working papers

UNIVERSIDAD CARLOS III DE MADRID

PRICE DYNAMICS, INFORMATIONAL EFFICIENCY AND WEALTH DISTRIBUTION IN CONTINUOUS DOUBLE AUCTION MARKETS [1]

Javier Gil-Bazo*, David Moreno** and Mikel Tapia***

Abstract

This paper studies the properties of the continuous double auction trading mechanishm using an artificial market populated by heterogeneous computational agents. In particular, we investigate how changes in the population of traders and in market microstructure characteristics affect price dynamics, information dissemination and distribution of wealth across agents. In our computer simulated market only a small fraction of the population observe the risky asset's fundamental value with noise, while the rest of agents try to forecast the asset's price from past transaction data. In contrast to other artificial markets, we assume that the risky asset pays no dividend, so agents cannot learn from past transaction prices and subsequent dividend payments. We find that private information can effectively disseminate in the market unless market regulation prevents informed investors from short selling or borrowing the asset, and these investors do not constitute a critical mass. In such case, not only are markets less efficient informationally, but may even experience crashes and bubbles. Finally, increased informational efficiency has a negative impact on informed agents' trading profits and a positive impact on artificial intelligent agents' profits.

Keywords: Artificial financial markets, information dissemination, artificial neural networks, heterogeneous agents.

JEL Classifications: G12, G14, C63.

*Universidad Carlos III. Dpt. of Business Administration. C/Madrid, 126. 28903 Getafe (Madrid) Spain. javier.gil.bazo@uc3m.es
** Universidad Carlos III.Dpt. of Business Administration. C/Madrid, 126. 28903 Getafe (Madrid) Spain jdmoreno@emp.uc3m.es
*** Universidad Carlos III. Dpt. of Business Administration. C/Madrid, 126. 28903 Getafe (Madrid) Spain mtapia@emp.uc3m.es

1 Introduction

The continuous double auction trading system, where any agent can submit orders to buy or sell to an automatic limit order book, is becoming the standard trading mechanism in financial markets throughout the world. Many exchanges that formerly worked as pure price-driven markets are moving to order-driven or double auction markets. Examples of this transition are Nasdaq's Supermontage, NYSE's Hybrid Market[SM], or London Stock Exchange's electronic order-driven system, which now coexist with traditional quote-driven trading systems. However, despite the popularity of electronic double auction financial markets, the superiority of this trading system over other mechanisms in terms of price efficiency or higher level of liquidity remain open questions.

The financial literature on market microstructure attempts to explain the process of price formation and how this process is affected by the market structure and design.[1] While the traditional asset pricing literature is based on the ideal price setting concepts of Walrasian auction and Rational Expectations Equilibrium, the market microstructure literature recognizes that price discovery and information dissemination are complex processes that take place through a sequence of many bilateral trades potentially affected by frictions and the conditions of the market mechanism itself.[2] Contributions to this literature have mainly been theoretical (based on simplified economic models) or empirical. More recently, however, researchers have relied on experiments conducted in laboratories to explore price formation under realistic trading mechanisms, either with human agents (see, for instance, Plott and Sunder, 1982, and Forsythe et al., 1982) or employing computational agents.[3]

In this paper, we use a computer simulated market populated by boundedly rational agents to investigate information dissemination from informed to uninformed traders in a continuous

---

[1] O'Hara (1995) or Madhavan (2000) are two excellent surveys of this literature.
[2] In a Walrasian auction, individual demand and supply functions are aggregated so the price that clears the market summarizes investors' valuations of the asset. In a rational expectations equilibrium, on the other hand, traders rationally update their beliefs upon observing the Walrasian equilibrium price, so prices become fully efficient in that they perfectly reflect all information about the asset.
[3] See Lebaron (2000, 2001) for a detailed survey of agent-based financial markets.

double auction market and how market characteristics affect this process. In order to achieve this goal, we construct a market where only a small fraction of the population observe the risky asset's fundamental value with noise, while the rest of agents try to forecast-in different ways-the asset's price from past transaction data.

Our paper is closely related to others, such as Gode and Sunder (1993), who construct a simple double auction market where investors submit random orders, and show that the market price converges to the equilibrium price as long as traders are not allowed to buy or sell at a loss.[4] In a more recent study, Chan et al. (2001) study information dissemination in a market populated by boundedly rational investors with heterogeneous trading strategies, which include technical analysis rules and Bayesian learning. They find prices converge fast to the RE equilibrium price when investors have homogeneous preferences, but the model fails if this condition does not hold. Our market departs from Chan et al. (2001) in that the risky asset pays no dividend so agents cannot learn from past transaction prices and subsequent dividend payments. Consequently, convergence to the asset's fundamental price can only be attributed to private information being disseminated in the market. This distinction is important because Yang (2002) has shown that uninformed agents with no private information but endowed with neural learning capabilities can bring market prices to the RE equilibrium price.

Our artificial market captures well the main stylized facts of real financial markets, such as nonnormality of returns, skewness, kurtosis and volatility clustering. Moreover, by changing different market characteristics we are able to explore how the population of agents or market microstructure features such as portfolio constraints, explicit transaction costs or tick size can affect the dynamics and efficiency of market prices as well as distribution of wealth across investors.

The paper's main conclusion is that private information can disseminate effectively in a continuous double auction market populated by heterogeneous boundedly rational investors. We find that, that convergence to the fundamental path is possible as long as informed investors'

---

[4] In Gode and Sunder (1993) unaccepted limit orders are deleted with each transaction, unlike in a real market where limit orders stay in the book until matched by subsequent market orders or canceled by the trader who submitted them.

trades constitute a critical mass capable of correcting price departures from the fundamental path. Otherwise, the market price may diverge from the asset's fundamental value, creating crashes or bubbles. A bubble, in particular, arises in our market when the price moves above the fundamental path and continues to rise despite informed investors' initial orders to sell the asset. If informed investors could short sell the asset, such orders would eventually bring prices down, however, in the presence of short-selling constraints informed investors cannot submit new asks to the limit order book as long as their budget constraint becomes binding. Only the randomness of market dynamics can then make the asset price return to the fundamental path and bring it once again under the discipline of informed traders.

Market characteristics have significant effects on price dynamics, market efficiency and distribution of wealth. Interestingly, as the market becomes more efficient due to informed agents' trading, the value of their private information decreases, and so does their ability to obtain profits at the expense of other investors. Conversely, agents endowed with artificial intelligence, learn the fundamental process better in more efficient markets, and therefore make a greater profit.

The rest of the paper is organized as follows: section 2 presents the features of our artificial market; section 3 describes the design of simulations; section 4 explains the results; and, finally, section 5 concludes.

