The Efficiency of an Artificial Stock Market with Heterogeneous Intelligent Agents

Jing Yang*

(First Draft)

February, 1999

* Economics Department, Concordia University. Please address all correspondence to Ms. J. Yang, 1160 St-Mathieu, Suite 1703, Montreal, Quebec, CA H3H 2P4 Tel: (514)938-8276 Fax: (514) 938-4511
E-mail : jingy@vax2.concordia.ca

The Efficiency of an Artificial Stock Market
with Heterogeneous Intelligent Agents

Jing Yang

Abstract

In this paper, we construct an artificial equity market using an Artificial Neural Network (ANN). Based on the heterogeneous beliefs and trading strategies prevailing in actual equity markets, we divide traders into three groups: value traders, momentum traders and noise traders. Characteristically, value traders are endowed with an ANN learning mechanism that allows them to forecast the dividend growth rate. A double-auction market setting is adopted as our market mechanism for simulating the market structure.

Our artificial market is able to replicate several important features of real markets, including excess volatility, volatility persistence, and serial autocorrelation of stock returns. The profitability of three trading strategies are compared. On average, value traders exhibit the highest Sharpe ratio and significantly positive excess returns. The rational expectations equilibrium emerges for only one of the three experiments.

I am grateful to my thesis supervisors Prof. Bryan Campbell and Prof. Lawrence Kryzanowski for their insightful advice and valuable supports. I would also like to thank Prof. Anastas Anastapoulos and Prof. Michael Sampson for their useful comments.

The Efficiency of an Artificial Stock Market with Heterogeneous Intelligent Agents

1. Introduction

A computational agent-based model of financial markets stresses interactions and dynamics among a very diverse set of traders. The growing body of research in this area relies heavily on computational tools which bypass the restrictions of an analytical method. Using the long history of building rational expectation models with heterogeneous agents,[1] this paper addresses a problem involving complex heterogeneity which goes beyond the boundary of what can be handled analytically.

Experiments involving simulated markets frequently appear in the literature. For example, Gode & Sunder (1993) run a computer experiment where agents have no learning abilities. They find that the budget constraint is critical to reaching higher allocational efficiency. Other studies include Arifovic (1996), Lettau (1997), Margarita & Beltratti (1993), Marengo & Tordjman (1995), Rieck(1994), and Steiglitz, Honig & Cohen (1996). LeBaron (forthcoming 1998a) provides a very useful survey of some portions of this literature.

Inspired by the existing literature, we make three major extensions in this paper.

- A double auction is employed as the market mechanism. Many stock, commodity, currency, and others markets are organized according to a double-auction setting, moreover, some experimental double-auction markets with human traders are known

[1] Some of the origins of this research are reviewed in Grossman (1976) and Grossman & Stiglitz (1980). A good example of this approach is Bray(1982). This paper is part of the large body of literature on bounded rationality which was introduced by Simth, and was recently summarized in Conlisk (1996) and Sargent (1993).

to yield data which could approximate the equilibrium predictions of economic theory in a variety of environments.[2]

- Second, we introduce an ANN mechanism into the market simulation. In our model, value- (or fundamental-) traders are endowed with Elman Recurrent Networks to forecast equity dividend growth rates, and hence, the stock price. We also introduce noise traders into the stock market. Noise traders do not act consistently and strategically, but rather randomly post market orders at the bid or ask. More importantly, their behavior increases market volatility, and drives price down to make others hold risky assets.[3]

- Third, instead of focusing our interest simply on price equilibrium convergence alone, we also examine stock market dynamics. Our motivation comes from the possibility that situations may exist in which the market never really converges to what we call equilibrium. Therefore, we analyze the time series of prices and returns from this market to see whether the market exhibits a broad range of characteristics cited in the empirical literature, including excess volatility, volatility persistence, and market inefficiency. We also examine the portfolio performance of each group of traders in the market.

The adoption of a multi-agent model for financial markets is driven by a series of empirical puzzles which remain to be explained within the traditional structure of a representative agent. These include stock return predictability, volatility persistence and high excess return.[4] The series that are generated by our experiments, surprisingly, reproduce and hence confirm most of these puzzles. However, each puzzle occurs under different circumstances. In our Experiment I, where value traders are the only type of market participant, price convergence occurs and trading volume approaches zero after

---

[2] Gode and Sunder (1993) show that a double auction can sustain high levels of allocative efficiency even if agents do not maximize or seek profits.
[3] See De Long, Shleifer & Summers (1990) for a 'noise trader' model.
[4] See Bollerslev et al (1990) for stock price and trading volume's serial correlation. Stock returns also contain small, but significant serial correlations. Fama (1991) and Campbell, Lo and MacKinlay (1996) provide good surveys of this large body of research.

150 trading periods. This confirms Tirole's (1982) 'no-trade' theorem. In Experiments II and III, when we first allow for momentum traders and then noise traders to enter into the market, a vivid phenomena of "market psychology" becomes observable. Momentum trading appears as a profitable activity, and temporary bubbles and crashes occur over time. Trading volume tends to be high, and price volatility exhibits persistence.

The remainder of the paper is organized as follows. Section 2 gives a literature review. Section 3 describes our basic asset pricing model, market structure and the double auction trading mechanism. Section 4 provides the design and the results from the computational experiments. The market simulation results are evaluated, tested and compared with the empirical findings in the literature. Section 5 concludes and suggests future extensions to the work reported herein.

2. Literature Review

Laboratory asset markets are a useful tool for understanding relative performance because they permit the controlled manipulation of the rules and procedures that constitute a trading mechanism.

Many hypotheses can be investigated under laboratory environment. The Santa Fe stock market (Arthur, etl.(1997), and LeBaron, etl.(forthcoming 1998)) experiment investigate market efficiency and price convergence with a rational expectation asset pricing model, genetic algorithm learning and Walrasian tatonnement for market clearing mechanism. This section reviews some other experiments.

After conducting of many real trading experiments in laboratories, Gode & Sunder(1993) were interested in just how much "intelligence" was necessary to generate the results they were seeing. They ran a computer experiment with agents who are almost completely random in their behavior. They use a double auction market similar to those used in many laboratory experiments. Their results show that the budget constraint is critical to achieve the market efficiency.

Arifovic (1996) considers a general equilibrium foreign exchange market in a overlapping- generations environment involving genetic algorithm(GA) learning. The results are very similar to those in the experiments. The exchange rate fails to settle to any constant value, but the first-period consumption-level is quite stable.

Lettau (1997) implements many of the ideas of evolution and genetic algorithm(GA) learning in a very simple setting which provides a useful benchmark. Results show that the GA is able to learn the optimum portfolio weight, but with an interesting bias. This bias implies that they generally will be holding more of the stock than their preferences would prescribe.

In Beltratti, Margarita & Terna (1996), the setup is different from earlier markets in that the traders trade on a decentralized market in an one-to-one random matching basis and agents forecast future prices using an artificial neural network (ANN). Traders are allowed to buy a more complicated neural network at a given cost. They analyze stock price dynamics and the fraction of agents for each type. They find what appear to be cost levels at which both types can coexist, and other cost levels for which “smart” and “naïve” dominate the market. The results show that in the early stages of the market when prices are very volatile it pays to be “smart”.

Chan, LeBaron, Lo and Poggio (1998) construct an asset market experiment to investigate information dissemination and aggregation. They find in tests of information dissemination, price can convergence to the fundamental value in almost all cases. But convergence is difficult to achieve in price aggregation in many cases. To test the robustness of the results, technical traders and noise traders are included.

The Santa Fe Institute (SFI) Stock Market is one of the most complex artificial markets in existence. It is described in detail in Arthur, Holland, LeBaron, Palmer & Tayler (1997), and LeBaron, Arthur & Palmer (forthcoming 1998). This model attempts to integrate the trading mechanism into a well-defined economic structure, along with inductive learning using a classifier-based system. One feature of the SFI model is that it allows agents to explore a fairly wide range of possible forecasting rules. Agents have flexibility in using and changing their trading strategies. A major property (and also a major restriction) in their market is the assumption of homogeneity among agents. Except for the divergence of their rule sets, the agents are quite identical, regardless of their information set, signal processing ability, and learning mechanism (a Genetic Algorithm or GA). Moreover, their model adopts the Walrasian auctioneer mechanism,[5] which provides a simple and elegant way for viewing the price-setting process, but is not a typical mapping for real stock markets.[6] To simulate stock markets more realistically, a more complicated trading mechanism needs to be used.

3. Market Structure

3.1 Rational expectations in an asset market

We set up a model of an asset market along the lines of Grossman and Stiglitz (1980) and Hussman (1992), which is used by SFI stock market (Arthur, etl.(1997), and LeBaron, etl.(forthcoming 1998)). Consider a market in which N traders, indexed by j, decide on their desired asset composition between a risky stock paying a stochastic dividend and a risk-free bond. These traders formulate their expectations separately, but are identical in other respects.

Traders are assumed to have the constant absolute risk aversion (CARA) utility function. They communicate neither their expectations nor their buying or selling intentions to each other. Time is indexed by t and the horizon is indefinite. The risk free bond is in

[5] The price formation process could be captured by the general representation of a Walrasian auctioneer who aggregates traders' demand and supplies to find a market clearing price. This begins with each trader submitting his demand function (not necessarily the demand) to the auctioneer. The auctioneer announces a market-clearing price where aggregate demand equals aggregate supply. Then traders determine their optimal demand or supply according to this price. There is no trading allowed outside of equilibrium.

[6] Only the London gold market approximately resembles the Walrasian framwork.

infinite supply and pays a constant interest rate r. Stock is issued in $\overline{Q}$ units, and pays a dividend, d_t, which follows an exogenous stochastic process $\{d_t\}$ not known to the traders.

