in **D**. No matter what strategy the other agents in the market might be using, it’s always advantageous for the agent to include technical trading in his market forecasting rules.

Second, recall that *each* agent in the market faces decision problem described in Table 1, which creates the multi-person simultaneous-move game we described above. Since including technical trading is each agent’s dominant strategy, the strategy of including technical trading is the one and only symmetric Nash equilibrium of the simultaneous-move game. The state in which everyone excludes technical trading is unstable. Imagine the market is temporarily in that state. Then, since the expected payoff in situation **B** exceeds that in situation **D**, it is in each agent’s advantage to switch to including technical trading. Rational

decision theory drives the market to the situation in which everyone includes technical trading.

Third, note that the expected payoff in situation **A** is less than the expected payoff in situation **D**. Thus, the expected aggregate wealth is less if everyone includes technical trading than if everyone excludes technical trading. In other words, everyone is better off if no one includes technical trading. When everyone follows the same strategy, it is socially optimal for no one to engage in technical trading. So, engaging in technical trading leads the market to a sub-optimal state. The market gets locked into a less desirable equilibrium.

Thus, technical trading creates a prisoner's dilemma problem in the market. Although it is to the social advantage if everyone foregoes technical trading, each individual has an incentive to cheat. In the aggregate, then, if everyone does what is rational for her, all will engage in technical trading and thus make themselves all worse off.

Figures 1–4 show time series data from typical simulations of each of the four situations in the decision matrix of our agents. The top of each figure compares how the accumulated wealth of the individual agent compares with that of the rest of the traders. The middle and bottom of each figure show the extent of technical trading in the market. Specifically, they represent those bits in the agents' forecasting rules that are set to non-null (i.e., non-# values) The percentage of those bits that are set to conditions recognized by fundamental rules are shown in the middle plot, and the percentage of those bits that are set to conditions recognized by technical rules are shown in the bottom plot.

These figures illustrate the market's behavior in the four situations. We see the significant advantage in accumulated wealth that technical trading creates in situations **B** and **C** (Figures 2 and 3), and we see illustrations of the different final wealth reported in Table 1. It is clear that agents take advantage of technical trading when they can. Note that 80% of those bits used in the agents' trading rules are technical bits in Figure 1, with similar levels of technical trading evident in those agents that include technical trading in Figures 2 and 3. Thus, it is precisely the occurrence of technical trading that explains the different expected payoffs in Table 1. When agents are given the opportunity to include technical trading in their market forecasts, they overwhelmingly do so.

5 Discussion

These results raise two important questions: (i) Why are agents led to an equilibrium in which everyone uses technical trading? (ii) Why is everyone worse off when everyone engages in technical trading?

We are attracted to the following answer to question (i). Assume that the price stream contains some definite trends. (In the present case, the price trends are due in part at least to the autoregressive form of the dividend stream; recall section 2.1 above. But the same argument applies no matter what causes the price trends.) Assume, further, that technical trading rules can detect these trends. Then if only a single agent discovers the technical trading rules, she