2 A continuous double-auction artificial market

In this section, we describe in detail the features of our artificial market structure: Trading mechanism, the risky asset's fundamental price, and the characteristics of agents participating in the market.

2.1 Market structure

In this artificial market, the agents' portfolio may contain two assets: a riskless asset with zero net return (numeraire) and a risky asset. For the risky asset, we model a trading mechanism with

no market maker, where orders are submitted by any investor and matched automatically. More specifically, our market is similar to the market presented in Chan et al. (2001).

Each market simulation consists of 1,000 trading periods (one trading period represents one day), and there are 5 rounds of trading in each trading period, which implies that each agent is granted the opportunity to trade a maximum of 5 times. Participating agents are sorted randomly[5] and the event of each agent participating in a given round is random with probability 0.7.[6] When an agent's turn arrives, he may submit a limit order, i.e., he may quote a price to buy (bid) or sell (ask) one unit of the risky asset. Alternatively, the agent may introduce a market order, which means that he may buy one unit of the risky asset at the best ask price or sell at the best bid price. As in Chan et al. (2001) or Chiarella and Iori (2002) we restrict the order size to one unit of the risky asset in order to limit the dimension of the investor's problem.

In addition, the agent must satisfy two different budget constraints at the time of submitting an order. First, he cannot submit market orders which imply short-selling the risky asset or borrowing. Second, he cannot submit a limit order such that if all his limit orders were matched at some point in the future, he would have to short sell or borrow. Moreover, every time agents go to the market they can revoke suboptimal limit orders introduce in the book in the past. These orders are defined as those that would result in a loss given the agent's new beliefs about the risky asset's price. Finally, agents' orders to buy (sell) are rounded downwards (upwards) in order to comply with the exchange's tick size requirements. We have set the tick size equal to 0.01 dollars for the majority of the experiments but we have changed it to 10^{-5} dollars in some experiments to be able to analyze the effects of this variable on the market dynamics.

2.2 Fundamental price

Our market departs from the standard approach in the agent-based market literature in that the risky asset pays no dividend at the end of each trading period. Instead, the risky asset's

---

[5] A permutation of agents is randomly sampled from all possible permutations according to a uniform discrete distribution.
[6] It is based on the idea that, in real financial markets, agents do not trade everytime: even active agents usually trade a few times in a trading session.

fundamental price process is exogenously determined. This assumption serves four different purposes. First, potential convergence to the asset's fundamental price can only be attributed to private information being disseminated in the market, and not to uninformed traders learning from the observed dividend process as in Yang (2002). In this way, the degree of informational efficiency becomes a good indicator of the extent of private information dissemination through prices. Second, the requirement that the asset pays dividends on a regular basis may be unrealistic for some assets such as growth stocks or zero-coupon bonds. Third, the exogenous fundamental process is directly comparable to the path of transaction prices. Finally, and perhaps more importantly, this assumption enables us to model asymmetric information in the most simple way and focus exclusively on differences in expectations about future prices across agents. In particular, while some investors observe a noisy signal of the fundamental price, others only observe the history of past transaction prices and form their expectations accordingly.

We asume that the risky asset's fundamental price dynamics is governed by a geometric Brownian motion:

$$dP_{f,t} = \mu P_{f,t} dt + \sigma P_{f,t} dW_t$$

where $P_{f,t}$ denotes the risky asset's fundamental price at time t and W_t is a standard Brownian motion. Under this process, the fundamental price is conditionally lognormally distributed:

$$\ln P_{f,t+\Delta} \big| P_{f,t} \sim \phi\left(\ln P_{f,t} + \left(\mu - \frac{\sigma^2}{2} \right) \Delta, \sigma \sqrt{\Delta} \right) \qquad (1)$$

where Δ is the time interval between two trading periods (we set $\Delta = 250^{-1}$) and ϕ denotes the normal density function. The continuously compounded fundamental return between any two periods, computed as the difference in logs of prices, is normally distributed, serially independent and has constant variance. Consequently, possible rejection of normality or serial independence in real transaction return series can only be caused by the trading process.

2.3 Agents and trading

All agents are risk neutral and myopic. Their objective function is therefore the expected value of their wealth at the end of each trading period. The myopic agent assumption is quite common in computer simulated markets, since it enables the researcher to abstract from investment horizon effects. Risk-neutrality, on the other hand, is necessary in our analysis, because the fundamental price is exogenously determined, and transaction prices would never converge to the fundamental price if risk-averse investors required a risk premium to buy the asset. Moreover, under this assumption, investors only have an incentive to trade when their expectation about the price level diverges from the current market price.

Given their information set and using a previously determined algorithm, the i-th agent forecasts the risky asset's price at the end of the current trading period, P_t^i, and buys as long as the best ask price is at least S dollars below his forecast. Similarly, he sells as long as the best bid price is at least S dollars above his forecast. This assumption permits us to incorporate explicit transaction costs such as broker commissions, other fees or taxes on trading: investors will trade only if the expected profit from trading is high enough to offset the cost of trading. Note, however, that a higher value of S will decrease the volume/frequency of trading. In our simulations, the exact value of S for each agent at each round is a realization from a uniform distribution.[7]

The agent can also submit a new bid (P_t^i-S) or ask (P_t^i+S). The agent's limit order is subsequently added to the limit order book in the corresponding position. Table 1 displays the agents' decision process.

Agents differ from each other in the way they obtain their forecast of the risky asset's price. In our market we consider four different classes of agents:

i) Zero Intelligence (ZI) agents. In our market, ZI agents serve as liquidity providers, since they introduce discrepancy in the risky asset's valuation. Their prediction for the risky asset's price is

---

[7] It try to reflects the fact that in real financial markets there are a huge heterogeneity in the transaction cost among different agents (i. e. between private investors and institutional investors).

a realization of a random variable which is uniformly distributed around the last transaction price:

$$P_t^i \sim U(0.9P_\tau, 1.1P_\tau)$$

where τ denotes the last time the risky asset was traded. This way of modeling ZI agents' behavior departs from that of Gode and Sunder (1993), since in our market ZI agents modify their prediction with every transaction. Consequently, if the market were exclusively populated by ZI agents, we would not expect convergence to the "equilibrium price" since a high (low) transaction price would shift expectations for all agents upwards (downwards). Moreover, the presence of ZI agents enables us to assess the effect of informed-based trading on uninformed traders' wealth.

ii) Informed agents. They observe the risky asset's fundamental price plus some noise. Their forecast of the asset's price is therefore:

$$P_t^i \sim U(0.99P_{f,t}, 1.01P_{f,t})$$

We could think of informed agents as fundamentalist traders (more noise) or as insider traders (less noise). It should be noted that only to the extent that informed traders' orders can drive market prices towards the risky asset's fundamental value, prices have the potential to become informationally efficient.