Each trader attempts to maximize next period's wealth by optimizing the allocation between the risky asset and the risk-free asset as follows:

$$\underset{Q_t^j}{Max}\, E_t^j[-\exp(-\lambda W_{t+1}^j)], \tag{1}$$

subject to

$$W_{t+1}^j = (1+r_t)W_t^j + Q_t^j[P_{t+1} + D_{t+1} - (1+r_t)P_t] \tag{2}$$

where λ is the constant coefficient of risk aversion, P_t is the stock price and E_t^j denotes the expectation of trader j conditional on the set of information Ω_t^j available to j at time t, including historical information on market prices and dividends. And r_t is the risk free rate at time t, which is fixed in this paper, i.e., $r_t = r$.

Assume for the moment that the agent's conditional predictions at time t of the next period's price and dividend are normally distributed with known mean and variance. It is well known that under CARA utility and Gaussian distributions for predictions, the agent's demand, Q^j_t, for holding risky asset is given by:

$$Q_t^j = \frac{E_t^j[P_{t+1} + D_{t+1} - (1+r)P_t]}{\lambda \sigma^2_{P_{t+1}+D_{t+1}} \mid \Omega_t^j} \tag{3}$$

The solution to this problem has a simple return/risk interpretation. Traders hold positions that are linear in the difference between the expected returns of the stock and the rate of interest. Risk aversion and the existing forecast error prevent traders from taking an infinite position based on expected return differentials.

The dividend process, thus far, is arbitrary. In the experiments we carry out below, we specialize it to an AR(1) process[7] .

$$D_{t} = rD_{t-1} + x_{t} \qquad x_t \sim N(0,\ s_x^2) \qquad (4)$$

Under the full information and homogeneous expectation, every trader observes each of these components in dividend process and hold identical expectations. The denominator in the demand function in equation (3) may be abbreviated to σ^2, and the index *j* may be omitted from the expectations operator.

The total amount of stock is denoted $\overline{Q}$. Market clearing then implies (Hussman (1992))

$$P_t = \frac{1}{1+r}(E_t[P_{t+1} + D_{t+1}] - (\overline{Q}\, l\, s^{\,2} / N)) \qquad (5)$$

This can be solved forward by Law of Iterated Expectations, to yield

$$P_t = E_t[\sum_{k=1}^{\infty} \frac{D_{t+k}}{(1+r)^k}] - y\ /r \qquad (6)$$

where $y = l\, s^{\,2} \overline{Q} / N$.

Note that the equilibrium price is simply the present discounted value of the expected future dividends, less a constant risk factor y.

Finally, as a benchmark for evaluating the equilibria in simulation sections, it will be useful to derive analytical expressions for price and dividend variances under full information. Using (4) and (5), the expected return on the stock in equilibrium is found to be

$$E_t\ (P_{t+1} + D_{t+1}) = \frac{\rho}{1+r-\rho}\ D_t + \lambda\sigma^2\ \overline{Q} / N \qquad (7)$$

the variance of the forecast error is

---

[7] There are two reasons for choosing this AR(1) process. First, it is common and simple specification among practitioners. For example, Santa Fe artificial stock market uses same dividend specification with

$$\sigma^2 = (\frac{\rho}{1+r-\rho}+1)^2\sigma_\xi^2 \quad (8)$$

In this section, it was possible to derive an analytical solution for price because all variables were directly observable and traders share same information set. In the following section, when the underlying components of the dividend process are not directly observable, the key question is how the individual next-period price and dividend expectations might be formed. We will consider this issue in the following section.

3.2 Heterogeneous agents and their expectations

We now have a simple two-asset market. We assume this market is populated by three types of traders: value traders, momentum traders and noise traders. Among them, value traders use the rational expectation model outlined in section 2.1 to determine their demands. The goal of value traders is to build forecasts of future prices and dividends to incorporate in their demand functions as described above.

Value traders are fundamental ones who believe that the current stock price reflects the discounted sum of the stream of future dividends. Their existence is expected to keep market prices close to its rational expectation equilibrium. Momentum traders are technicians who believe that the future price movement can be determined by examining the past patterns represented by various moving averages. Noise traders are not the rational ones, they trade only for liquidity.

3.2.1. Value traders' estimation

Value traders assume the expected dividend grows with a constant rate, *g*, as in Gordon dividend growth model, which is

---

us. Second, this assumption is consistent with the Gordon dividend growth model that we assume value trader uses.

$$E_t(D_{t+1}) = D_t(1+g) \tag{9}$$

Substituting value traders' expectation equation (9) into alternative expression of market clearing price equation (6), we have value traders' forecasting for the market price

$$\hat{P}_t = D_t \frac{1+\hat{g}}{r-\hat{g}} - \frac{y}{r} \tag{10}$$

If the dividend expectations are unbiased, dividend forecasts will be upheld on average by the market and, therefore, the price sequence will be in rational expectation equilibrium. Thus, the price fluctuates as the information $\{\Omega_t\}$ fluctuates over time, and it reflects "correct" or "fundamental" value.

This forecasting equation is one of the crucial items that traders need to estimate. It is used as a benchmark for our experiments.

For value traders, dividend growth rate g is the only parameter that they need to learn.

We assume they employ artificial neural network (ANN) to learn this dividend growth rate, g. We assume each value trader possess a ANN model, similar structure but different in parameters. Traders then learn by training their model according to newly revealed market information at the end of each period, and then use their newly updated model to forecast the next period state. This structure will offer better desirable properties comparing to traditional learning method (say, least squares): It avoid any presumption on beliefs, but it will allow the individuality of expectations to converge over time. And it will better mirror actual cognitive process, in which different individual might well "cognize" different patterns and arrive at different forecast from the same market information.

In this market, the value traders use Elman Recurrent Neural Networks[8] to forecast dividend growth. Such ANN- based procedures have become quite popular among practitioners. In option pricing, Hutchinson, Lo and Poggio (1994) use four different types of ANN models to estimate the pricing formula of a derivative asset. Brown, Goetzmann and Kumar (1998) use Elman Recurrent Neural Net modeling to replicate Hamilton's timing strategies. They do find evidence to support the forecasting ability of this most famous Wall Street chartist.

The dividend growth rate is the only parameter that value traders need to learn using an Artificial Neural Net (for more details, see Appendix A). Value traders use the information set $\Omega_{t-1}=\{ g_{t-1}, g_{t-2}, g_{t-3}, g_{t-4}, P_{t-1}, P_{t-2}, D_{t-1}, D_{t-2} \}$ as the input to forecast the next dividend growth rate, g. They then calculate the next period's dividend and price using the Gordon constant dividend growth model and equation (9). This implies that the value trader believes that the current market price will reflect the cash flows as represented by future dividends. The value trader compares this forecasted price with the existing market price to form her best ask or bid, and then transacts according to different scenarios. For example, a value trader enters into the market and finds the current best ask is lower than his forecasted price, and will conclude that the stock is under-valued. Therefore, he will buy the stock at that best ask. On the other hand, if he finds the current best bid is higher than his forecasted market price, then he will sell at this best bid with a inference that the stock is over-valued.

3.2.2 Momentum trader

Momentum traders are technicians who believe that future price movement can be determined by examining past price patterns as represented by various moving averages (MAs). Several studies support the notion that the price history can be used to predict

[8] Elman (1990) develops a neural net model that allows the intermediate series, or the so-called 'hidden layer', to be used as original inputs. This simply allows for a richer potential nonlinear interaction among series. It is well suited for time-series analysis where the objective is to identify temporal patterns.

market movements,[9] partly due to the “self-fulfilling”. The moving average trading rule states that when the short-term (usually 1-to-5 day) moving average is greater than the long-term moving average (usually more than 50 days) [c.f. Brown et al.(1998)], a rising market is indicated. Correspondingly, the trading rule then generates a buy signal. Based on such market trends, the momentum trader decides to enter or exit the market.

The traders are divided into two groups according to their choice of trading rules. The first group of momentum traders compares the current market price P_t with MA(5). That is,

If P_t>MA(5), they buy shares according their demand function.

If P_t=MA(5), they hold their current position.

If P_t<MA(5), they sell their stock inventories.

The second group of momentum traders identifies a trading opportunity by comparing MA(5) with MA(10). Specifically:

If MA(5) > MA(10), they buy shares according their demand function.

If MA(5) = MA(10), they hold their current position.

If MA(5) < MA(10), they sell their stock inventories and exit the market.

3.2.3. Noise trader

Noise traders do not act consistently and strategically. Ignoring all the available information, they have no memory or learning abilities. Instead, they randomly post orders to buy or to sell at the market price.

3.3 Trading Mechanism: Double Auction

Any trading mechanism can be viewed as a type of trading game in which players meet at some venue and act according to some rules. The rules are used in determining the

[9] In a study by Brock, Lakonishok, and LeBaron (1992) which uses bootstrapping techniques, the ability of two simple trading rules are used to predict the movements of the Dow Jones Industrial Average (DJIA).

operation of the market. These rules dictate who can trade and when and how orders can be submitted, who may see or handle the orders, how orders are processed, and how prices are set eventually. It has been recognized that these rules are very important because they can affect bidding incentives, and therefore the terms and the efficiency of an exchange.

The literature has identified many different auction mechanisms in the world. In the following, we describe two different mechanisms for determine price: Walransian tatonnement and double auction. Walransian tatonnement auction is a path-independent process and no transaction will occur outside the equilibrium. While in a continuous double auction we used in this paper, the price is path-dependent. The time upon which the orders arrive is relevant. Transaction is allowed to occur outside the equilibrium and the trading volume is determined by the short side of the demand and supply.