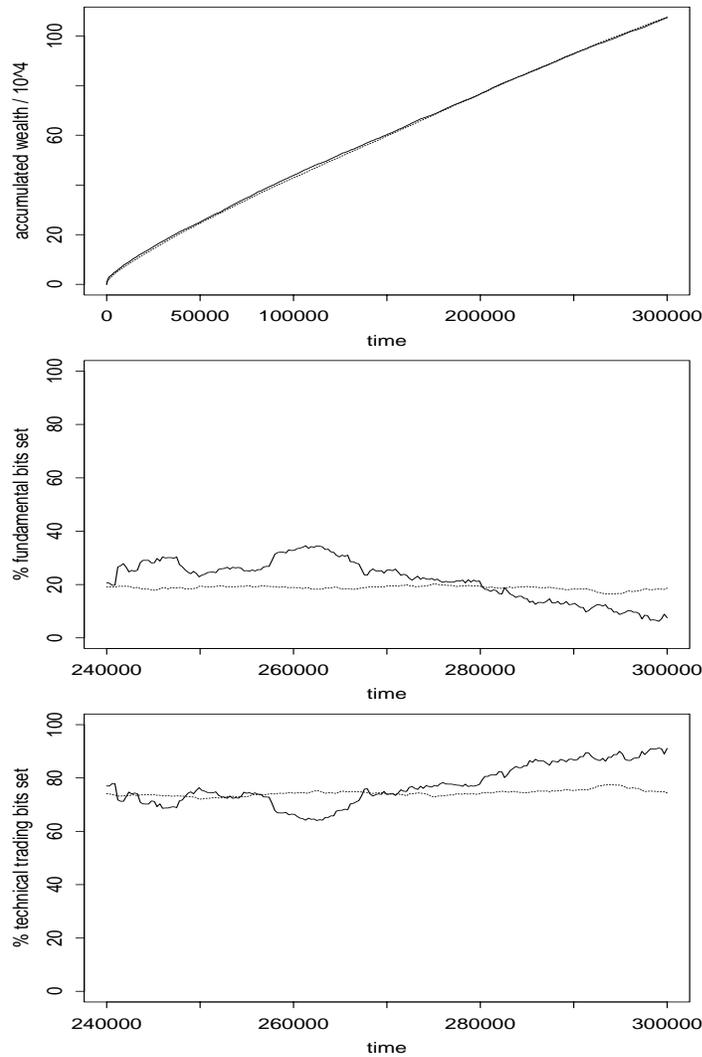


Figure 1: Time series data from a typical simulation of situation **A**, in which all agents *include* technical rules. The solid lines are data from the single agent and the dotted lines are data *averaged* from all other agents. (The accumulated wealth plot shows the entire duration of the simulation, but the plots of fundamental and technical bits set are blow ups of the last fifth of the simulation.) **Top:** The wealth of the agents. **Middle:** The percentage of the bits set in trading rules (of all agents in the market) that are *fundamental* bits in the final fifth of the run. **Bottom:** The percentage of bits set that are *technical* bits in the final fifth of the run. The number of technical and fundamental bits set reflects the number of technical and fundamental ‘market states’ an agent can recognize. Note that the number of fundamental and technical bits set for the single agent is close to the mean for the rest of the population. (The deviations from this mean are artifacts of the smoothing caused by averaging the data for all the other agents.)