iii) ANN agents. These investors use a nonlinear model (Artificial Neural Networks (ANN)) to forecast next period's return. Artificial Neural Networks have been usually considered as complex models inspired in the structure of the brain. ANNs inherit three basic characteristics of the biological neurons: they are intrinsically parallel; they provide nonlinear responses to stimulus; and they process the information through several layers of interconnected neurons. One of the main characteristics of ANNs is their capacity to "learn" and "generalize" using real data, that is, an ANN learns the relationship between a set of inputs and their corresponding outputs, and for this reason they have been widely used in several financial problems, especially in forecasting issues on different financial assets (see, Hansen et al. 1999 among others). An

ANN is formed by a number of *processing units* which are hierarchically organized in *layers*. The input layer consists of a set of nodes that receive the information from the outside world. The *hidden layer* processes the information while the output layer sends the signal to the outside.[8] The most widely used structure is that of a *feedforward neural net* in which the information is hierarchically processed in a single way from the input layer to the output through the hidden layer(s). The units are connected through a synaptic weight which determines quantitatively the influence of one unit on the other. The ANN agents in our market use a feedforward neural net with a unique hidden layer[9] with "g" units, "h" units in the input layer, and a unique unit in the output layer to form their expectation about next period's return. It is important to highlight that the net is not fixed during a simulation, but trained each "m" trading periods More specifically, we select the parameters, $g = 8$, $h = 10$, $m = 50$, and the number of epochs used in the training of nets is 200. We employ the log-sigmoid transfer function in the hidden layer and the linear transfer function in the output layer, as in most common feedforward ANNs. The training algorithm used is the batch Levenberg-Marquardt backpropagation algorithm (see Hagan and Menhaj, 1994), which like the quasi-Newton methods was designated to approach second-order training speed without having to compute the Hessian matrix. We use this training algorithm because it is one of the fastest methods for feed-forward neural networks and it is also very efficiently implemented in the software used for this study, MATLAB. It must be noted that an interesting characteristic of our market is that as the market becomes more efficient, the ANN agents are capable of training the net in a more efficient way and therefore learn the true fundamental process (geometric Brownian motion). In turn, as ANN agents learn the fundamental process, their orders can potentially contribute to making the market more efficient.

iv) Technical Analysts (TA). These agents attempt to identify trends in price series. In particular, they compute the average close price of last "s" trading periods and the average close price of

---

[8] Kuan and White (1994) provide an introduction to ANNs in an econometric context, showing that these models are, in fact, quite familiar to the econometrician, see it for a more detailed description.
[9] Hornik et al. (1989) showed that an ANN with a single hidden layer with enough hidden logistic units and linear outputs can approximate arbitrarily well any measurable function.

last "l" periods, where "s" is either 1, 2 or 5 periods with equal probability and "l" can take the following values: 50, 150 or 200, also with equal probability.[10] If the short moving average is higher than the long moving average the TA will think the market is bullish and will want to buy the risky asset. Otherwise, he will sell.[11] TA's actions therefore differ slightly from those of the rest of agents. In particular, in a bullish market, a TA will buy at the prevailing ask price, a, if there is at least one ask price with $P_t^i > a + S$, where P_t^i is taken to be price of the last transaction. Otherwise, they will submit a limit order to buy with bid equals to P_t^i-S. Analogously, in a bearish market, TAs will sell at the prevailing bid price, b, if there is at least one bid price with $P_t^i < b - S$ and will submit a limit order to sell with ask equal to P_t^i+S, otherwise.

At the beginning of each simulation, each agent receives an identical endowment consisting of 3,000 dollars and 30 units of the risky asset. Since TAs and ANN agents require a minimum number of transactions before they start to trade, wealth is reinitialized again when all agents operate simultaneously. For this reason, in the first 200 periods only informed agents and ZI agents participate and we do not take into account these results in our further analysis.

3 Simulation design

In order to analyze the sensitivity of market's efficiency to market parameters we have simulated 10 paths of fundamental values according to (1) with a different μ for each path (uniformly distributed between +/-5%), and $P_{f,0} = 100$. The reason why we let μ change is to ensure that results do not depend on a specific market trend. Next, for each fundamental path, we have simulated a total of 1,000 trading periods, changing one of the market parameters at a time.[12] More specifically, we have considered the parameter sets shown in Table 2.

---

[10] These parameters for the long and short mean moving averages have been taken from Brock et al. (1992).

[11] In future extensions, TA agents might use some kind of evolutionary algorithm (as a genetic algorithm) to select among competing forecast rules as in Arthur et al. (1997) or Lettau (1997).

[12] Since each trading period contains 5 trading rounds, each simulation consists of 5,000 trading rounds.

To study convergence of the price dynamics to the fundamental path, we compute the following

measures proposed by Theissen (2000):

a) Mean Absolute Error

It is the average of absolute deviations of the transaction close price from the fundamental price:

$$MAE = \frac{1}{T}\sum_{t=1}^{T}\left|P_t - P_t^f\right|$$

b) Mean Relative Error

It enables comparisons across different paths.

$$MRE = \frac{1}{T}\sum_{t=1}^{T}\frac{\left|P_t - P_t^f\right|}{P_t^f}$$

c) Root Mean Squared Error

This measure weights extreme relative errors more heavily:

$$RMSE = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\frac{P_t - P_t^f}{P_t^f}\right)^2}$$

4 Results

In this section, we perform a sensitivity analysis based on controlled experiments where the

value of one parameter of the market is changed at a time. In particular, we study the effects of

changes in the population of agents and market microstructure on price dynamics, informational

efficiency and distribution of wealth across traders.

4.1 The benchmark market

Table 3 displays descriptive statistics of end-of-period returns averaged across all simulations

for each market. In our benchmark market, the table shows that trading itself (or microstructure

noise) accounts for more than half the standard deviation of market returns: 25.13%, as opposed

to 10% in the fundamental series. Moreover, Value at Risk (VaR), i.e., the maximum relative

decline in the risky asset's price between two consecutive periods is 4 times higher for actual

market returns than for the fundamental path.

Table 3 also shows that market returns in the benchmark market are characterized by excess kurtosis and negative skewness.[13] Moreover, the null hypothesis of normal returns is rejected in 6 of the 10 simulated series at the 1% significance level, according to the Jarque-Bera normality test.

