A Prisoner's Dilemma Causes Technical Trading

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

[†]Reed College

[‡]Santa Fe Institute

Abstract

We examine the use and profitability of technical trading rules in financial markets by studying the Santa Fe Stock Market, an agent-based model of a stock market, in which traders make investment decisions by forecasting stock prices using technical and fundamental rules. We show that individual traders earn more by using technical rules, no matter what other traders do, so the use of technical trading rules is everyone's dominant strategy. We also show that agents would collectively earn more if nobody used technical trading rules than if everyone used them. Thus, technical trading becomes widespread due to a multi-person prisoner's dilemma. This prisoner's dilemma arises because technical trading generates positive-feedback, destabilizes prices and makes everyone's market forecasts less accurate over time.

Corresponding author:

Shareen Joshi

Santa Fe Institute

1399 Hyde Park Rd.

Santa Fe NM 87501

voice: (505) 984-8800, ext. 249

fax: (505) 982-0565

1 Introduction

The widespread use and proven profitability of technical trading rules in financial markets has long been a puzzle in academic finance (e.g., Campbell, Lo and MacKinlay 1997). While the usefulness of fundamental trading rules can be explained by the standard theory of efficient markets (Samuelson 1965, Cootner 1967, Malkiel 1992), this is not true for technical trading rules. Nevertheless, there is ample evidence for the use and profitability of these rules (Shiller 1989, Frankel and Froot 1990, Kirman 1991, Brock, Lakonishok and Lebaron 1992, Soros 1994, Reick 1994, Werner and Thaler 1995, Keim and Madhaven 1995, Schwager 1995, Acar and Satchell 1997).

In recent years, a wide variety of theoretical, empirical, computational and experimental models have focussed on technical trading rules, explaining their widespread use and their effects on asset prices. A central theme in this literature is that technical trading causes positive feedback, which in turn causes asset prices to deviate from their fundamental values (Arthur 1988, Arthur 1989, Shiller 1989, Delong, Schleifer, Summers and Waldman 1990a, Schleifer and Summers 1990, Delong, Schleifer, Kirman 1991, Summers and Waldman 1991, Reick 1994, Soros 1994, Farmer 1998, Youssefmir, Huberman and Togg 1998).[1] It has been argued that the presence of technical traders induces higher price volatility (Shiller 1989, Delong et al. 1990a, Kirman 1991, Arthur et al. 1997), and that technical traders may earn more wealth than the more 'rational' fundamental traders (Delong et al. 1991, Soros 1994, Reick 1994, Farmer 1998).

In this paper we use simple game theory and an agent-based model of a financial market to argue that technical trading can both cause and result from a prisoner's dilemma. The key characteristic of the agent-based model we use—which is the Santa Fe Artificial Stock Market (Palmer, Arthur, Holland, Lebaron and Tayler 1994, Arthur, Holland, Lebaron, Palmer and Tayler 1997, Lebaron, Arthur and Palmer 1998)—is that traders choose their forecasting rules from an evolving set of such rules, depending on which ones have proved to be the most successful predictors of recent stock-price changes. We first show that, regardless of whether other traders are using fundamental or technical trading rules, an individual agent in the market always gains by adding technical trading rules to her repertoire of forecasting techniques, so technical trading is a trader's dominant strategy. Second, we show that the use of this strategy by all agents in the market drives the market to a symmetric Nash equilibrium at which all traders earn less than

[1] There is a significant literature on trend following and destabilizing speculation, dating back to at least Bagehot (1872). Reviews of this literature may be found in Delong et al. (1990a), Campbell et al. (1997) and Shiller (1989).

they would than in a hypothetical equilibrium where all agents use only fundamental rules. So, once the use of technical rules becomes an option, the relative payoffs of using technical rules compared with not using these rules creates a decision matrix with payoffs just like those in a prisoner's dilemma; i.e., the possibility of technical trading creates a situation formally equivalent to a prisoner's dilemma, in which the classic prisoner's dilemma choice between cooperating and defecting is replaced with the choice between using technical trading rules or not doing so. Just as in the classic prisoner's dilemma it is rational to defect, in our prisoner's dilemma analogue it is rational to engage in technical trading. For simplicity of exposition, in this paper we will say that any decision problem with a payoff structure just like a prisoner's dilemma *is* a prisoner's dilemma. It is in this sense that technical trading both causes and results from a prisoner's dilemma.