Walransian tatonnement auction has long served the need of general equilibrium price theory. In this type of auction, the time upon which the orders arrive is not relevant, since all the order will be executed at a single equilibrium price. And indeed, such a market approximates the London gold market. The auctioneer suggests an opening price. The traders then indicate whether their firm is a net buyer or seller at that price. If at the starting price there are no sellers, the price is raised by varying amounts until one or more of the traders indicates that he is a seller at the standing price. Similarly, the price is moved down if there are no buyers at that price. At this juncture, the auctioneer asks for the net quantities each trader wishes to buy or sell. If the total indicated purchase quantities do not match the quantity offered by the traders the price is further adjusted until a match occurs. Then all the order will be execute at this ‘fixed’ price. No trades occur out of equilibrium.

In a double auction (DA), transaction prices arises from a two-sided auction where buyers improve (raise) bids and sellers improve (lower) asks until one of the buyers and one of the sellers reach agreement. The orders are cleared with the priority of price, size and time of arrival.

We chose this mechanism for two reasons. First, major stock, commodity, currency, and many other markets are organized as double auctions. A classic example of a double auction market is the NYSE. Second, laboratory double auctions with human traders are known to yield data that approximate the equilibrium predictions of economic theory in a variety of environments[10].

In our simplified double auction market, traders can either submit a bid/ask, or accept an existing bid/ask. If a bid exists for the stock, any subsequent bid must be higher than the current one. Similarly, a subsequent ask to an existing ask must be lower than the current one. A transaction occurs when an existing bid /ask is accepted. Given the absence of opportunity to borrow or short sell, traders have to trade subject to their budget constraints.

At the beginning of each trading round, we suppose a random permutation of the traders, which determines the subsequent order of traders. Initially, the value traders come to the market with their own expectations about price and attempt to post or accept a bid (ask) order by comparing their expected price with the existing best ask (best bid). They can only observe the best bid and the best ask price. Momentum traders go to the market with their "buy", "sell" or "hold" signals and submit market orders according to these signals. Noise traders enter the market without any expectation and trading strategy, and rather randomly submit a market order of buying or selling for some liquidity reason.

These traders are endowed with certain stock shares and cash at the beginning. They submit bid and ask order and trade with each other according to the following scenarios :

- If a best bid, b, and a best ask, a, exist on the market, trader compare his expected price $E(P)$ with these bid and ask prices,

  --If $E(P) > a$, he will post a market order, buy at this ask price;

  --If $E(P) < b$, he will post a market order, sell at this bid price;

[10] Chan, LeBaron, Lo and Poggio (1998) use the same auction mechanism.

-- If $b< E(P)< a$ and *E(P) is lower than the midpoint of the bid and ask*, he will post a sell order at a price uniformly distributed as *[E(P), E(P)+ S]*, where S is the maximum spread from the expected price.

-- If $b< E(P)< a$ and *E(P) is higher than the midpoint of the bid and ask*, he will post a buy order at a price uniformly distributed as [E(P) -S, E(P)]

- If only the best ask, *a*, exists

-- If $E(P) > a$, he will post a market order, buy at this ask price;

-- If $E(P)< a$, he will post a buy order at a price uniformly distributed as *[E(P) -S, E(P)]*

- If only the best bid, *b*, exist

-- $E(P) < b$, he will post a market order, sell at this bid price;

-- $E(P) > a$,he will post a sell order at a price uniformly distributed as *[E(P), E(P)+S]*

- If no bid and ask exist, he will has an equal chance to post a buy or a sell order at prices uniformly distributed as *[E(P)-S, E(P)] or [E(P), E(P)+S]* respectively.

The final issue should be considered is quantity each trader will trade. In full general equilibrium enviorenment, trader determined their demand according to their utility function. In this paper, the amount traded is fixed to be an unit share[11]. Initial endowment of each trader will also be fixed.

3.4 Experiments

To isolate the effect of market rules and agent behavior on market performance, we conduct three experiments. In experiment I, we select only value traders as market

[11] It is important to remember that the quantity of transactions is an important aspect of the market. It could be associated with agents' risk aversion and determined by the demand function we presented in the last section.

participants. In experiment II, we observe the performance of a double auction with both value traders and momentum traders, and compare their relative portfolio performance. In experiment III, we introduce 'noise traders' into the market and analyze their effect on the market dynamics.

The rational expectation equilibrium is given by equation (5), which will be used as our benchmarks for our three experiments. The traders are assumed to work in the following environment:

(1). The risk-free interest rate, *r*, is fixed to be 0.1; risk aversion parameter, λ , is set to be 0.5.

(2). Throughout the experiments, we specify the dividend process as

$$D_t = \rho D_{t-1} + \xi_t \qquad \xi_t \sim N(0, s_x^2) \tag{9}$$

Where $\rho = 1.002$, $\sigma^2_\xi = 0.007$.

(3). Traders are assumed to hold endowment of 2 units of and 2000 dollars in cash.

(4). Value traders are assumed to be equipped with an ANN learning mechanism, Elman recurrent network. The neural network is programmed to learn the value of the parameter in dividend process. The initial starting value is randomly draw from an uniform distribution. The details are given in an appendix.

(5). Momentum traders are initially endowed with "hold" signals until they get enough information to form their trading rules. Moving averages of the stock price are calculated with equal-weighted parameters. Momentum traders using MA(5) or MA(10) wait for 5 or 10 periods respectively to enter the market. They enter the market with their "buy" or "sell" signals generated from their trading rules.

(6). Noise traders only trade for liquidity, who randomly buy or sell at the market price.

(7). An unit of time correspond to a trading period. Each trading period begins with permutation of the traders to determine the order of the participation of the market. The first trader submit a sell or buy order with an equal probability at the price which draw uniformly from *[E(P), E(P)+S]* or *[E(P)-S, E(P)]* respectively. The next trader faces this buy or sell price and act according with the rule described as above. The trading period is completed when each trader is considered in turn.

(8). The market is run for 500 periods for experiment I and II.

- Experiment I consists of 20 value traders, with constant order size.

  The purpose of this experiment is to test the convergence of prices to the REE prediction and "no trade theorem". Value traders serve as the only force for the market price to reach the share's fundamental price.

- Experiment II consists of 10 value traders and 10 momentum trader with constant order size.

  The momentum traders should add some noise to our market. In this experiment, we focus on the effects of adding the momentum traders to the market. It is interesting to observe whether the value traders can maintain an efficient market given the existence of some erroneous signals caused by the presence of momentum traders.

- Experiment III consists of 10 value traders, 10 momentum traders and 5 noise traders, all with constant order size.

  In this experiment, we allow for entrance of noise traders into the market. These traders do not follow any consistent trading rule. Instead, they randomly submit bid or ask market orders.

  In this experiment, we are interested in observing the noise trader's performance. In particular, we are interested in whether or not, and for how long noise trader survive in the market.

4. Computational results and time series features

4.1 Computational results

In our time-series generating experiments, the market is run for 500 time periods for the first and second experiments, and for 1000 time periods for the third experiment. Convergence to the REE occurs in Experiment I, the identical trading strategy case. Such convergence occurs with difficulty or is unattainable in experiment II and III that involve more than one type of traders. This can be explained by the fact that in the first experiment, all the traders are fundamental traders. Their forecasts are supervised by the fundamental price. In the cases that include momentum and noise traders, no price convergence occurs. The reason is that momentum traders only obey their own trading rules (e.g. MA(5) or MA(10)), and no rule exists for noise traders. Both of these trader groups disregard the views of value traders. In other words, the fundamental value of our stock is not common knowledge.

As is seen from Figure 1, the price series appears to be nearly identical to the price in the rational expectation regime. When the homogeneous rational expectation is reached, traders have no incentive to trade. At this point, trading volume approaches zero (c.f. Figure 2). When market participants move from value traders to a combination of 10 value traders and 10 momentum ones, we find that more complex patterns form in the collection of expectations. Market price displays characteristics that differ from the Rational Expectation Equilibrium (REE), and do not exhibit the tendency to converge (Figure 3). Comparing the market prices generated from the first two experiments, we find that momentum traders bring more volatility into the market (Figure 4).

A closer inspection of the results in Experiment III identifies richer and more complex market dynamics, especially in the first 500 periods. Prices diverge substantially from theoretical (fundamental) prices. The differences between the two price series provide systematic evidence of temporary price bubbles and crashes (Figure 5). This appearance of bubbles and crashes suggests that, momentum and noise trading have affected the market. We further examine market evolution in Experiment III by observing the changes in the accumulated wealth of traders. Value traders apparently have more forecasting power that makes it easier for them to outperform the other market players (c.f. Figure 7). The momentum traders perform better than the noise traders. Most of them are able to accumulate wealth and to survive much longer in the market. Some of the MA(10)-type momentum traders continue trading until the end of the experiment (c.f. Figures 8 and 9). Based on the wealth change for 5 noise traders presented in Figure 10, we see that these traders are very active during the beginning periods but exit the market after 150 periods because they run out of financial resources. In order to compare relative portfolio management performance, the wealth levels of the best and the worst traders are plotted on the same scale (c.f. Figure 11). This generates one of our most interesting experimental outcomes. A more rigorous comparison of this relative trading performance is conducted in section 3.3.

4.2 Time series features

It is interesting to contemplate what these price series would look like if they converged closely with the REE benchmark. Market price would always have an extra amount of variability relative to REE. It is difficult to estimate how much this would be. From Eq.(5), it is clear that increased price variability by itself merely lowers the price by a constant amount. We plot the difference between the actual and R.E.E. prices in Figure 5. In some periods, the price deviation is close to zero. However, in other periods, the market price breaks away from the theoretical price and generates some bubbles and crashes. In other words, the market is not always efficient over time.

This analysis lacks rigor. We now present some basic statistical features of and tests on the time series generated in our three markets. These results are presented in Tables 2 and 3.

Table 2 reports basic summary statistics for stock returns.