Another typical characteristic of financial series is volatility clustering: periods of high volatility tend to be followed by periods of high volatility and viceversa. Autoregressive Conditional Heteroscedasticity (ARCH) captures such effects. In order to test for the presence of ARCH effects, we implement Engle's (1982) test the null hypothesis of no serial correlation in conditional volatility. Table 4 displays the number of simulated series in each market for which the null hypothesis is rejected. For the benchmark case, as many as 4 of all the series display ARCH effects in the disturbance term of the return process.[14] Volatility clustering has also been documented in other artificial markets (Lux, 1995, and LeBaron, 1999, among others). To summarize, our benchmark market therefore captures stylized facts of real financial market price dynamics.

Next, we analyze informational efficiency in the benchmark market. As Table 5 shows, mean relative error equals about 3% on average, a small figure relative to the standard deviation of market returns. In order to show the divergence between actual market returns and the fundamental path, we have chosen a single simulation and displayed fundamental and market prices and returns on Figure 1. The Figure also includes the series of transactions per period. We conclude that private information disseminates efficiently in the benchmark market, which gives uninformed agents endowed with artificial intelligence the opportunity to learn from transaction prices. This becomes evident in Table 6, where we display the average investor's return by agent class between the first and last period. While the average informed agent exploits his informational advantage to the extent of obtaining a total portfolio return of 31%, the average ANN agent, who enjoys no private information, makes a profit of 14.1%. Even TA

---

[13] These resutls have also been found in some other artificial markets (LeBaron, 1999 among others)
[14] We have previously fitted an autoregressive process to conditional mean returns.

can exploit market trends to make a small positive return of 1.56%. Such profits are made at the expense of ZI agents, who lose 17.19% of their initial wealth on average.

4.2 Effect of changes in agent population

When the number of informed agents is halved (Market 1), price dynamics depart more noticeably from the fundamental distribution as can be seen from Table 3: volatility increases from 25.13% in the benchmark market to 65.87%, negative skewness, kurtosis, and Value at Risk also rise significantly. Normality is rejected in all simulations. This is accompanied by less liquidity as measured by the average number of transactions period (last column of Table 3). Also, Table 4 suggests that ARCH effects are always present in series of returns generated in this market. Finally, as shown in Table 5 price efficiency drops dramatically in Market 1 with respect to the benchmark market, with mean relative errors close to 20%.

To understand these results, we consider a single fundamental path and analyze price dynamics in the benchmark market and Market 1. Figure 2 shows the evolution of fundamental and transaction prices and returns for a single path in Market 1. This Figure is in contrast to Figure 1 that corresponds to the same single path under the benchmark market. First, the high volatility of returns is apparent from the bottom panel (returns) of Figure 2 in all periods. This suggests that informed investors not only ensure convergence to price fundamentals, but also reduce excess volatility in transaction prices. This result suggests that turbulence in actual financial markets may actually be the consequence of less precise aggregate information in the market about fundamentals.

Figure 2 also reveals the presence of an interesting phenomenon. While transaction prices in Market 1 tend to track the fundamental path in most periods, from period 200 to period 350, the market experiences a crash. This crash starts with a sharp decline in the asset's price followed by a relative stability, and finally, a quick return to the fundamental path that is not abandoned again. Interestingly, the medium panel shows that trade volume peaks when prices fall and rise, and remains relatively low in between.

In order to understand the crash in Market 1, we look at the evolution of agents' average positions per investor class, which are displayed in Figure 3. The Figure clearly shows that as Technical Analysts sell their positions, possibly following a minor downward trend, the price falls below its fundamental value, which further reinforces their belief that the market has started to decline. Informed investors then increase their positions in the asset. By doing so, however, they run out of cash, and since borrowing is not permitted, they cannot submit new bids to the system, so the decline in transaction price continues. At some point, because TA are also constrained by short selling restrictions, they cannot continue selling and prices end up stabilizing. Towards period 300, only ANN agents and ZI agents can trade. As prices eventually rise above their recent mean, TA start to buy from ANN and ZI agents, and prices start to rise again. Around period 350, prices have already reached the fundamental path and keep rising, so informed traders intervene again by selling the asset. This time, however, their orders manage to bring prices down to the fundamental path. Informed agents must therefore create a critical mass for information to disseminate in the market. The crash therefore teaches us that informed agents must constitute a critical mass to disseminate their private information throughout the market and correct price movements. The size of their critical mass is a function of both their number and their aggregate wealth. When they are not enough in number and have constrained budgets, herding TA can drive market prices away from fundamentals.

Also, it is interesting to note that informed agents in this inefficient market enjoy a greater informational advantage that, as Table 6 shows, enables them to make an average portfolio return of 84.73%, mostly at the expense of TA. Clearly, less information dissemination also means that private information is more valuable. In informationally efficient markets with fully revealing prices, private information would have no value, so it would not be profitable to collect it, as postulated by Grossmann and Stiglitz (1980) in a Rational Expectations Equilibrium framework. Another consequence of price inefficiency is that ANN agents learn less about the fundamental process, and therefore make lower profits.

In Markets 2 and 3, we investigate the effect of a reduction in the number of ANN agents and TA, respectively. Table 3 suggests that reducing the number of ANN agents has an effect on

market dynamics similar to that of reducing the number of informed agents: Volatility, Value at Risk, negative skewness and kurtosis increase with respect to the benchmark. Also, ARCH effects are present in 8 of the 10 simulated series as can been seen in Table 4. Eliminating TA from the market, however, affects market dynamics only slightly although VaR, kurtosis, volatility clustering and nonnormality reduce marginally, and skewness becomes positive. These results imply that TA contribute to nonnormality in returns while ANN agents induce the opposite effect. The presence of both, however, reduces the relative weight of informed agents, and therefore their ability to influence prices as evidenced by Table 5: When ANN agents or TA are eliminated, Mean Relative Errors reduce to only 2% and 1%, respectively. Large sustained deviations of transaction prices from the fundamental process, however, are still possible. Figure 4 shows market dynamics for a selected simulation under Market 2 where ANN agents do not operate. The top panel indicates that around the 800th period, prices start to increase very fast above the risky asset's fundamental value to reach a peak and decrease to the fundamental path before the 840th period. Volume increases both at the beginning and the end of the bubble and stays lower than average in the meantime. Figure 5 sheds light on the reason behind such anomalous behavior: the bubble starts when TA detect an upward trend in prices and start to buy the asset. Informed agents then react by selling the asset until they have no units of the asset left. A few periods later, TA run out of cash and the bubble starts to burst. The situation is therefore very similar but inverse to that of Figures 2 and 3. Again, the combination of a limited budget for informed traders, together with short-selling constraints and the presence of TA can trigger a major mispricing episode.

As for wealth distribution, Table 6 indicates that while ZI agents' returns are not sensitive to the presence of ANN agents and TA, informed agents make a higher return (47.5% as opposed to 31%), especially when there are no ANN agents, possibly due to decreased competition for ZI agents' money.