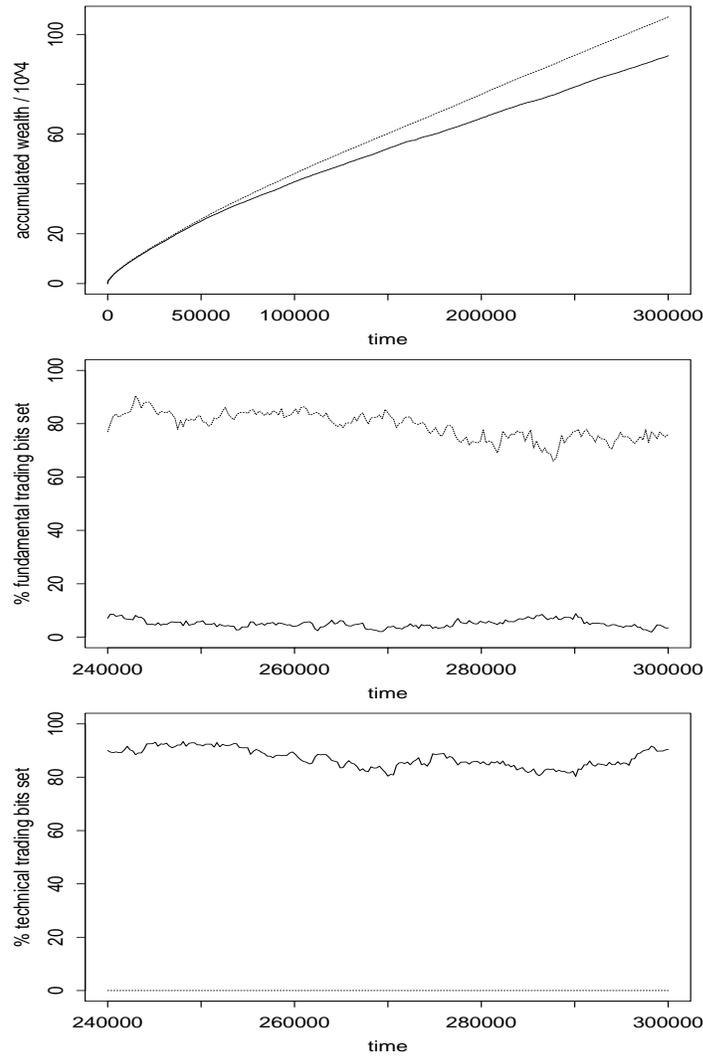


Figure 2: Time series data from a typical simulation of situation **B**, in which one agent *includes* technical rules while all others *exclude* them, analogous to Figure 1 (see caption above). Note that the singular agent using technical rules accumulates significantly more wealth than those agents using only fundamental rules almost all through the run, and that this difference grows over time.

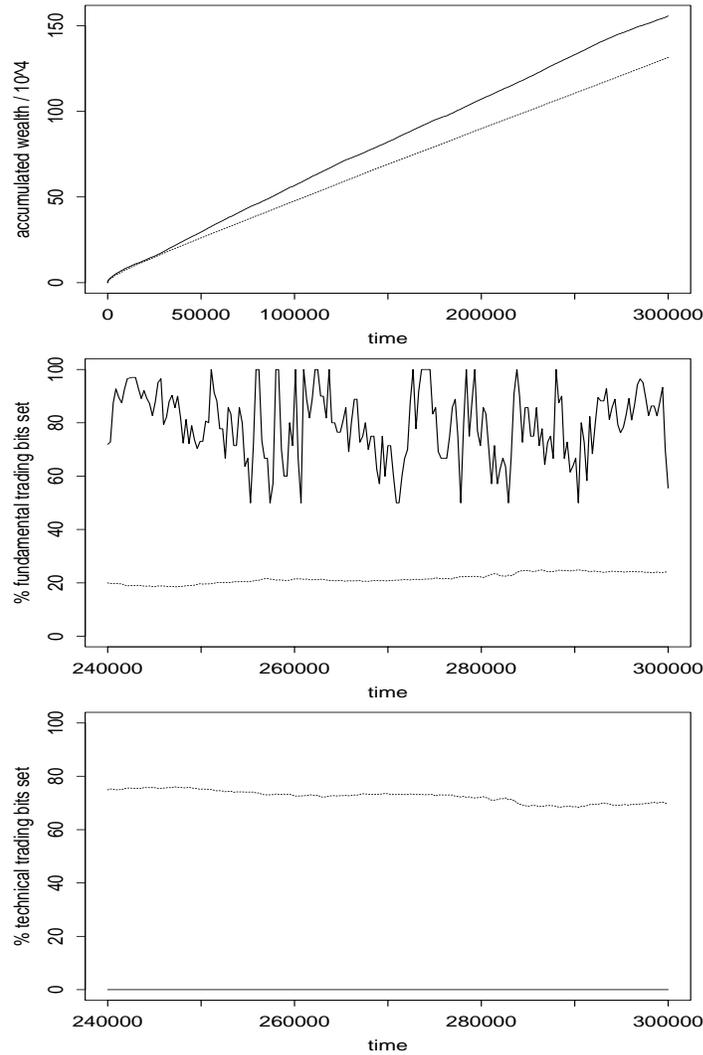


Figure 3: Time series data from a typical case simulation of situation **C**, in which one agent *excludes* technical rules while all others *include* them, analogous to Figure 1 (see caption above). Note that, since the singular agent has only fundamental rules, almost all of the bits set in her rules are fundamental bits. The higher variance of the percentage of bits set for the single are due to the fact that this data is not averaged.