The mechanism behind this prisoner's dilemma is that widespread technical trading increases the variability of prices, making the market more noisy. This makes it more difficult for any trader to predict stock-price movements, and thus lowers the wealth earned by all. This explanation generally meshes with the recent work on technical trading cited above. The primary novelty of our work is to show that technical trading causes and results from a prisoners' dilemma.

Section 2 below describes the Santa Fe Artificial Stock Market model and explains how we use it to study technical trading, section 3 presents and explains the results of our experiments, and section 4 concludes by explaining the relevance of these results to financial markets in general.

2 A Method for Studying Technical Trading

Our goal is to study technical trading without assuming that all traders are perfectly rational and have homogeneous expectations. An obvious context for such a study is an agent-based model of a financial market (Holland and Miller 1991, Sargent 1993). In this section, we describe the agent-based market model we use and the framework we use to investigate the causes and effects of technical trading in this model.

2.1 The Santa Fe Artificial Stock Market

We study the emergence of technical trading in the Santa Fe Artificial Stock Market, which was developed by Brian Arthur, John Holland, Blake LeBaron, Richard Palmer, and Paul Taylor (Palmer et al. 1994, Arthur et al. 1997, Lebaron et al. 1998). This section briefly describes this model. More detailed descriptions are available elsewhere (Palmer et al. 1994, Arthur et al. 1997, Lebaron et al. 1998). 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 of a financial market in which

agents continually explore and develop market forecasting rules, buy and sell assets based on the predictions of their best performing rules, and revise or discard these rules based on their past performance. Each agent acts independently, but the returns to each agent depend on the decisions made simultaneously by all the other agents in the market.

The market contains a fixed number N [25] of agents each of whom is endowed with an initial sum [10000] of money (in arbitrary units). Time is discrete. At a given time period each agent decides how to invest between a risky stock and 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.

At each time step each agent 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. The rule used at each time is chosen from the agent's set of [100] 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:

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

The values of the variables a and b in an agent'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 agents discard ineffective forecasting rules and try out new forecasting rules, so the values of a and b in an agent's set of rules evolves, as described in the next section.

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 0s and 1s match more narrow and
specific market states. As the agents' sets of forecasting rules evolve, the number of 0s and 1s 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 agents is the number of 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.

The market conditions defining market states fall into two main categories: *technical* condi-
tions 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 agents 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 only current prices 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).

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 technical rules. But in our model the
agents do not know the parameters of the dividend process, so to improve their forecasts they
must experiment with alternative fundamental and, perhaps, technical rules.

A simplified example might help clarify the structure of market forecasting rules. Suppose
that there is a three-bit market descriptor.[2] The first bit corresponds to the fundamental market
condition in which the price is 75% higher than its fundamental value, the second bit corresponds

---

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

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.

Each rule is assigned a measure of *accuracy*, where the accuracy is defined as the moving-average of the variance of the error (the difference between the forecasted price and the true price). An accuracy updating parameter [100] controls the length of time over which the moving-average is calculated.

If the market state in a given period matches the descriptor of a forecasting rule, the rule is said to be *activated*. A number of an agent's forecasting rules may be activated at a single time, thus giving the agent many possible forecasts to choose among. An agent decides which of the active forecasts to use by choosing at random among the active forecasts, with a probability proportional to the rule's accuracy. Once the agent has chosen a specific rule to use, the rule's a and b values determine the agent's investment decision at that time.