4.3 Effect of Portfolio Constraints

Since price efficiency appears to be limited by the presence of restrictions on short-selling and borrowing together with a binding budget constraint for informed traders, it is an interesting exercise to investigate how the market behaves when such constraint is never binding, as is the case in Market 4. As we can see in Table 3, Market 4 is characterized by lower volatility and marginally lower Value at Risk than the benchmark market, although returns display slightly more kurtosis and negative skewness. Absence of ARCH effects is also rejected more frequently than in the benchmark case as Table 4 suggests.

Significant differences become evident, however, when we analyze market efficiency. Table 5 shows that mean relative errors in this market are below 1% on average. Clearly, when informed investors are not constrained by their budget, their ability to influence prices is much stronger and the probability of bubbles and crashes becomes very low. The idea that short sale or borrowing constraints can prevent information or opinions from being reflected in stock prices is not new. Miller (1977) or Diamond and Verrechia (1987) models express this point formally. Our paper shows that the same perverse effect of short selling and borrowing constraints is present in more complex markets populated by heterogenous boundedly rational investors.

4.4 Effect of transaction costs

Several papers (Keim and Madhavan, 1997; Barclay, et al., 1999 or Domowitz, et al., 2001 among others) in the financial literature postulate that trading cost represents a relevant variable in the trading process, and they also play a very prominent role from a practitioner's perspective. Keim and Madhavan (1997) affirm that transaction cost are central in determining investment performance and can substantially reduce the expected value created by an investment strategy. Lesmond et al. (1999) points out that transaction cost estimates are not always available, or where available, are cumbersome to use and expensive to purchase. They

propose a new way to estimate them and find that total transaction cost (explicit plus implicit) range from 10.3% for small firms and 1.2% for the large firms.

In a more recent paper, Domowit et al. (2001) analyze the interactions between cost, liquidity, and volatility, and their determinants, using panel-data for 42 countries from September 1996 to December 1998. They offer empirical evidence of the huge variability of trading cost across countries (emerging and developed ones) and that the composition of global efficient portfolios could change severely when cost and turnover are taken into account.

The trading costs are usually decomposed into two components: explicit costs and implicit costs. Explicit costs are the direct costs of trading, such as broker commission costs, or even taxes, etc. Implicit costs represent indirect trading cost (the more important is the price impact of the trade). According to Domowitz et al. (2001) in general, explicit costs are the major element in the total trading costs (they represents roughly two-thirds of total cost). Therefore, in this section, we generate two new markets (market 5 and 6) to analyze the effects of the explicit transaction cost, on the previous results.

We achieve this goal by increasing the upper bound of S to 2.5 and 3.5 in markets 5 and 6, respectively. According to Table 3, in market 5 the volatility in returns and, most notably, VaR decreases with respect to the benchmark market. Moreover, normality in this market is only rejected in 30% of the simulations. This outcome is confirmed by the results found for the market 6 (Table 3). In this case we can observe that the volatility decreases to 19.8%, the VaR is reduced to half of the VaR found in the Benchmark market, and the kurtosis value is very close to the kurtosis in a normal distribution.

Although not shown in the paper for the sake of brevity, our analysis of the limit order book suggests that higher explicit transaction costs make trading less likely and increases the depth of the limit order book. The natural consequence is a reduction in the sensitivity of market prices to minor shifts in investors' expectations, and hence in volatility and volume. We can observe this fact in the last column of Table 3 (volume decreases in 16% and 26% in markets 5 and 6, respectively).

In Table 4 we also analyze the probability of achieving volatility clustering in these two markets. We observe the hypothesis of homoscedasticity is rejected only in 20% of the simulations.

Perhaps more interesting are the results reported in Table 5, where we can observe that there is a positive relationship between the level of transaction cost and market efficiency. Specifically, in the case of the market 6, the value of the RMSE is almost half of the RMSE in the benchmark cases. This interesting result suggests that explicit transaction costs are a potentially powerful tool for market designers to increase market efficiency.

4.5 Effect of Tick Size

Finally, we study how a specific feature of the market microstructure, the tick size, affects price dynamics, and agents' portfolio profits. Harris (1994) shows that tick size plays an important role in determining the quality of double auction trading system. Tick size determines the minimum spread that can be quoted and as a consequence the distribution of volume in different price levels. The previous evidence related with trading activity shows some interesting results. First, Harris (1994) is the only paper finding a positive relationship between volume or trading activity and reduction of tick size. The rest of the papers find no significant evidence.[15]

The relationship between tick size and volatility is more conclusive. Ronen and Weaver (2001) and Bessembinder (2000) find that tick level is associated with intra and interday volatility for different markets.

In Tables 3 and 5 we can see that if the tick size (minimum price increment allowed) is reduced to 10^{-5} (market 7) from $0.01 (benchmark market) the effects on the basic characteristics of the returns distribution and the efficiency are slightly different. Thus, the volatility and VaR increase slightly, and so does the volume. Moreover, we find that the change in tick size does

---

[15] See Lau and McInish (1995), Anh, et al. (1998), Ronen and Weaver (1998) and Bourghelle and Declerck (2004). The argument to justify these result is that relationship between activity and tick size cannot be isolated from other variables.

not affect significantly the market's efficiency. However, in Table 4 we observe that the probability of finding volatility clustering increases in market 7 compared to the benchmark. Finally, Table 6 shows that informed agents are able to improve their performance in a market with lower tick size, because they can introduce more accurate prices in the system to increase their profits. The wealth achieved by the rest of the agents is not changed significantly.

5 Summary and conclusions

Experiments with humans and with computational agents have demonstrated that that private information can disseminate well in continuous double auction markets. The exact mechanism through which such transfer of information is possible or the limits to information dissemination deserve a deeper look. This paper attempts to shed light on these issues which have important implications for both academics and market structure designers. In this paper, we investigate further the way information disseminates from informed to uninformed traders in a market populated by heterogeneous boundedly rational agents. In order to achieve our goal, we construct a computer simulated market where only a small fraction of the population observe the risky asset's fundamental value with noise, while the rest of agents try to forecast the asset's price from past transaction data.

We find that the more relevant stylized facts in financial market can be obtained in an artificial financial market even when investors are risk-neutral, myopic, and do not behave strategically. Another interesting finding is that if informed traders constitute a critical mass of population, their private information can be properly disseminated through the market. However, when they do not represent a critical mass, then irrational dynamics such as bubble or crashes can emerge. This impossibility to drive transaction prices to the risky asset's fundamental value seems to decrease when borrowing and short-selling restrictions are eliminated. In this case, informed agents' ability to influence prices is much stronger and the probability of bubbles and crashes becomes almost zero. However, as the market becomes more efficient the value of private information decreases, and so does the profit of informed agents. Interestingly, agents endowed

with artificial intelligence obtain higher trading profits in more efficient markets given that transaction prices contain more precise information about the true fundamental process that agents are able to learn.