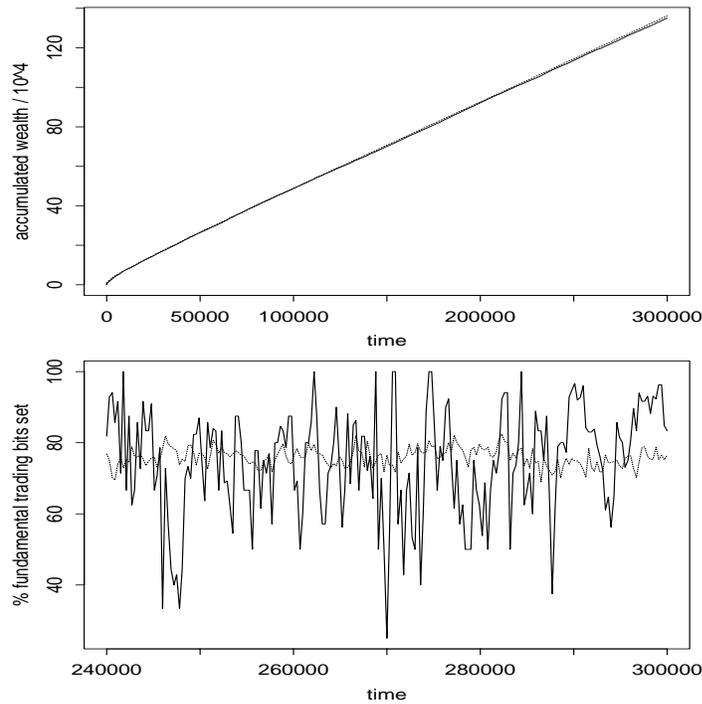


Figure 4: Time series data from a typical simulation of situation **D**, in which all agents *exclude* technical rules, analogous to Figure 1 (see caption above) except that technical bits are not shown since no agents can use them. Note that all agents accumulate equivalent wealth and have similarly structured rules. The higher variance of the percentage of fundamental bits set for the single are due to the fact that this data is not averaged. Deviations of the data for the single agent from the mean of the rest of the agent are entirely accidental and characteristic only of this run.

can exploit these trends without dissipating them and thus “beat the market,” earning huge profits. But now, as more agents begin to adopt technical rules, the incentives for technical trading can reinforce themselves in a new way. Detailed descriptions of the mechanisms for this are provided elsewhere [1, 2, 10, 11, 15, 22, 28]. 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. This self-reinforcement process can make technical trading rules more accurate than fundamental rules that generally predict that the price will revert to its true value. (Evidence for this positive-feedback in the Santa Fe stock market has been provided elsewhere [3].)

The mechanism behind this process in the Santa Fe Artificial Stock Market would be the genetic algorithm by which agents’ trading rules evolve. If technical trading rules become more successful, even if merely because they happen to be self-fulfilling prophecies, they will be likely to survive the culling process of the GA, and new rules introduced by the GA, their “offspring”, will also be technical trading rules.

This answer to question (i) implies an answer to question (ii). The self-fulfilling prophecies created by technical trading dramatically increase the volatility of prices in the market, causing bubbles and crashes [27, 3, 20, 21]. This increased noise in the market decreases the accuracy of the forecasting rules being used. The decreased accuracy of forecasting rules, in turn, drives down the agents’ wealth; 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, the use of a technical trading rule in the market poses a negative externality. It worsens everyone else’s strategies by driving prices away from the fundamental value and increasing noise. When all agents choose to perform high technical trading, they worsen each others strategies, there is a loss of efficiency and the average returns in the market are lowered.

These explanations fit well with the results of our experiments. In situation **A**, high technical trading by all agents lowers everyone’s wealth, presumably because everyone’s predictors are less accurate. In situation **B**, in which only one agent engages in technical trading, she accumulates significantly more wealth than the other agents, but since only one agent is cashing in on price patterns, everyone else’s forecasting rules are not rendered inaccurate, so the price patterns do not dissipate in noise. This lack of noise makes the single agent’s trend detectors stronger, which is reflected in her high final wealth (see Figure 2).

If one agent uses only fundamental rules but everyone else uses technical rules (situation **C**, Figure 3), the fundamental trader is worse off than the other agents. The market is so noisy that fundamental strategies have little value; technical traders are driving short-term price patterns so prices do not obey her fundamental predictions and she ends up worse off.

Situation **D** (Figure 4) is the best global state. All agents in this case rely

solely on fundamental rules. The absence of technical trading rules reduces the noise in the market, strengthening the accuracy of agent’s predictors, thus leading them to accumulate higher levels of wealth over time.