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

$$U(W_{i,t+1}) = -\exp(-\lambda W_{i,t+1}) \tag{2}$$

where $W_{i,t+1}$ is the wealth of agent i at time $t+1$, and $0 < \lambda[0.5] \leq 1000$. 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:

$$W_{i,t+1} = x_{i,t}(p_{t+1} + d_{t+1}) + (1 + r)(W_{i,t} - p_t x_{i,t}) \tag{3}$$

where $x_{i,t}$ is agent i's demand for the stock at time t. If we assume 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 $\sigma^2_{i,t,p+d}$, and if we assume that the distribution of forecasts is normal, then, as Arthur et al. (1997) explain, agent i's demand for the stock at time t should be:

$$x_{i,t} = \frac{E_{i,t}(p_{t+1} + d_{t+1}) - p_t(1 + r)}{\lambda \sigma^2_{i,t,p+d}} \tag{4}$$

6

where p_t is the price at time period t and λ is the relative degree of risk aversion. The bids and offers submitted by agents need not be integers; the stock is perfectly divisible. The aggregate demand for the stock must equal the number of shares in the market.

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, repeats this process. When the market clears, the 'trial price' becomes the current period's market price.

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 its accuracy 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 and crossover 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 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 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.

2.2 Experimental Methods

In this paper, we study one particular aspect of an agent's strategy for trading in the market: whether technical rules should be included in her collection of market 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.

We investigate whether it is advantageous for an agent in our model market to include technical trading rules; that is, we investigate what happens when a single agent must choose between the two strategies explained above. We investigate this in two steps. First, we consider what happens

7

when 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. Second, we consider a more general situation and ask what strategy the agent should choose when she assumes that the population of other traders might be following some mixed strategy, i.e., the percentage of other traders including technical trading varies somewhere between 0% and 100%.

To make a rational decision in the 2×2 decision problem, 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 and all other traders *include* technical rules.

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 and all other traders *exclude* technical rules.

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 (Bierman et al., 1993). 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 each of 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 sets of simulations.

---

[3] The 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).

[4] A dominant strategy is defined as one that outperforms all other strategies *regardless* of the strategies being used by other agents (Bierman and Fernandez, 1993).

We next considered a more general problem and checked for the existence of a mixed-strategy equilibrium by varying the percentage of agents with access to technical trading rules. Then we observed the difference in the wealth earned by technical traders compared to the mean wealth earned in the population of all traders in the market. Once again, we simulated each of these market condition situations 45 times (using the same 45 random number sequences as before), and except for the percentage of agents with access to technical trading rules, all model parameters were held constant, at the same values as before.

Previous work (Palmer et al. 1994, Arthur et al. 1997, Lebaron et al. 1998, Joshi and Bedau 1998) has shown that the evolutionary learning rate is a crucial parameter controlling the behavior of this model. All our simulations here were carried out 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 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 (Joshi, Parker and Bedau 1999) has shown 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.

3   Causes and Effects of Technical Trading

Table 1 shows the expected payoffs to an 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. Figure 1 shows how much this payoff exceeds the mean payoff the subpopulation of traders using only fundamental rules, as the percentage of technical traders varies between 0% and 100%. (Be aware that, as the proportion of traders engaging in technical trading varies, so does the population mean payoff of the subpopulation of traders using only fundamental rules.)

These data support three conclusions. First, note that the agent's dominant strategy is to *include* technical trading rules. In the 2×2 decision, the payoff in **A** exceeds that in **C** and the payoff in **B** exceeds that in **D**. More generally, if the agent does not assume that other traders all follow the same strategy but makes her decision in a context in which the market can include an arbitrary mixture of technical and non-technical traders, Figure 1 shows that an agent always improves her position by adopting technical trading rules. No matter what strategy the other agents in the market are using, it is always advantageous for the agent to include technical trading rules to her repertoire of market forecasting rules. Engaging in technical trading is the agent's dominant strategy.

Second, recall that *each* agent in the market faces the decision problem analyzed in Table 1 and Figure 1, because of the multi-person simultaneous-move game we described above. Since the inclusion of technical trading rules is a single agent's dominant strategy, the only symmetric Nash equilibrium of the simultaneous-move game occurs when all agents include technical trading rules. The state in which some fraction of traders exclude technical trading is unstable. Imagine the market is temporarily in such a state. Then, since the expected payoff in situation **B** exceeds that in situation **D** and, more generally, since including technical trading rules always improves an agent's payoff when some percentage of traders eschew these rules (Figure 1), it is advantageous for those not including technical trading rules to change and start using them in addition to their fundamental rules. So 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 rules than if everyone excludes these rules. In other words, everyone is better off if no one includes technical trading rules. 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, the option of engaging in technical trading creates a prisoner's dilemma in the market, and this same prisoner's dilemma causes all rational traders to include technical trading rules to their repertoire of forecasting rules. Although it is to the social advantage if everyone foregoes the use of such rules, each individual has an incentive to cheat. In the aggregate, then, if everyone does what is rational for her, all will engage in the use of technical trading rules and thus make themselves all worse off.