Changes in the fraction of technical analysts and Artificial Neural Network agents alter market dynamics but not the main conclusions regarding the ability of private information to disseminate in a double auction market. In addition, we also analyze the effect of changing some specific features of the market microstructure: the explicit transaction cost and tick size.

We find that there is a positive relationship between the level of explicit transaction cost and market efficiency. This striking result suggests that explicit transaction costs could be a potentially powerful tool for market designers to attempt to increase market efficiency in some cases.

Finally, tick size in this market does not seem to affect significantly price dynamics or market efficiency. However, we find that the informed (or insiders) agents are able to improve their performance in a market with lower tick size, because they can introduce more accurate prices in the system and profit more from their informational advantage.

Bibliography

Ahn H-J., C Cao and H. Choe. 1998. Decimalization and Competition among Stock Markets: Evidence from the Toronto Stock Exchange Cross-Listed, Journal of Financial Markets, 1, 51-87.

Arthur W.B., J.H. Holland, B. LeBaron, R.G. Palmer, and P. Tayler. 1997. Asset Pricing under Endogenous Expectations in an Artificial Stock Market. Economics Notes, 26(2), 297-330.

Barclay, M.J., W.G. Christie, J.H. Harris, E. Kandel, and P.H. Schultz. 1999. The Effects of Market Reform on the Trading Costs and Depths of Nasdaq Stocks. Journal of Finance, 54, 1-34.

Bessembinder H. 2000. Tick Size, Spreads and Liquidity: an Analysis of Nasdaq Securities Trading near Ten Dollars, Journal of Financial Intermediation, 9, 213-239.

Bourghelle D. and F. Declerck. 2004. Why markets should not necessarily reduce the tick size, Journal of Banking and Finance, 28 (2), 373-398

Brock, W. A., J. Lakonishok and B. LeBaron. 1992. Simple Technical Rules and the Stochastic Properties of Stock Returns. Journal of Finance, 24(1), 66-77.

Chan, N. T., B. D. LeBaron, T. Poggio, and A. W. Lo. 2001. Agent-Based Models of Financial Markets: A Comparison with Experimental Markets. *In* MIT Sloan Working Paper No. 4195-01.

Chiarella, C. and G. Iori. 2002. A Simulation Analysis of the Microstructure of Double Auction Markets. Quantitative Finance, 2, 346-353

Diamond, D. W. and R. E. Verrecchia. 1987. Constraints on Short-Selling and Asset Price Adjustment to Private Information. Journal of Financial Economics, 18(2), 277-312.

Domowitz, I., J. Glen and A. Madhavan, 2001. Liquidity, Volatility and Equity Trading Costs across Countries and Over Time. International Finance, 4(2), 221-255.

Engle, Robert F. 1982. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation. Econometrica, 50, 987-1007.

Forsythe, R., T. R. Palfrey and C. R. Plott. 1982. Asset Valuation in an Experimental Market. Econometrica, 50(3), 537-568.

Gode, D. K. and S. Sunder. 1993. Allocative Efficiency of Markets with Zero Intelligence Traders. Journal of Political Economy, 101, 119-137.

Grossman, S. J., and J. E. Stiglitz. 1980. On the Impossibility of Informationally Efficient Markets. American Economic Review, 70, 393-408.

Hagan, M. T., and M. Menhaj. 1994. Training Feedforward Networks with the Marquardt Algorithm. IEEE Transactions on Neural Networks, 5 (6), 989–993.

Hansen, J. V., J. B. McDonald and Ray D. Nelson. 1999. Time Series Prediction with Genetic-Algorithm Designed Neural Networks: An Empirical Comparison with Modern Statistical Models. Computational Intelligence 15 (3), 171-184

Harris L. 1994. Minimum Price Variations, Discrete Bid-Ask Spreads and Quotation Sizes, Review of Financial Studies, 7, 149-178.

Hornik, K., M. Stinchcombe and H. White. 1989. Multilayer Feedforward Networks Are Universal Approximators. Neural Networks, 2, 359-366

Keim, D. B. and A. Madhavan. 1997. Transactions Costs and Investment Style: An Interexchange Analysis of Institutional Equity Trades. Journal of Financial Economics, 46(3,Dec), 265-292.

Kuan, C.M. and H. White. 1994. Artificial Neural Networks: An Econometric Perspective. Econometric Reviews, 13, 1-91.

Lau S. and T. McInish. 1995. Reducing Tick Size on the Stock Exchange of Singapore, Pacific-Basin Financial Journal, 3, 485-496.

LeBaron, B. 2000. Agent-Based Computational Finance: Suggested Readings and Early Research. Journal of Economic Dynamics and Control, 24, 679-702.

LeBaron, B. 2001. A Builder's Guide to Agent-Based Financial Markets. Quantitative Finance, 1, 254-261.

LeBaron, B., W.B. Arthur, and R. Palmer. 1999. Time series properties of an artificial stock market. Journal of Economic Dynamics and Control, 23, 1487–1516.

Lesmond, D. A., J. P. Ogden and C. A. Trzcinka. 1999. A New Estimate of Transaction Costs. Review of Financial Studies, 12(5,Winter), 1113-1141.

Lettau, M. 1997. Explaining the Facts with Adaptive Agents: The Case of Mutual Fund Flows. Journal of Economic Dynamic and Control, 21, 1117-1147.

Lux, T. 1995. Herd behavior, bubbles and crashes. Economic Journal, 105 (431), 881–896.

Madhavan, A. 2000. Market microstructure: A survey. Journal of Financial Markets 3, 205-258.

Miller, E. M. 1977. Risk, Uncertainty, and Divergence of Opinion. Journal of Finance 32, 1151-1168.

O'Hara, M. 1995. Market Microstructure Theory (Cambridge, MA: Blackwell Publishers).

Plott, C. R. and S. Sunder. 1982. Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models. Journal of Political Economy, 90(4), 663-698.

Ronen T. and D. Weaver. 2001. Teenies Anyone? The Effect of Tick Size on Volatility, Trader Behavior, and Market Quality. Journal of Financial Markets, 4,231-60.

Theissen, Erik. 2000. Market Structure, Informational Efficiency and Liquidity: An Experimental Comparison of Auction and Dealer Markets. Journal of Financial Markets, 3(4,Nov), 333-363.