- As expected, the variance of the stock returns increases from Experiment I to III, as more noise exists in the third market. In all three experiments, especially in Experiment III, stock returns exhibit significant amounts of excess Kurtosis. This is consistent with the empirical observation that real asset returns are leptokurtotic.
- In the last column of Table 2, the average trading volume over 500 or 1000 periods is reported for each experiment. In Experiment I, 10 value traders use an ANN learning mechanism to forecast the next period's price. They are identical in terms of their

information set, information processing ability, and the fundamental model they employ to map dividends to price.

- After a certain period of learning so that homogeneous rational expectations are reached, trading ceases. This confirms Tirole's (1982) famous 'No-trading Theorem'.
- The normality tests provide strong evidence against the null hypothesis for all three experiments.

Table 2. Basic summary statistics for the stock returns and trading volumes for each of the three experiments

Experiments	Mean	Variance	Skewness	Kurtosis[1]	Trading Vol.[2](Ave.)
Exp. 1	0.12	0.32	0.8	3.2	0.46
Exp. 2	0.18	0.46	1.4	5.2	1.67
Exp. 3	0.21	0.62	2.6	18.71	2.52

1. The kurtosis reported is excess kurtosis.
2. The trading volume is the average value for each period over all trading periods.

In Table 3, we report the results from tests on normality, autocorrelations and volatility persistence. All of the tests are imposed on the stock return series.

Table 3. Diagnostic Tests on Stock Returns for Each of the Three Experiments

Returns	Normality	$\rho(1)$ (%)	Q_5	ARCH(5)
R1	10.39(0.00)*	3.67	27.8(0.00)*	1.64(0.94)

R2	16.2(0.00)*	-7.8	81.0(0.00)*	63.6(0.000)*
R3	45.6 (0.00)*	35	142 (0.00)*	226(0.000)*

==

Diagnostic tests of normality, unit roots, autocorrelations and ARCH for stock returns for each of our three experiments for 1000 trading periods are reported in this table. The test statistics are reported in the main row, and P-values are given in the parenthesis immediately beside each test statistic. Under the null hypothesis, the test statistics marked with asterisks indicate that the P-values are significant at the 5%.

- For the predictability of stock return, we examine the autocorrelations and the Box-Pierce Q-statistics for daily stock returns. Column (2) in Table 3 reports that the first-order autocorrelations of returns in Experiment III is 35%. Under the IID random walk null hypothesis, $\rho(1)$ is asymptotically normally distributed with mean 0 and standard deviation $1/\sqrt{T}$ (c.f. Campbell, Lo and Mackinlay (1997)). This yields a standard error of 3.1% for $\rho(1)$. Hence an autocorrelation at the 35% level is clearly statistically significant.
- Moreover, the Box-Pierce Q-statistic with 5 autocorrelation in column (3) is significant at all conventional significance levels. To develop a sense of the economic significance of the autocorrelations that are reported in this table, we observe that the R^2 of a regression of returns on a constant and its first lag is 12.3%. Therefore, an autocorrelation of 35% implies that 12.3% of the variation in the daily return is predictable using the preceding day's return.
- The last column of Table 3 delivers the results of a simple test for ARCH dependence in the residuals.[12] The ARCH effect is accepted only for Experiments II and III. This implies that the volatility of asset returns appears to be serially correlated.

4.3 Portfolio performance

In this section we evaluate portfolio performance over groups. In our artificial economy, each trader manages a portfolio composed of two assets, a stock and a risk-free bond.

[12] This is the test for conditional heteroskedasticity, as proposed by Engle (1982)

Each trader decides on his desired position between these two assets in every period. The return earned on their portfolios is measured by the growth of their wealth from the beginning to the end of each experiment. We conduct the evaluation of portfolio performance only for the 25 traders in Experiment III.

To do such, we first calculate the portfolio return based on changes in wealth over time. The average return for each group is reported in Table 4. Not surprisingly, the value trader group outperforms the other two groups in terms of return. The market returns are measured as the average of the 25 traders over each of the 1000 periods. This reflects average portfolio management ability. Based on the first column in Table 4, the momentum trader group slightly underperforms the market average, and the noise trader group underperforms the market. Since the noise traders trade very actively during the early periods when market prices experience high volatilities, noise traders run out of the money and exit the market after about 150 periods on average.

We then conduct a risk-adjusted performance evaluation based on the mean-variance criterion and the Capital Asset Pricing Model (CAPM). Our measures include the Sharpe measure, the Treynor measure and the Jensen measure. Both the Sharpe and Treynor measures examine excess return per unit of risk. While the former uses total risk, the later uses systematic or market risk.

Based on the results reported in the second column of Table 4, value traders have the highest excess return per unit of total risk, and noise traders earn a significantly negative excess return for each unit of risk. The Sharpe ratios are higher for the portfolios of value relative to momentum traders. The reason is that the volatility of returns is lower for the portfolios of value traders. In contrast, the ten momentum traders demonstrate quite different performances that result in high volatility.

The third column in Table 4 reports the Treynor values, where the values of b are extracted from Table 5. Value and momentum traders earn positive excess returns in terms of market risk, and noise traders earn negative market-risk-adjusted excess returns.

The CAPM model relationships are estimated across each type of trader. The estimated parameter for the alphas and betas are reported in Table 5. The value trader group demonstrates positive excess returns on their portfolio, which according to the EMH should not be significantly different from zero. But given the standard deviation in the parenthesis of 0.011, the value traders' excess return of 0.08 is statistically significant.

Table 4. Portfolio Returns for Value, Momentum and Noise Traders

Trader Type	Portfolio Returns	Sharpe Ratio* $(R_i - r)/s_i$	Treynor $(R_i - r)/b_I$
Value Trader (10)	20.42%	64.11%	5.18%
Momentum (10)	10.87%	8.06%	0.072%
Noise Traders (5)	-5.6%	-70.46%	-7.4%
Market **	**11.39%**	**7.64%**	**------**

*The Sharpe ratio divides average portfolio excess return over the sample period by the standard deviation of returns over that period. It measures reward-to-total volatility.

** Market returns are calculated as the average over all 25 traders for each of the 1000 periods.

Table 5. The Jensen alpha values and beta values of the portfolio returns of the value, momentum and noise traders

Traders Type	Jensen alpha (α)	Portfolio Beta (β)
Value Trader	0.08 (0.011)	2.01 (0.076)
Momentum Trader	-0.005 (0.001)	1.21 (0.015)
Noise Trader	-0.17 (0.09)	1.42 (0.016)

The estimated parameters are reported in the main row and the standard deviations of the estimators are reported in the parentheses immediately beside each test statistic.

5. Conclusions and Directions for Future Research

Our results show that the artificial stock market is able to replicate certain empirical puzzles observed in real markets. Among these are predictability, excess volatility, volatility persistence and positive excess returns. Our model also is capable of generating behavior close to REE under certain circumstances. This shows that, when the heterogeneity of traders increase, the market is driven by more noisy factors than just fundamental trading. As a result, it becomes more difficult for traders to forecast each other's forecasts. Therefore, market prices do not converge.

One restriction in our market is a fixed order size. In future work, when this assumption is relaxed, we should be able to observe more complex market behaviors. Another interesting extension to this initial work is related to the choice of market mechanism. In this paper, we use a double auction. It would be appealing to simulate different types of auctions or to endow some participants with more market power to observe whether and how they could corner the market.

Other directions for model enrichment include stochastic behaviors of interest rates, dividend growth rates, and risk aversion parameters. It would be interesting to examine if and how interest rate changes add to market complexity, and if and how they affect the learning process of agents. Moreover, it would be interesting to determine the impact on trading behavior of endowing different agents with different risk aversion levels.

This set of computer experiment should be viewed as an earlier inquiry in this growing area of research into market dynamics. The prototypes are probably still somewhat below more generally accepted analytical methods. However, to adequately explain the growing list of real market puzzles, analytical methods also will need to become more advanced.

Appendix (Artificial neural networks)

Recently, Artificial Intelligence (AI) techniques have received much attention in the economics and business communities. Artificial Neural Networks (ANN) may be viewed as being a generalized, and nonparametric estimation technique.

ANN [13]algorithms are statistical procedures to fit a reduced functional form to the data. They are similar in content to stepwise regressions in that they try a wide range of model specifications to reduce in-sample residual variation. But, unlike stepwise regressions, their specifications are not necessarily linear. The form of ANN model need not be explicitly specified. In fact, the innovation of ANN models is that they offer a parsimonious but flexible nonlinear specification. Campbell, Lo and Mackinlay (1997) provide a general overview of the applicability of neural net modeling to financial series prediction problems. Recent applications of ANN methods to financial markets suggest that nonlinear dynamics are potentially important characteristics of markets. Current research shows that ANN models have the potential to uncover sophisticated nonlinear processes that lead to price changes in financial markets.

ANNs are simply a good tool to fit in-sample. Indeed, given enough calculation time and enough hidden layers, ANNs can fit in-sample perfectly. However, the resulting functions may have no economic interpretation. Not only are there no "standard errors" for the coefficient estimates, but there may be no identifiable coefficients since input data are recombined as intermediate variables in the hidden layers. Finally, as with all curve fitting, there is no guarantee for out-of-sample performance.

[13] Smith (1993) is a good introductory book that concentrates on one type of neural network; namely, feedforward neural networks, one of the most commonly used NNs. Wasserman (1993) is an intermediate level book, which provides a good introduction to different types of neural nets. Other learning paradigms are also discussed. Hertz, Krogh and Palmer (1991) provide a clear and concise description of the theoretical foundations of neural networks using a statistical mechanical framework.