Statistics of the price stream in the Santa Fe Artificial Stock Market provide further support for these explanations. When all agents use fundamental trading strategies, agents show behavior that is consistent with the theory of rational expectations. When the price is over-valued, agents predict that the price will fall and thus drive the price down. Consequently, the volatility of prices is low and prices stay close to fundamental values. Trading still occurs because the market is constantly changing. But when agents include technical rules in their pool of forecasting rules, the market becomes unstable. Bubbles and crashes occur frequently. The volatility of prices roughly doubles and prices deviate from fundamental values for extended periods of time, having about a third the correlation compared to when only fundamental trading rules are used.

An alternate explanation of our results reported above is that, when only one agent exploits these patterns in the market, this agent beats the market (as we described above), but if all agents use technical trading rules, they dissipate the patterns, thereby making the market more efficient and allowing the agents to accumulate less wealth. However, we find it difficult to reconcile this explanation with the bubbles, crashes and positive-feedback observed in the market [3, 20].

We should reiterate that the observed advantage enjoyed by a singular technical trader is no surprise. The autoregressive dividend stream creates structure in the price stream that fundamental traders cannot detect, so a single technical trader can exploit this structure without destroying it. What *is* notable is that the wholesale adoption of technical trading worsens everyone’s earnings so much that a prisoner’s dilemma is created. Furthermore, the explanation for this result in no way depends what causes price patterns that technical trading exploits. Both real and artificial markets can have many kinds of patterns in prices, and in general these are not driven by external structure in dividends. No matter how these patterns arise, our results suggest that, while a single trader who discovers these patterns can profit significantly, if all traders discover the patterns they dissipate them by exploiting them, thus lowering profits for all.

6 Summary and Conclusion

Our simulations using the Santa Fe Artificial Stock Market suggest that financial markets can end up in a prisoner’s dilemma, creating a sub-optimal strategic equilibrium in which extensive technical trading creates market volatility and thus reduces earnings. We show that each agent will choose to include technical trading rules in his repertoire of forecasting rules, even if other traders use only rules based on stock-price fundamentals. Including technical rules is the dominant strategy of the game because it makes each agent better off regardless of what strategy other traders in the market follow.

Because this singular agent’s decision is mirrored by a decision for every other trader in the market, we have a multi-person game in which each agent

has a dominant strategy. The use of this dominant strategy by all agents in the population, however, drives the market to a symmetric Nash equilibrium at which the average final wealth of agents in the market is lower than in the hypothetical equilibrium in which everyone uses only fundamental trading rules. Our explanation of this reduced wealth is that the wide-spread use of technical trading rules worsens the accuracy of the predictions of all agents by reinforcing price trends, augmenting volatility, and making the market more noisy.

Though the model considered in this paper is an extreme simplification of real-world stock markets, we believe that it captures some essential elements of such markets. Moving away from assumption of rational expectations, with its implication that agents *know* the underlying structure of the stochastic processes driving the model, allows us to mimic the kind of asymmetric model uncertainty and learning that we observe in actual markets. Our market models the process of searching for the ideal forecasting rule explicitly through a mechanical, yet quite sophisticated, learning process. Our analysis leads to an equilibrium outcome of this process—a volatile market in which the use of technical trading rules is pervasive—that mirrors some key aspects of real markets that are contrary to the predictions of some of the most widely accepted models of stock markets.

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. However, the results obtained in our early explorations point to a conclusion of great potential importance: that technical trading might be inevitable, yet traders would end up better off if it were possible to prevent it.

Acknowledgments. 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, thanks to Doyme Farmer, Blake LeBaron, John Moody, Norman Packard, Richard Palmer, and Robert Seymour.

References

- [1] Arthur, W.B. 1988. Self-reinforcing mechanisms in economics. In: *The Economy as an Evolving, Complex System I: Proceedings of the Evolutionary Paths of the Global Economy Workshp.* P.W. Anderson, K. J. Arrow and D. Pines (eds.). Redwood City, CA: Addison-Wesley.
- [2] Arthur, W.B. 1989. Positive feedbacks in the economy. *Scientific American* (February): 92–99.
- [3] 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: *The Economy as an Evolving, Complex System II.* W. B. Arthur, D. Lane, and S.N. Durlauf (eds.). Menlo Park: Addison-Wesley. Also published as Santa Fe Institute Paper 96-12-093.