Yang, J. 2002. The Efficiency of an Artificial Double Auction Stock Market with Neural Learning Agents. En Evolutionary Computation in Economics and Finance, Edited by S.H. Chen, 79–98. Springer-Verlag.

Tables

Table 1. Order generation procedure.

Scenario	Order
Case 1: There is at least one ask price and one bid price in the limit order book*.	
$P_t^i > a + S$	Market order to buy
$P_t^i \leq b - S$	Market order to sell
$a + S > P_t^i > b - S$ & $\left\| a - P_t^i \right\| > \left\| P_t^i - b \right\|$	Limit order to sell with ask = $P_t^i + S$
$a + S > P_t^i > b - S$ & $\left\| a - P_t^i \right\| \leq \left\| P_t^i - b \right\|$	Limit order to buy with bid = $P_t^i - S$
Case 2: There are no bid prices	
$P_t^i > a + S$	Market order to buy
$P_t^i \leq a + S$	Limit order to buy with bid = $P_t^i - S$
Case 3: There are no ask prices	
$P_t^i < b - S$	Market order to sell
$P_t^i \geq b - S$	Limit order to sell with ask = $P_t^i + S$
Case 4: There are no ask or bid prices	
With probability ½	Limit order to buy with bid = $P_t^i - S$
With probability ½	Limit order to sell with ask = $P_t^i + S$

* The lowest ask prevailing in the limit order book when the agent takes his decision is denoted by a and the highest bid is denoted by b.

Table 2. Sensitivity Analysis.

Market	Number of Agents			Budget		Explicit Transaction Costs (S)	Ticks
	Infor-med	ANN	TA	Cash	Asset Units		
Benchmark	**40**	**10**	**40**	**3,000**	**30**	**U (0,2)**	0.01
Market 1	**20**	10	40	3,000	30	U (0,2)	0.01
Market 2	40	**0**	40	3,000	30	U (0,2)	0.01
Market 3	40	10	**0**	3,000	30	U (0,2)	0.01
Market 4	40	10	40	∞	∞	U (0,2)	0.01
Market 5	40	10	40	3,000	30	**U (0,2.5)**	0.01
Market 6	40	10	40	3,000	30	**U (0,3.5)**	0.01
Market 7	40	10	40	3,000	30	U (0,2)	**10^{-5}**

* The number of ZI agents is 120 in all markets and the standard deviation of the fundamental prices is fixed to 10% for all experiments.

Table 3. Descriptive Statistics for returns.

Return is the mean annualized return in percentage averaged across all simulations in a given market. *Std. Dev.* denotes the average mean annualized standard deviation in percentage. *Skw* indicates the mean skewness, *Kur* is kurtosis, *VaR* denotes the maximum price decrease with a 99% probability, *Reject. norm* shows the number of times (out of 10 simulations) that the null hypothesis of normality is rejected using the *Jarque-Bera* test at 1% significance level. The value of the Jarque-Bera statistic is shown in the next column, and the last column (*Volume*) denotes the average number of transactions in a market per trading period. We also include descriptive statistics for fundamental prices.

	Return (%)	Std. Dev. (%)	Skw	Kur	VaR	Reject. norm	Jarque-Bera Statistic	Volume
Fundamental	-0.83	9.97	0.016	2.9	2.02	0	--	--
Benchmark	-2.09	25.13	-0.009	7.1	8.46	6	3032.1	288.8
Market 1	1.09	65.87	-0.15	17.0	31.70	10	9733.0	254.6
Market 2	-0.67	34.66	-0.44	21.6	17.37	9	47868.0	284.8
Market 3	-0.27	27.85	0.03	4.0	6.43	4	154.9	281.9
Market 4	-0.49	18.09	-0.22	8.7	7.04	8	2540.8	309.6
Market 5	-0.29	21.92	0.29	7.8	4.58	1	7577.3	242.3
Market 6	-0.42	19.81	-0.039	3.1	4.35	1	36.6	228.2
Market 7	-0.44	26.5	-0.07	9	9.21	6	2958.9	290.2

Table 4. The analysis of ARCH effects.

This table shows the number of simulations the null hypothesis of homoscedasticity is rejected for each market using Engle's (1982) test.

	Rejected at 1% significance level	Rejected at 5% significance level
Fundamental	2	0
Benchmark	4	3
Market 1	10	10
Market 2	8	7
Market 3	3	3
Market 4	5	5
Market 5	1	1
Market 6	2	2
Market 7	6	6

Table 5. Analysis of informational efficiency.

This table shows the average values of three different measures of market efficiency. MAE denotes the mean absolute error (defined as the difference between the transaction price and the fundamental price), MRE is the mean relative error (the absolute error divided by the fundamental price), and RMSE is the root of the mean squared relative error.

	MAE	MRE	RMSE
Benchmark	3.63	3.17%	5.96%
Market 1	17.54	17.24%	28.02%
Market 2	2.00	2.04%	4.68%
Market 3	1.03	1.04%	1.32%
Market 4	0.81	0.82%	1.13%
Market 5	3.54	3.09%	5.84%
Market 6	2.48	2.1%	3.52%
Market 7	3.3	3.06%	6.68%

Table 6. Agents' returns.

This table shows the average agent's return achieved by each class of agents at the end of each market. A positive (negative) return implies that on average this kind of agents has gained (lost) wealth in the trading process in that market.

	ZI agents	Informed agents	Technical agents	ANN agents
Benchmark	-17.19%	31.01%	1.56%	14.10%
Market 1	-8.60%	84.73%	-23.42%	5.48%
Market 2	-16.13%	47.51%	-3.21%	--
Market 3	-16.91%	43.72%	--	20.84%
Market 4	--	--	--	--
Market 5	-22.65%	19.12%	-9.80%	3.66%
Market 6	-9.35%	21.14%	0.95%	11.01%
Market 7	-15.28%	40.21%	-0.24%	13.31%

Figure 1. Price dynamics in the benchmark market.
The dotted line corresponds to transaction prices/returns and the solid line corresponds to fundamental prices/returns.
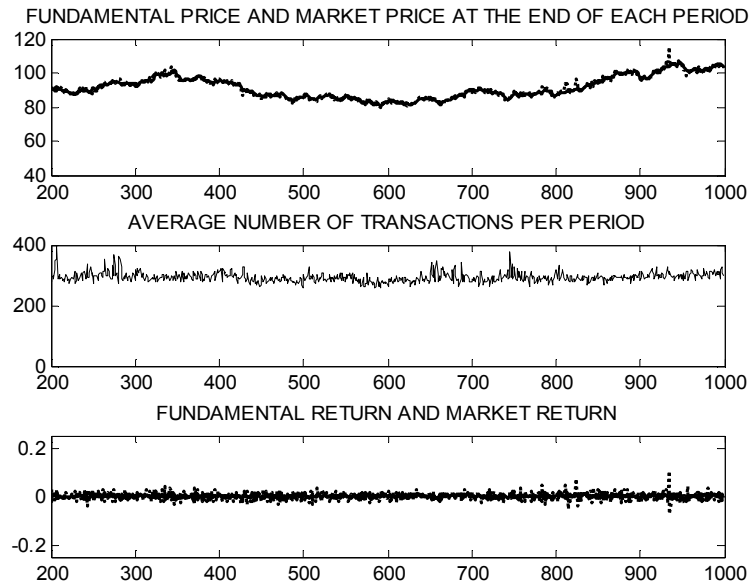
Figure 2. Price dynamics in Market 1.