To eliminate these drawbacks associated with the standard NN modeling procedure, we run two diagnostic tests on our computer based version of the model before we conduct our experiments. In the first diagnostic test, we examine whether an ANN can replicate the Rational Expectation Equilibrium (REE). We do this by calculating the analytical value of the market clearing price. Then, we run the computations with all predictors "clamped" to these calculated REE parameters. We do find such predictions are upheld. This assures us that the computerized model is working correctly. In the second test, we endow value traders with a given dividend series and a calculated corresponding R.E.E. price series. We then test whether the ANN individually learns the correct parameters. Since the ANN does with some variation, this assures us that our neural nets are learning properly.

Our setting for value traders uses a 10-4-1 Elman recurrent neural network. This involves 10 inputs, one hidden layer with 4 neutrons, and one output, g_t. The Elman recurrent network feeds half of the hidden layer neutrons back to the input layer.

The neural net forecasting can be divided into four steps. First, at the beginning of each period, we include the most recent 8 previous pieces of information on growth, the last two period's market price and dividend, and the updated forecasting errors, that is, g_{t-1} to $g_{t-4}, P_{t-1}, P_{t-2}, d_{t-1}, d_{t-2}$ in the input layer. Each trader invokes his own network to forecast current growth E(g). Second, we calculate the expected price and dividend using the Gordon dividend model. Third, at the end of this period when the market price and dividend is revealed, we extract the targeted g by solving the inverse of Eq.(9). Fourth, we retrain the neural nets by using this target and get a new network with weights updated by minimizing the forecasting errors by using a back-propagation algorithm.[14] This is employed to do next period's forecasting. For each training, the error target is set to be 0.0001 and the training epoch is set to be 50 in each period.

[14] White (1989) proves that backpropagation can learn arbitrary mappings.

References

Admati, A.R.(1991), The informational role of prices, *Journal of Monetary Economics* 28, 341-361.

Arifovic, J.(1996), The behavior of the exchange rate in the genetic algorithm and experimental economics, *Journal of Political Economy* 104, 510-541.

Arthur, Holland, LeBaron, Palmer & Tayler (1997), Asset pricing under endogenous expectations in an artificial stock market, in W.B. Arthur, S. Durlauf & D. Lane, (eds), *The Economy as an Evolving Complex System II*, Addison-Wesley, Reading, MA, 15-40

Back, T.(1996), *Evolutionary Algorithm in Theory and Practice*, Oxford University Press, Oxford, UK.

Belttratti, Margarita and Terna (1996), *Neural Networks for Economic and Financial Modeling*, International Thomson Computer Press, London, UK

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

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

Brock,W.A., Dechert,W.D., Scheinkman, J.A. & LeBaron,B.(1996), A test for independence based on the correlation dimension, *Econometric Reviews* 15(3),197-235.

Brock, W.A. & Hommes, C.H.(1998), Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *Journal of Economic Dynamics and Control* 22(8-9),1235-1274.

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

Brock, W.A. & leBaron,B. (1996), A dynamic structural model for stock return volatility and trading volume, *Review of Economics and Statistics* 78, 94-110.

Brown, Goetzmann and Kumar (1998), The Dow Theory: William Peter Hamilton's

Track Record Reconsidered, *Journal of Finance,* Aug, 1998

Campbell, J. & Shiller, R. (1989a), Stock prices, earnings and expected dividends, *Journal of Finance* 63, 661-676.

Campbell, J.Y., Lo, A.W. & MacKinlay,A.C. (1996), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.

Casdagli, Martin, and Stephen Eubank, (1992), *Nonlinear Modeling and Forecasting,* A Proceedings Volume in the Santa Fe Studies in the Sciences of Complexity, Addison-Wesley Publishing Company, Reading, Mass.

Chan, LeBaron, Lo and Poggio (1998) Information dissemination and aggregation in asset markets with simple intelligent traders, Working Paper, MIT.

Cliff, D. & Bruten, J. (1997), Zero is not enough: On the lower limit of agent intelligence for continous double auction markets, Technical Report HPL-97-141, Hewelett-Packard Laboratories, Bristol, UK.

Conlisk, J. (1996), Why bounded rationality?, *Journal of Economic Literature* 34, 669-700.

De Long, J.B., Shleifer,A. & Summers, L.H.(1990), Noise trader risk in financial markets, *Journal of Political Economy* 98,703-738.

Elman, Jeffrey, 1990, Finding structure in time, *Cognitive Science* 14, 179-211.

Engle, R.F.(1982), Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987-1007

Fama, E.F.(1991), Efficient capital markets: II, *Journal of Finance* 45, 1575-1617.

Fogel, D.B. (1995), *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*, IEEE Press, Piscatway, NJ.

Forsythe, R., Palfrey, T.R. & Plott, C.R.(1982), Asset valuation in an experimental market, *Econometrica* 50, 537-567.

Gode and Sunder (1993), Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality *Journal of Political Economy, 1993 Vol.101.No.1*

Grossman, S. (1976), On the Efficiency of Competitive Stock Market Where Trades Have Divese Information, *Journal of Finance, Vol. XXL. No.2*

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

Grossman, S. (1989), *The Informational Role of Prices,* The MIT Press, Cambridge, MA.

Hertz, John, Anders Krogh, and Richard Palmer, 1991, An Introduction to the Theory of Neural Computation, Addison-Wesley Publishing Company, Reading, Mass.

Hussman (1992), Market Efficiency and Inefficiency in Rational Expectations Equilibria, *Journal of Economic Dynamics and Control*, *16.*

Hutchinson, J, A. Lo and T. Poggio, (1994), A nonparametric approach to to the pricing and hedging of derivative securities via learning networks, *Journal of Finance, 49, 851-889*

Jamal, K. & Sunder, S. (1996), Bayesian equilibrium in double auctions populated by biased heuristic traders, *Journal of Economic Behavior and Organization*31,273-91.

Kocherlakota, N. (1996), The equity premium: It's still a puzzle, *Journal of Economic Literature* 34(1), 42-71.

Kurz, M. (1994), On rational belief equilibria, *Economic Theory* 4, 859-876.

Kurz, M & Beltratti, A. (1995), The equity premium is no puzzle, Technical Report, Stanford University, Stanford, CA.

LeBaron, B. (1998b), Technical trading rules and regime shifts in foreign exchange, *in* E. Acar & S. Satchell, (eds), *Advanced Trading Rules*, Butterworth-Heinemann, 5-40.

LeBaron, B. (forthcoming 1998a), Agent based computational finance: Suggested reading and early research, *Journal of Economic Dynamics and Control.*

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

Lettau, M. (1997), Explaining the facts with adaptive agents: The case of mutual fund flows, *Journal of Economic Dynamics and Control* 21,1117-1148.

Lux, T.(1998), The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions, *Journal of Economic Behavior* *and Organization* 33, 143-165.

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

Margarita,S. & Beltratti, A. (1993), Stock prices and volume in an artificial adaptive stock market, in *New Trends in Neural Computation*: *International Workshop on Artificial Networks*, Springer-Verlag, Berlin, 714-719.

O'Hara (1995), *Market Microstructure Theory*, Blackwell Publishers, Cambridge USA.

Plott, C.R. & Sunder, S.(1988), Rational expectations and the aggregation of diverse information in laboratory settings, *Econometrica*56, 1085-1118.

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

Rust, J., Miller, J. & Palmer, R. (1992), Behavior of trading automota in a computerized double auction market, in D. Friedman & J. Rust, (eds), The Double Auction Market: institutions, Theories, and Evidence, Addison-Wesley, Reading, MA, 155-198.

Sargent, T. (1993) *Bounded Rationality in Macroeconomics*, Oxford University Press, Oxford, UK.

Simon, T. (1993), *The Sciences of the Artificial*, MIT Press, Cambridge, MA.

Smith, V. (1969) An experimental study of competitive market behavior, *Journal of Political Economics, 70, 1969*

Smith, Murray, (1993), *Neural Networks for Statisticians*, Van Nostrand Reinhold, New York, N.Y.

Steiglitz, K., Honig, M.L. & Cohen, L.M.(1996), A computational market based on individual action, *in* S.H.Clearwater, (eds.), *Market-Based Control: A paradigm for distributed resource allocation*, World Scientific, New Jersey, 1-27.

Tirole, Jean (1982), On the possibility of speculation under rational expectations. *Econometrica*, 50: 1163-1181

Wasserman, Philip, 1993, *Advanced Methods in Neural Computing* , Van Nostrand Reinhold, New York, N.Y.

Weigend, Andreas, and Neil Gershenfeld, (1994), *Time Series Prediction: forecasting the Future and Understanding the Past*, Addison-Wesley Publishing Company, Reading, Mass.

White, Halbert, (1989), Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings, *Neural Networks* 3, 535-549.