- [4] Beltratti, A., and S. Margarita. 1992. Simulating and artificial adaptive stock market. Mimeo. Turin University.
- [5] Bierman, S.H., and L. Fernandez. 1993. *Game Theory with Economic Applications*. Redwood City, CA: Addison-Wesley.
- [6] Bray, M. 1981. Futures trading, rational expectations, and efficient markets hypothesis. *Econometrica*, **49**: 575–96.
- [7] 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.
- [8] Cootner, P. 1967. *The random character of stock market prices*. Cambridge, MA: MIT Press.
- [9] DeLong, J.B., and A. Schleifer, L.H. Summers and J. Waldmann. 1990. Noise trader risk in financial markets. *Journal of Political Economy*, **98**: 703–738.
- [10] DeLong, J.B., A. Schleifer, L.H. Summers, and J. Waldmann. 1990. Positive feedback and destabilizing rational speculation. *Journal of Finance*, **45**: 379–395.
- [11] 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.
- [12] Diba T.B., and G. Grossman. 1988. The theory of rational bubbles in stock prices. *The Economic Journal*, **98**: 746–754.
- [13] Figlewski, S. 1978. Market efficiency in a market with heterogeneous information. *Journal of Political Economy*, **86**(4): 581–97.
- [14] Fama E.F. 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, **25**: 383–420.
- [15] Farmer, R.A. 1993. *The Macroeconomics of Self-Fulfilling Prophecies*. Cambridge, MA: MIT Press.
- [16] Frankel, J.A., and K.A. Froot. 1990. Chartists, fundamentalists, and trading in the foreign exchange market. *AEA Papers and Proceedings*, **80**: 181–185.
- [17] Grossman, S. 1978. Further results on the informational efficiency of competitive stock market. *Journal of Economic Theory*, **18**: 81–101.
- [18] Grossman, S., and J. Stiglitz. 1980. The impossibility of informationally efficient markets. *American Economic Review*, **70**: 393–68.
- [19] Grossman, S., and R. Shiller. 1981. The determinants of the variability of stock prices. *American Economic Review*, **71**: 222–7.
- [20] Joshi, S., and M.A. Bedau. 1998. An explanation of generic behavior in an evolving financial market. To appear in: *Proceedings of Complex Systems '98: Complexity Between the Ecos—from Ecology to Economics*, November 30–December 3, 1998, University of New South Wales, Sydney, Australia.
- [21] Joshi, S., J. Parker, and M.A. Bedau. In progress.

- [22] Kirman, A., 1991. Epidemics of opinion and speculative bubbles in financial markets. In: *Money and Financial Markets*, M. Taylor (ed.). London: Macmillan.
- [23] Kurtz, M. 1994. On the structure and diversity of rational beliefs. *Economic Theory*, **4**: 877–900.
- [24] LeBaron, B., W.B. Arthur, and R. Palmer. 1998. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, forthcoming.
- [25] LeBaron, B. 1997. Technical trading rules and regime shifts in foreign exchange markets. In: *Advanced Trading Rules*, E. Acar and S. Satchell (eds.). Woburn, MA: Butterworth-Heinemann.
- [26] Marengo, L., and H. Tordjman. 1995. Speculation, heterogeneity, and learning: A model of exchange rate dynamics. Working Paper, WP-95-17. IIASA.
- [27] 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.
- [28] Rieck, C. 1994. Evolutionary simulation of asset trading strategies. In: *Many Agent Simulation and Artificial Life*, E. Hillenbrand and J. Stender (eds.). Washington DC: IOS Press.
- [29] Shleifer A. and L.H. Summers. 1990. The noise trader approach to finance. *Journal of Economic Perspectives*, **4**: 19–33.
- [30] Schwager, J.D. 1995. *Technical Analysis*. New York: John Wiley and Sons.
- [31] Satchell, S. 1997. *Advanced Trading Rules*. Woburn, MA: Butterworth-Heinemann.
- [32] Shiller, R. 1989. *Market Volatility*. Cambridge, MA: MIT Press.
- [33] Soros, G. 1994. *The Theory of Reflexivity*. New York: Soros Fund Management.
- [34] Summers, L.H. 1986. Does the market rationally reflect fundamental values? *Journal of Finance*, **46**: 591–601.