The dotted line corresponds to transaction prices/returns and the solid line corresponds to fundamental prices/returns.

FUNDAMENTAL PRICE AND MARKET PRICE AT THE END OF EACH PERIOD

AVERAGE NUMBER OF TRANSACTIONS PER PERIOD
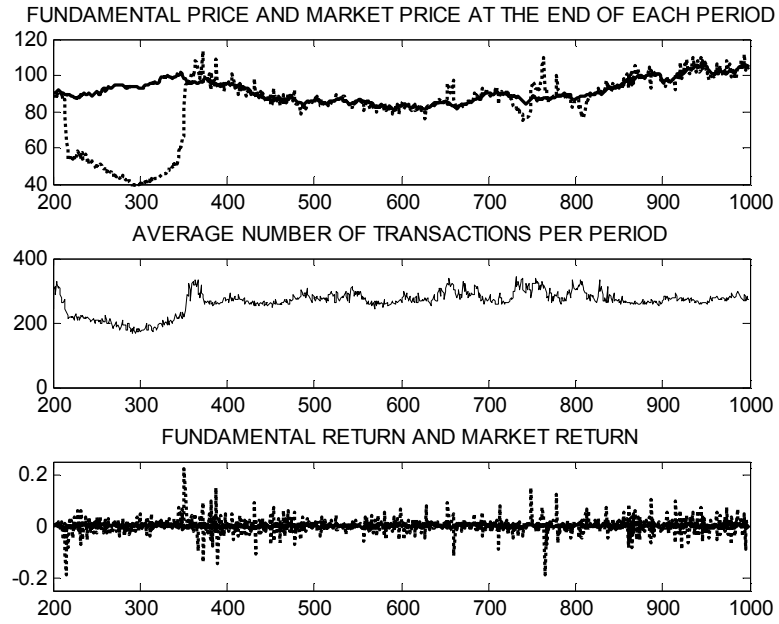
FUNDAMENTAL RETURN AND MARKET RETURN

Figure 3. Agents' average positions per investor category in Market 1.
This table shows the agent's position in cash, risk asset and also the total wealth in each trading period. The thick solid, thin solid, dashed and dotted lines correspond to informed investors, zero intelligence agents, technical analysts and ANN agents, respectively.
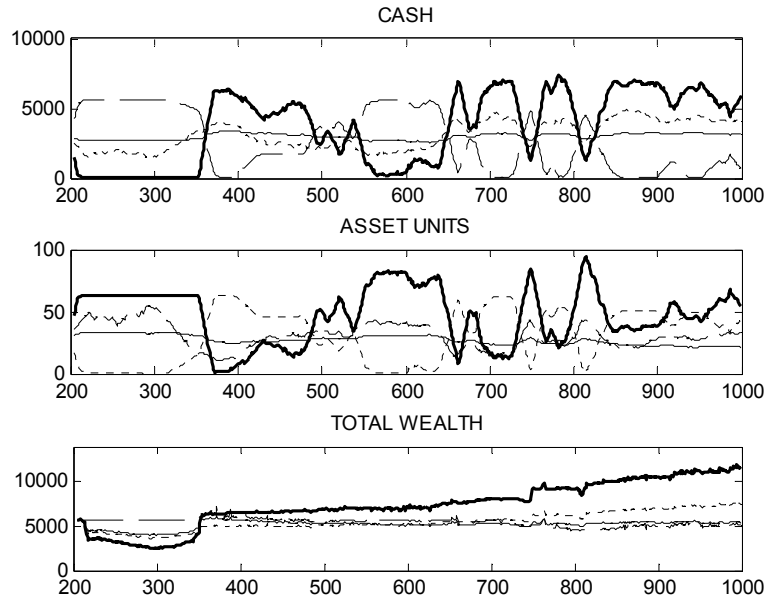
Figure 4. Price dynamics in Market 2.
The dotted line corresponds to transaction prices/returns and the solid line
corresponds to fundamental prices/returns.

FUNDAMENTAL PRICE AND MARKET PRICE AT THE END OF EACH PERIOD

AVERAGE NUMBER OF TRANSACTIONS PER PERIOD
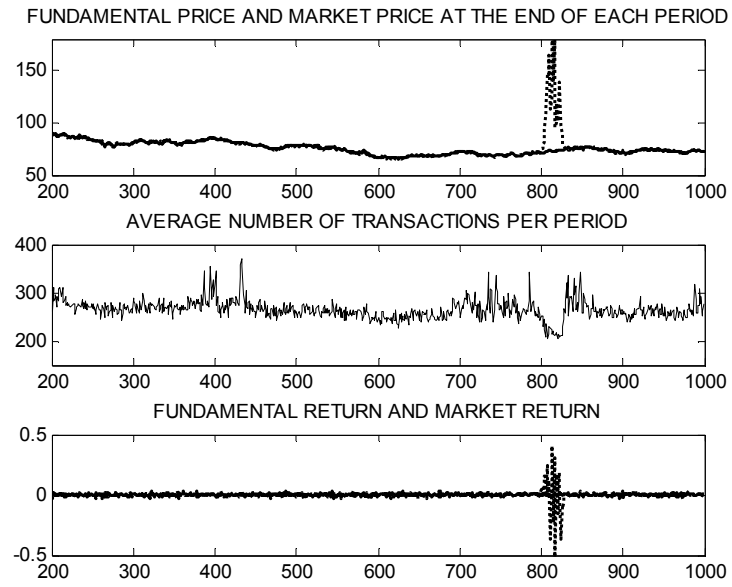
FUNDAMENTAL RETURN AND MARKET RETURN

Figure 5. Agents' average positions per investor category in Market 1.
This table shows the agent's position in cash, risk asset and also the total wealth in each trading period. The thick solid, thin solid, dashed and dotted lines correspond to informed investors, zero intelligence agents, technical analysts and ANN agents, respectively.