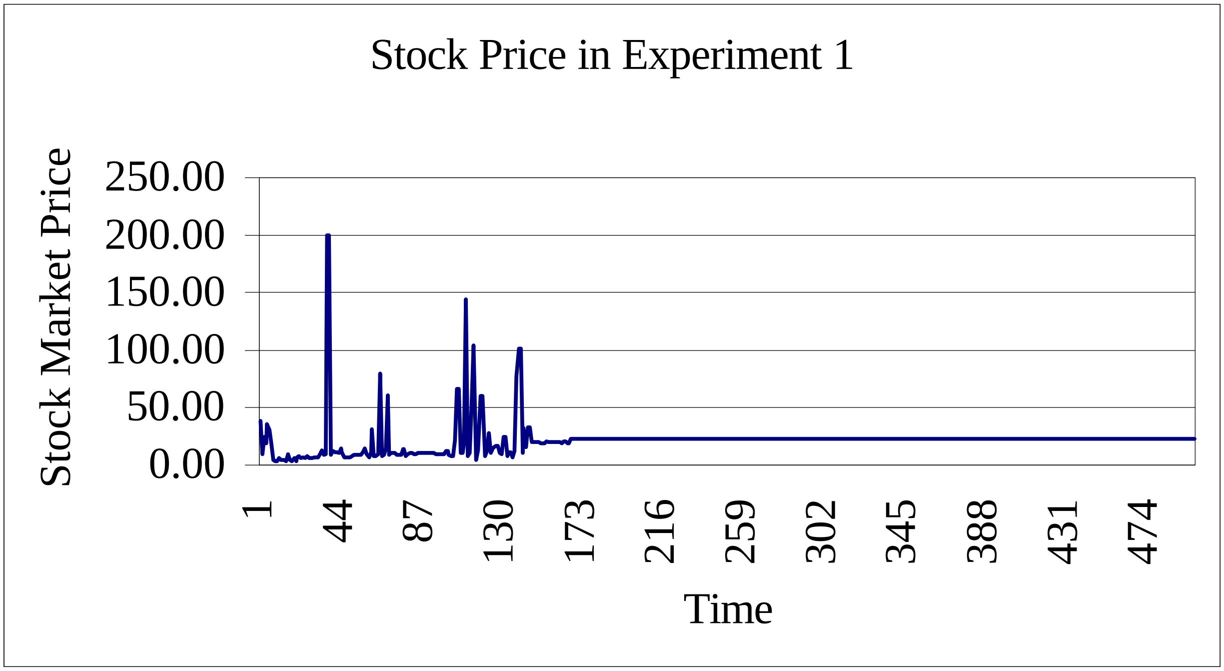

Figure 1: Stock Price in Experiment 1

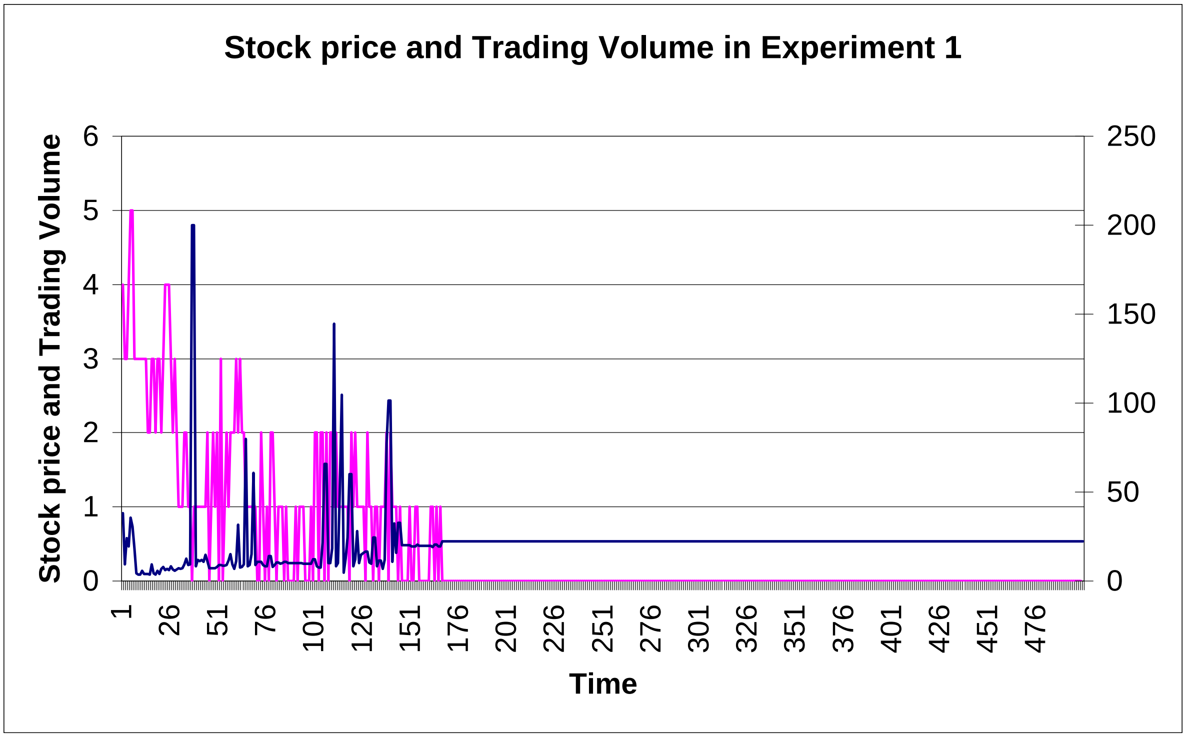

Figure 2. Stock Price and Trading Volume in Experiment 1

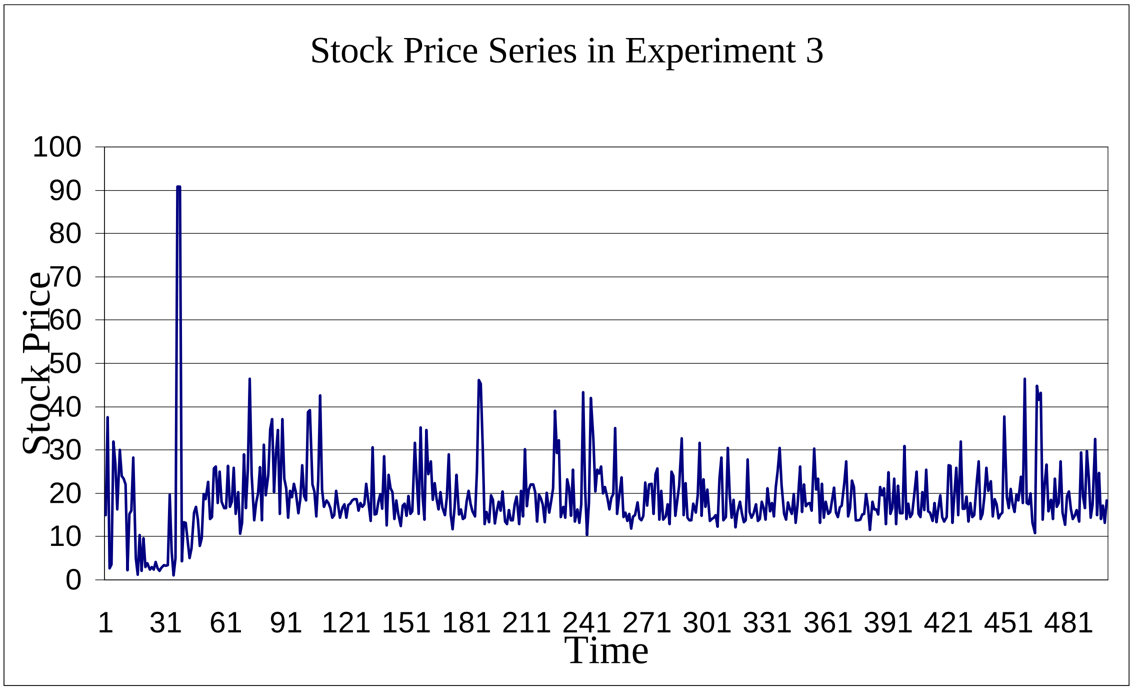

Figure 3. Stock Price Series in Experiment 2: 10 value traders and 10 momentum traders

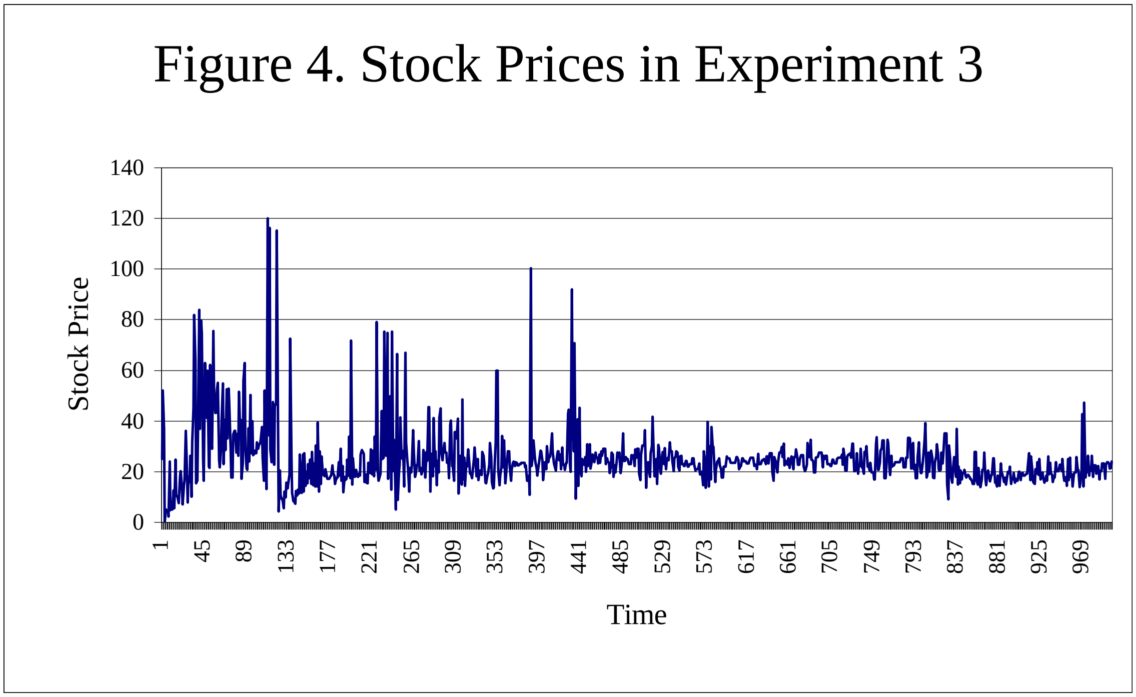

Figure 4: Stock Prices in Experiment 3: 10 Value Traders, 10 Momentum Traders and 5 Noise Traders