Figures 2–5 show time series data from typical simulations of each of the four situations in the 2×2 decision matrix of our agents. The top of each figure compares the accumulated wealth of the individual agent 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, with fundamental bits shown in the middle and technical bits shown at the bottom.

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 3 and 4), 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 2, with similar levels of technical trading evident in those agents that include technical trading in Figures 3 and 4. 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 take advantage of technical trading in their market forecasts, they overwhelmingly do so.

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). 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 argument we give here applies no matter what causes the price trends.) Assume that technical trading rules can detect these trends. If only a single agent exploits the patterns with technical trading rules, she 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 (Arthur 1988, Arthur 1989, Delong et al. 1990a, Delong et al. 1991, Kirman 1991, Farmer 1993, Youssefmir 1998). 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. In the short run, this self-reinforcing process can make technical trading rules more accurate than fundamental rules which presume that the price will revert to its true value. Arthur et al. (1997) provide evidence for this positive-feedback in the Santa Fe Artificial Stock Market.

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 increase the variability of prices in the market, causing bubbles and crashes. This aspect of the Santa Fe Artificial Stock Market has been detailed elsewhere (Palmer et al. 1994, Arthur et al. 1997, Joshi and Bedau 1998, and Joshi, Bedau and Parker 1999). 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 because less accurate rules tend to be less profitable.[5] 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.

---

[5] This is true of our model, but not be true in real world markets. Noise trading has been shown to be profitable in several situations (Delong et al. 1990b, Delong et. al. 1991).

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 other's 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 3).

If one agent uses only fundamental rules but everyone else uses technical rules (situation **C**, Figure 4), 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 the single agent's fundamental predictions and she ends up worse off.

Situation **D** (Figure 5) 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 variability 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 variability 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 (Arthur et al. 1997, Joshi and Bedau 1998).

We should note that the advantage enjoyed by a singular technical trader in this artificial stock market 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 on what causes the 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.

4   Conclusions

Our simulations using the Santa Fe Artificial Stock Market suggest that technical trading in financial markets both causes and results from a prisoner's dilemma. The use of technical trading rules in addition to fundamental rules always increases a trader's earnings, no matter what trading strategies are being used by other traders in the market, so rational traders will all be technical traders. But when all traders adopt this dominant strategy, the market is driven to a symmetric Nash equilibrium in which everyone earns less than they would in the hypothetical situation in which everyone eschewed technical trading. Thus, the possibility of technical trading creates a prisoner's dilemma. At the same time, the rational behavior in this prisoner's dilemma is to add technical trading rules to their collection of forecasting rules. So, this prisoner's dilemma causes the existence of technical trading. Our explanation of this reduced wealth is that widespread technical trading induces positive feedback which reinforces price trends and makes the market more noisy. This worsens the accuracy of every trader's predictions and thereby creates a sub-optimal strategic equilibrium.

Though the model considered in this paper is an extreme simplification of real-world stock markets, we believe that it captures enough of their essential elements that our results may well hold in many other markets, both artificial and real. By moving away from assumption of rational expectations, with its implication that agents know the underlying structure of the stochastic processes driving the model, we can mimic the kind of asymmetric uncertainty and learning observed in actual markets. Our model captures the search for ideal forecasting rule through a mechanical yet sophisticated learning process, and the model's equilibrium behavior— a noisy market with pervasive technical trading—mirrors some key aspects of real markets that contradict the predictions of widely accepted traditional theories.

More research needs to be done to exhaustively confirm the robustness of these results to variations in the model's parameters and structure and to check for similar effects in other models. Still, these initial results already point to a potentially important general conclusion about the

causes and effects of technical trading: that technical trading is inevitable even though traders would benefit if it could be prevented.