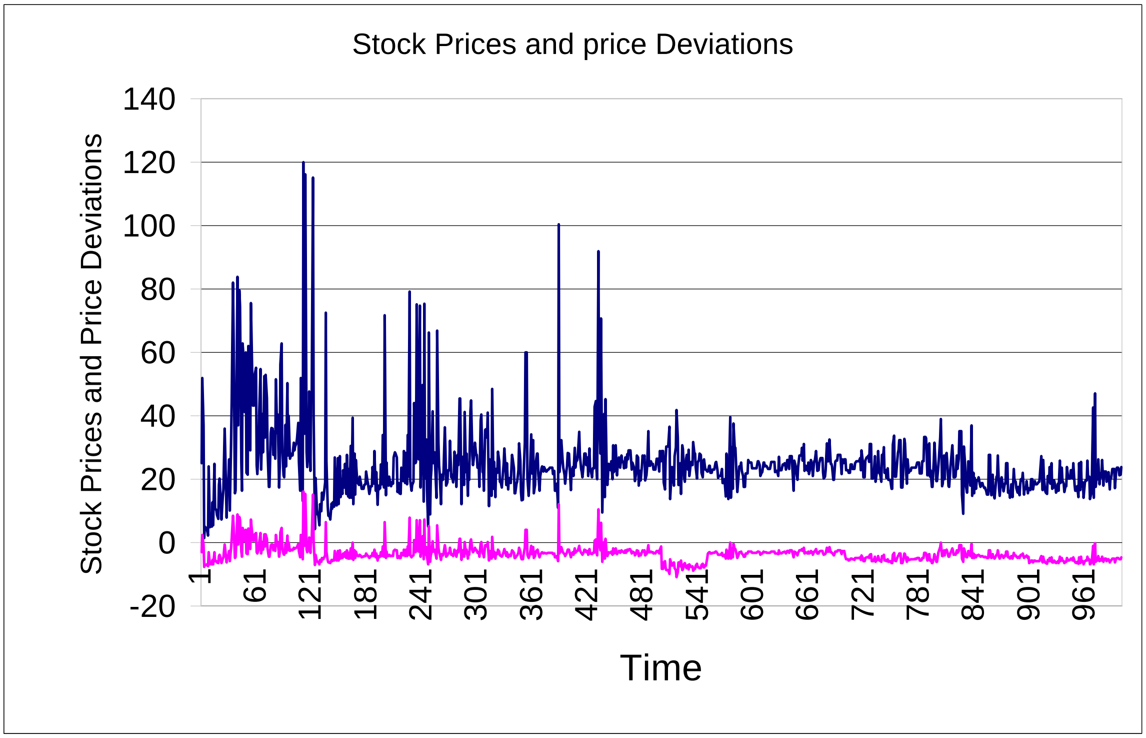

Figure 5: Stock Prices and Price Deviations in Experiment 3

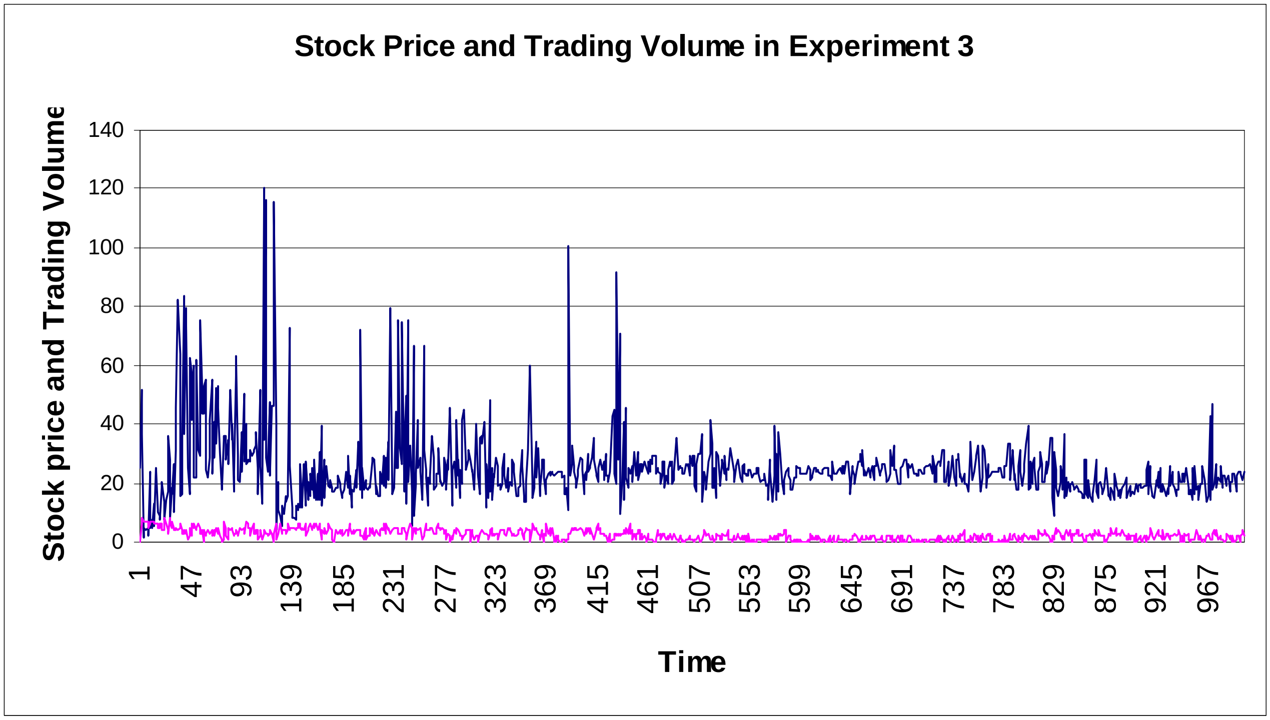

Figure 6. Stock Price and Trading Volume in Experiment 3

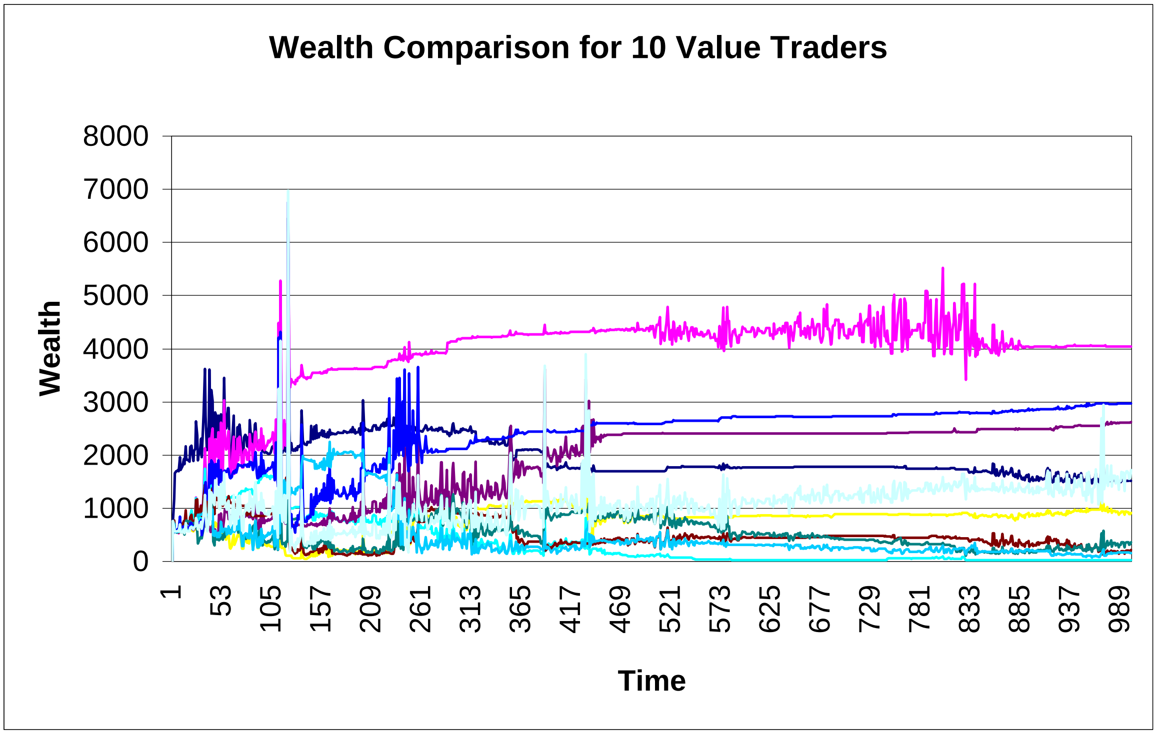

Figure 7. Accumulated Wealth For 10 Value Traders

Figure 8. Accumulated Wealth of 5 MA(5) Traders

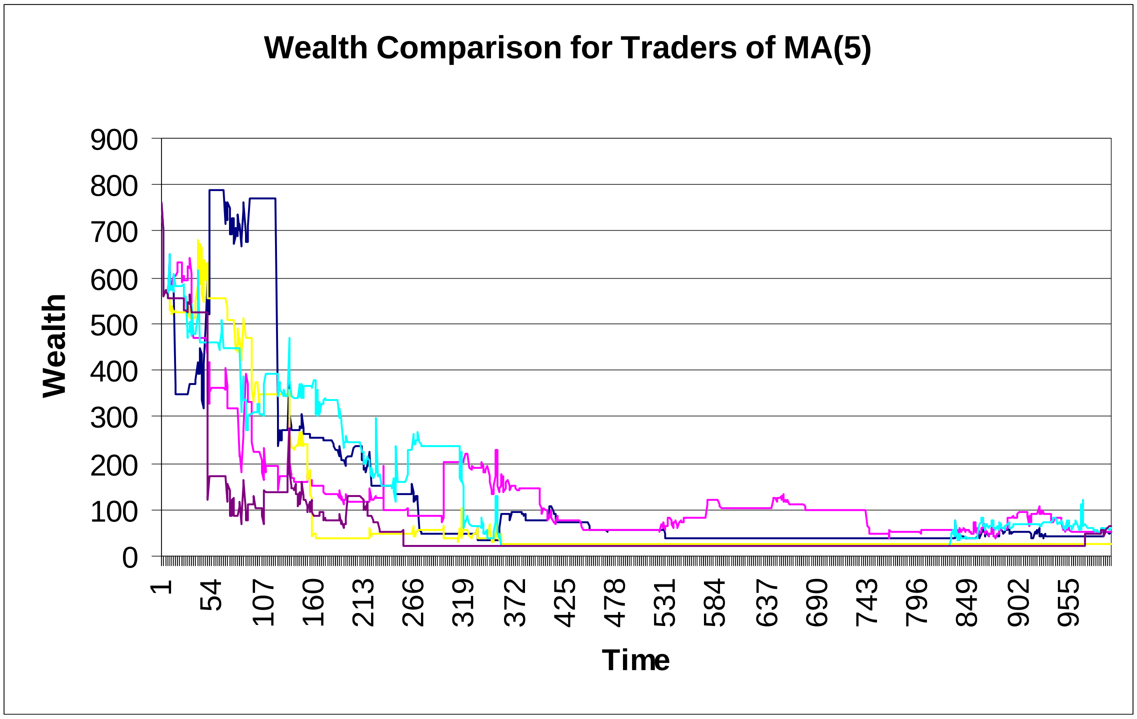

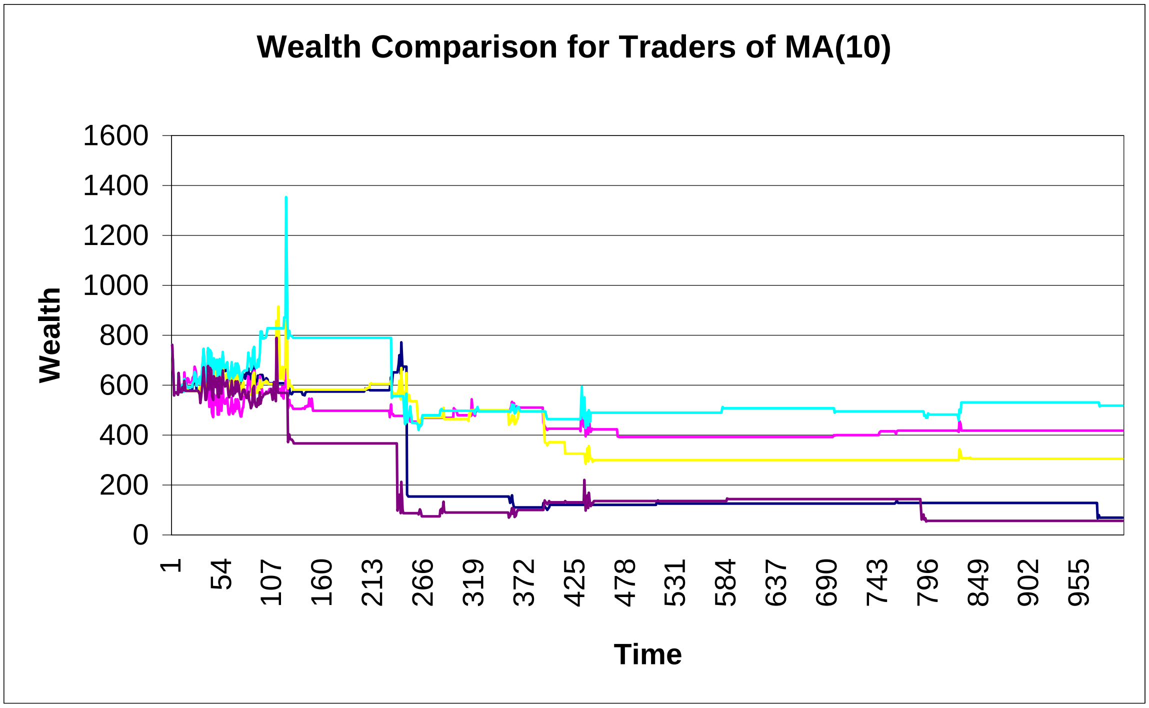

Figure 9. Accumulated Wealth for Momentum Traders of MA(10)

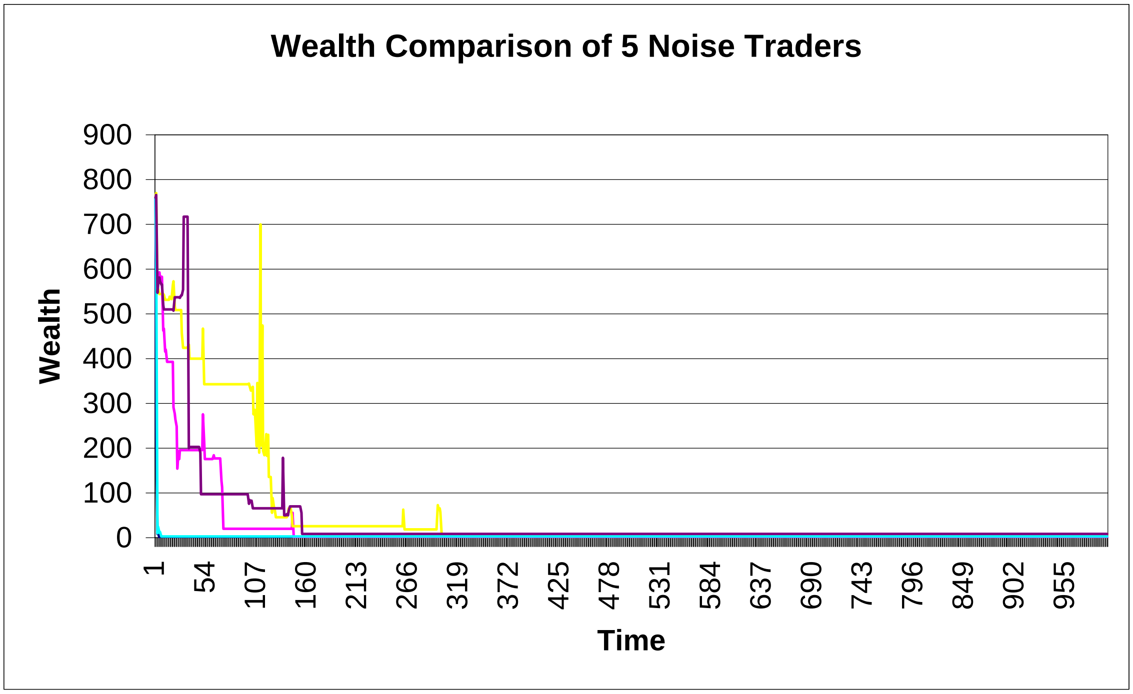

Figure 10. Accumulated wealth for 5 Noise Traders

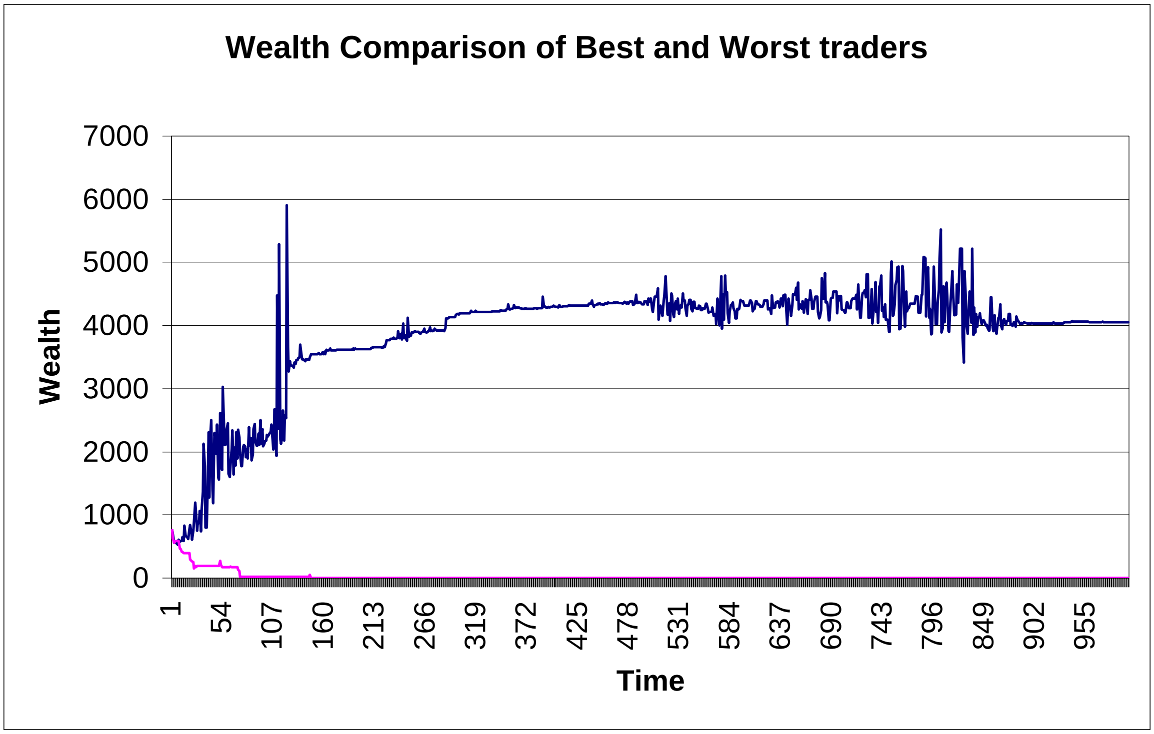

Figure 11. Wealth Comparison of the Best and the worst Traders