References

[1] Acar, E. and S. Satchell, 1997, Advanced trading rules (Butterworth-Heinemann, Woburn).

[2] 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.

[3] Arthur, W.B., 1989, Positive feedbacks in the economy, Scientific American 262 (February), 92–99.

[4] 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 (Addison-Wesley, Menlo Park) 15–44.

[5] Bagehot, W., 1872, Lombard Street (Smith Elder, London).

[6] Beltratti, A., and S. Margarita, 1992, Simulating and artificial adaptive stock market, mimeo, Turin University.

[7] Bierman, S.H., and L. Fernandez, 1993, Game theory with economic applications (Addison-Wesley, Redwood City).

[8] Bray, M., 1981, Futures trading, rational expectations, and the efficient markets hypothesis, Econometrica 49, 575–596.

[9] 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.

[10] Campbell, J. Y., A. Lo., and A.C. MacKinlay, 1997, The econometrics of financial markets (Princeton University Press, Princeton).

[11] Cootner, P., 1967, The random character of stock market prices (MIT Press, Cambridge).

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

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

[14] 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.

[15] Diba, T.B., and G. Grossman, 1988, The theory of rational bubbles in stock prices, The Economic Journal 98, 746–754.

[16] Figlewski, S., 1978, Market efficiency in a market with heterogeneous information, Journal of Political Economy 86, 581–597.

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

[18] Farmer, R.A., 1993, The macroeconomics of self-fulfilling prophecies (MIT Press, Cambridge).

[19] Frankel, J.A., and K.A. Froot, 1990, Chartists, fundamentalists, and trading in the foreign exchange market, American Economics Association Papers and Proceedings 80, 181–185.

[20] Grossman, S., 1978, Further results on the informational efficiency of competitive stock market, Journal of Economic Theory 18, 81–101.

[21] Grossman, S., and J. Stiglitz, 1980, The impossibility of informationally efficient markets, American Economic Review 70, 393–398.

[22] Grossman, S., and R. Shiller, 1981, The determinants of the variability of stock prices, American Economic Review 71, 222–227.

[23] Holland, J.H, and J.H. Miller, 1991, Artificial adaptive agents in economic theory, American Economic Review 81, 365–371.

[24] 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.

[25] Joshi, S., J. Parker, and M.A. Bedau, 1999, Financial markets exist at sub-optimal equilibria, in preparation.

[26] 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.

[27] Kirman, A., 1991, Epidemics of opinion and speculative bubbles in financial markets, in: M. Taylor, ed., Money and financial markets (Macmillan, London) 355-368.

[28] Kurtz, M., 1994, On the structure and diversity of rational beliefs, Economic Theory 4, 877–900.

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

[30] LeBaron, B., 1997, Technical trading rules and regime shifts in foreign exchange markets, in: E. Acar and S. Satchell, eds., Advanced trading rules, (Butterworth-Heinemann, Woburn) 5–40.

[31] Malkiel, B., Efficient Market Hypothesis, 1992, in: P. Newman, M. Milgate, and J. Eatwell, eds., New palgrave dictionary of money and finance (Macmillan Press Limited, London).

[32] 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.

[33] Rieck, C., 1994, Evolutionary simulation of asset trading strategies, in: E. Hillenbrand and J. Stender, eds., Many agent simulation and artificial life (IOS Press, Washington DC) 112–136.

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

[35] Sargent, T.J., 1993, Bounded rationality and macroeconomics (Clarendon Press, Oxford).

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

[37] Schwager, J.D., 1995, Technical analysis (John Wiley and Sons, New York).

[38] Shiller, R., 1989, Market volatility (MIT Press, Cambridge).

[39] Soros, G., 1994, The theory of reflexivity (Soros Fund Management, New York).

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

[41] De Bondt, W., and R. H. Thaler, 1995, Financial decision making in markers and firms: A behavioral perspective, in: R.A. Jarrow, V. Maksimovic, and W.T Ziemba, Handbooks in Operations Research and Management Science, Vol. 9 (North-Holland, Amsterdam).

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

Caption for Figure 1. Additional wealth earned by traders using technical rules as a function of the percentage of technical traders in the market. (We show observed wealth divided by 10^4, for better readability.) The technical traders' additional wealth is computed by subtracting the mean wealth earned in the subpopulation of agents using only fundamental rules from the mean wealth earned in the subpopulation of agents including technical rules. Error bounds are calculated using standard deviations of each set of 45 simulations. Note that, no matter what percentage of traders in the market use technical trading, the technical traders always earn more than those using only fundamental trading rules.

Caption for Figure 2. 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 (note the smoothing caused by this averaging). (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 accululated by 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 conditions that 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.)

Caption for Figure 3. 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 2 (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.

Caption for Figure 4. 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 2 (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.

Caption for Figure 5. Time series data from a typical simulation of situation **D**, in which all agents *exclude* technical rules, analogous to Figure 2 (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; in other runs of the model, the single agent shows different accidental deviation from the mean.

| | 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 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, for better readability), derived by averaging the results of 45 simulations of each situation. Errors bounds are calculated using standard deviations of each set of 45 simulations.
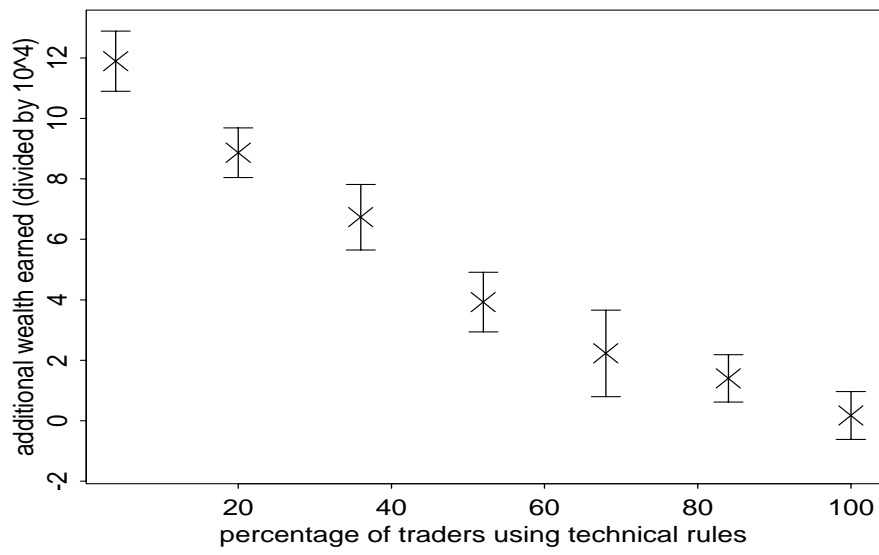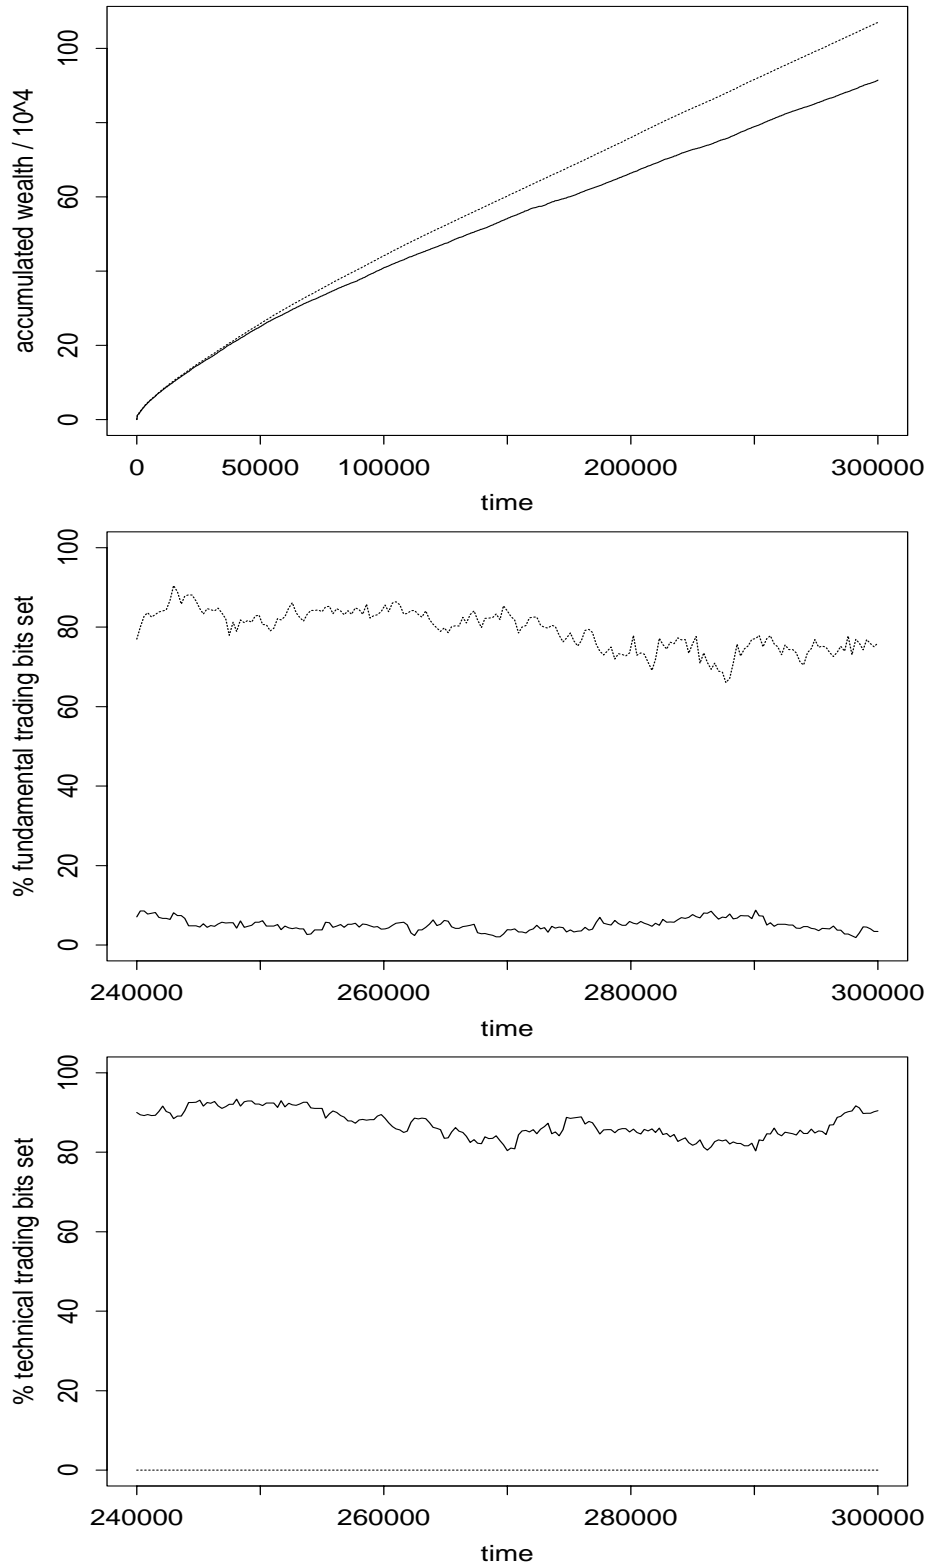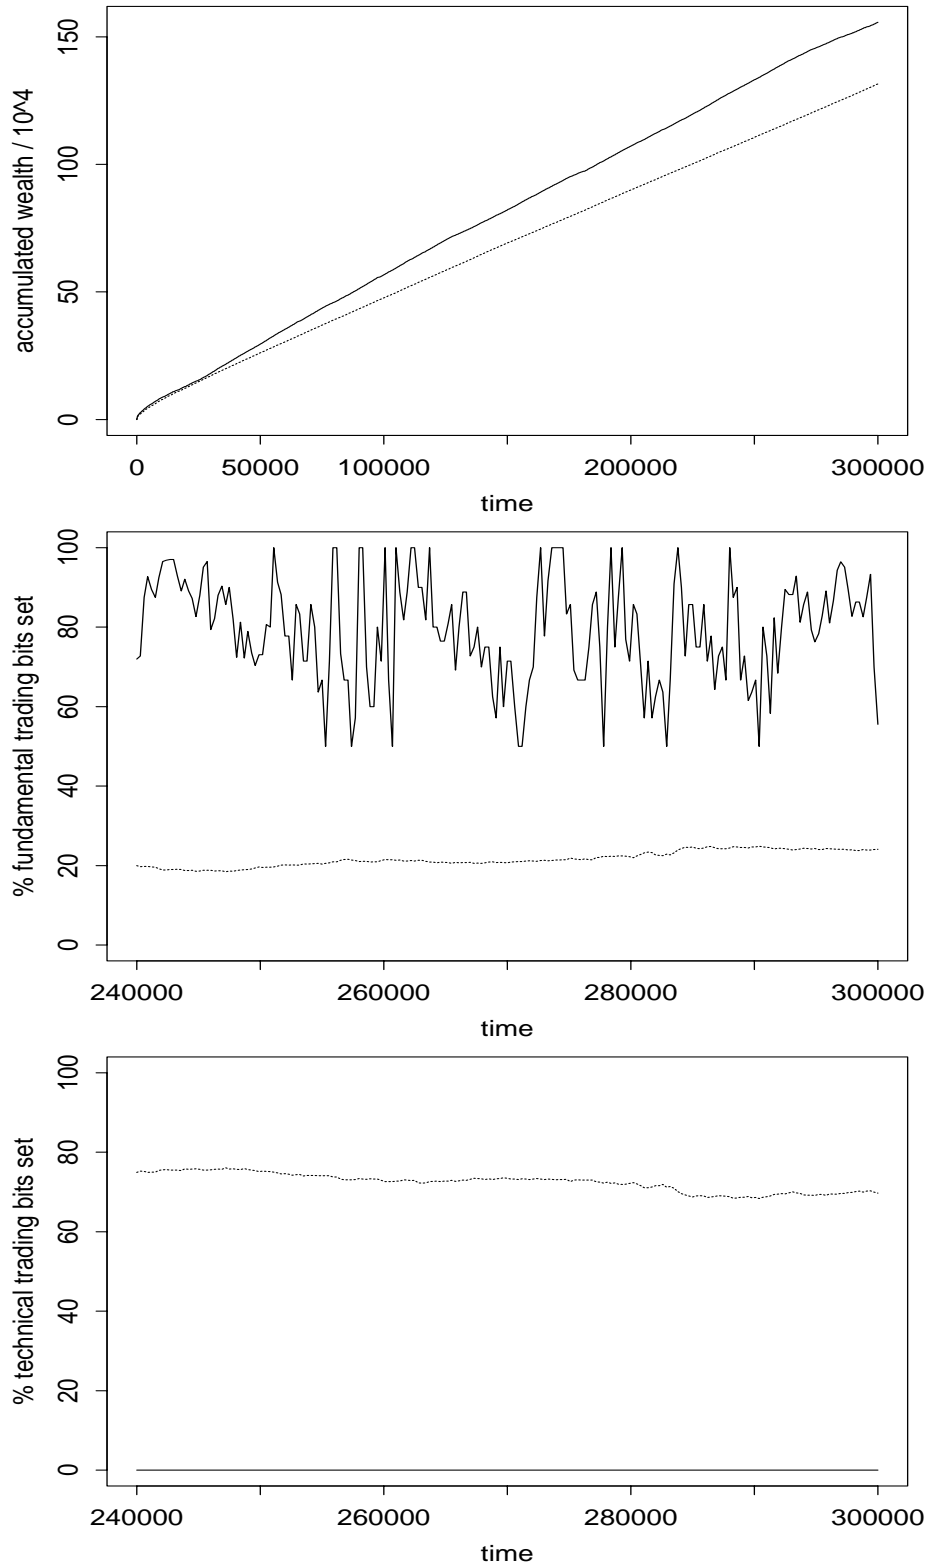
Figure 1:

Figure 2:

Figure 3:

Figure 4:

Figure 5: