An Agent-Based Model of Dealership Markets

Sanmay Das
Center for Biological and Computational Learning
and Computer Science and Artificial Intelligence Lab
Massachusetts Institute of Technology
sanmay@ai.mit.edu

Abstract

This paper describes an agent-based model of financial markets with monopolistic or competitive market-makers and analyzes some of the emergent properties of these markets, including time series properties. The artificial markets we discuss utilize models of "informed" trading agents who decide to trade based on received signals of the true or fundamental value of the stock, and "uninformed" trading agents (sometimes called liquidity traders) who trade for exogenous reasons and are modeled as buying and selling stock randomly. These simple models of traders, combined with more complex market-making agents who function as price-setters and inventory-holders in the market, lead to a rich array of market properties, many of which qualitatively replicate properties observed in real financial markets. For example, the bid-ask spread increases in response to uncertainty about the true value of a stock, average spreads tend to be higher in more volatile markets, and market-makers with lower average spreads perform better in competitive environments. The time series data generated by our market models demonstrate phenomena like volatility clustering and the fat-tailed nature of return distributions, without the need to specify explicit models for opinion propagation and herd behavior in the trading crowd. Our models of trading are simpler than most models in the literature that succeed in demonstrating such properties.

1 Introduction and Related Work

In the last decade there has been a surge of interest in describing equity markets through computational agent models. At the same time, financial markets are an important application area for the fields of agent-based modeling and machine learning, since agent objectives and interactions tend to be more clearly defined, both practically and mathematically, in these markets than in other areas. In this paper we describe market-making agents who play important roles in stock markets and who need to optimize their pricing decisions under conditions of asymmetric information while taking into account considerations such as portfolio risk. We then use these agents to set prices in markets populated by simple trading agents and analyze the properties of the markets thus formed.

1.1 Market Microstructure and Market-Making

Price formation in markets occurs through the process of trading. Market microstructure is concerned with the specific mechanisms and rules under which trades take place in a market and how these mechanisms impact price formation and the trading process. O'Hara [19] and Madhavan [18] present excellent surveys of the market microstructure literature. The simplest type of asset market is a standard double auction market, in which competitive buyers and sellers enter their prices and matching prices result in the execution of trades [11]. Some exchanges like the New York Stock Exchange (NYSE) employ *market-makers* for each stock in order to ensure immediacy and liquidity. The market-maker for each stock on the NYSE is obligated to continuously post two-sided quotes (a *bid* quote and an *ask* quote). Each quote consists of a size and a price, and the market-maker must honor a market buy or sell order of that size (or below) at the quoted price, so that customer buy and sell orders can be immediately executed.

The NYSE employs monopolistic market-makers. Only one market-maker is permitted per stock, and that market-maker is strictly regulated by the exchange to ensure market quality. Market quality can be measured in a number of different ways. One commonly used measure is the average size of the bid-ask spread (the difference between the bid and ask prices). An exchange like the NASDAQ (National Association of Securities Dealers Automated Quotation System) allows multiple market-makers for each stock with less regulation, in the expectation that good market quality will arise from competition between the market-makers[1].

Theoretical analysis of microstructure usually involves modeling order arrival as a stochastic process [12, 13], which is important for the derivation of optimal pricing strategies. The minimization of risk through inventory control is typically modeled mathematically through market-makers' preferred inventory positions and utility functions (see, for example, Amihud and Mendelson [1]). The main problem with theoretical microstructure models is that they are typically restricted to simple, stylized cases with rigid assumptions about trader behavior. This motivates the agent-based approach to modeling markets ([5, 8, 21] *inter alia*).

Trading agents can use heuristics, rules, and machine learning techniques to make trading decisions. Many artificial market simulations also use an evolutionary approach, with agents entering and leaving the market, and agent trading strategies evolving over time. Research in agent-based market models centers on modeling financial markets from the bottom up as structures that emerge from the interactions of individual agents.

Computational modeling of markets allows for the opportunity to push beyond the restrictions of traditional theoretical models of markets through the use of computational power. At the same time, the artificial markets approach allows a fine-grained level of experimental control that is not available in real markets. Thus, data obtained from artificial market experiments can be compared to the predictions of theoretical models and to data from real-world markets, and the level of control allows one to examine precisely which settings and conditions lead to deviations from theoretical predictions. LeBaron [15] provides a summary of some of the early work on agent-based computational finance.

There are two major strands of research on agent-based modeling of financial markets. The first of these focuses on the emergent properties of price processes and attempts to replicate observed properties of financial time series in real markets. For example, the work of Raberto *et al* [21] follows this approach, implementing simple traders who place limit orders, along with a model of opinion propagation among agents in the Genoa Artificial Stock Market. The results described by Raberto *et al* show that their model can capture some features of real financial time series, such as volatility clustering and the leptokurtic distribution of returns. Lux [17] also obtains leptokurtic return distributions in his model, which focuses on chaotic properties of the dynamical system derived from traders changing

[1] A detailed exposition of the different types of market structures is given by Schwartz [22].

between chartist and fundamentalist trading strategies[2].

The other strand of research in artificial markets focuses on the algorithms employed by individual traders in an attempt to understand the environments in which particular strategies are successful, and the resulting implications for market design. Research that follows this pattern includes the reinforcement-learning electronic market-maker designed by Chan and Shelton [4], recent work in the Genoa Artificial Market framework by Cincotti *et al* [6] that studies long-run success of trading strategies, and the NASDAQ-inspired simulations of Darley *et al* [8].

There is a paucity of work on market-making in the artificial markets literature. Some simulations of the NASDAQ stock market have been carried out, but none of them have focused on market-maker behavior or on adaptive agents [8, 3]. With the exception of the work of Chan and Shelton mentioned above, most research on market-making has been in the theoretical finance literature, such as the important paper of Garman [12] which was among the first to explicitly formulate the market-maker's decision problem. Amihud and Mendelson [1] introduced inventory control considerations for market-making. Glosten and Milgrom [13] solve the market-maker's decision problem under information asymmetry.

1.2 Contributions and Overview

We have developed a market-making algorithm that allows us to derive price processes in markets with very simple trading agents who only place market orders [9]. This greatly simplifies the interaction between traders, and allows us to analyze market properties with different populations of traders.

The data from our market simulations yield interesting insights into the behavior of price processes. We compare the time series properties of the price data generated by our simulations to the known characteristics of such data from real markets and find that we are able to replicate some important features of real financial time series, such as the leptokurtic distribution of returns, without postulating explicit, complex models of agent interaction and herd behavior[3] as has previously been done in the literature [17, 21][4].

The rest of this paper is structured as follows. Section 2 provides necessary background information on market microstructure, introduces the market model, and briefly summarizes the price setting algorithm used by market-making agents. Section 3 presents empirical results, including the important time series properties of our model. In section 4, we summarize and suggest avenues for future work.

2 The Market Model

2.1 Market Microstructure Background

There are two main types of orders in stock markets. These are *market orders* and *limit orders*. A market order specifies the size of the order and whether the order is a buy or sell

[2]Interestingly, Lux does not actually implement a multi-agent simulation, but restricts his model to a level of simplicity at which he can model the entire market as a system of nonlinear differential equations.

[3]Some explicit models of herd behavior are presented in the economics literature by Banerjee [2], Cont and Bouchaud [7] and Orléan [20] *inter alia*.

[4]It is worth noting that the true value process can induce behavior similar to that induced by herd behavior through informed traders all buying or selling simultaneously based on superior information. However, the mechanism is a much weaker assumption than the assumption of explicit imitative behavior or mimetic contagion.

order. A limit order also specifies a price at which the trader placing the order is willing to buy or sell. Market orders are guaranteed execution but not price. That is, in placing a market order a trader is assured that it will get executed within a short amount of time at the best market price, but is not guaranteed what that price will be. Limit orders, on the other hand, are guaranteed price but not execution. That is, they will only get executed at the specified price, but this may never happen if a matching order is not found.

Often, markets in which traders can place both market and limit orders are organized as *double auction markets*. The limit orders taken together form an order book, in which the buy orders are arranged in decreasing order of price, while the sell orders are arranged in increasing order of price. Orders that match are immediately executed, so the highest buy order remaining must have a lower price than the lowest ask order remaining. Market orders, when they arrive, are executed against the best limit order on the opposite side. So, for example, a market buy order would get executed against the best limit sell order currently on the book.

Double auction markets are effective when there is sufficient liquidity in the stock. There must be enough buy and sell orders for incoming market orders to be guaranteed immediate execution at prices that are not too far away from the prices at which transactions executed recently. Sometimes these conditions are not met, typically for stocks that do not trade in high volume, and immediately following particularly favorable or unfavorable news, when everyone wants to be either on the buy or sell side of the market, leading to huge order imbalances.

Market-makers are traders designated by markets to maintain immediacy and liquidity in transactions. Market-makers are obligated to continuously post two-sided quotes (*bid* (for buying) and *ask* (for selling) quotes) and honor these quotes. Market-makers are expected to smooth the transition when the price of a stock jumps dramatically, so that traders do not believe they received unfair executions, and to maintain a reasonable bid-ask spread. Exchanges with monopolistic market-makers like the NYSE monitor the performance of market-makers on these categories, while markets like NASDAQ use multiple market-makers and expect good market quality to arise from competition between market-makers.

2.2 Detailed Market Model

In our market model, we only allow the trading crowd to place market orders, which get executed against the limit orders continuously posted by market-makers. We analyze is a discrete time dealer market with only one stock. The market-maker sets bid and ask prices (P_b and P_a respectively) at which it is willing to buy or sell one unit of the stock at each time period (when necessary we denote the bid and ask prices at time period i as P_b^i and P_a^i). If there are multiple market-makers, the market bid and ask prices are the maximum over each dealer's bid price and the minimum over each dealer's ask price. All transactions occur with the market-maker taking one side of the trade and a member of the trading crowd (henceforth a "trader") taking the other side.

We assume that a stock has an underlying true value (or fundamental value) at all points in time. The price at which the stock trades is not necessarily close to this value at all times (for example, during a bubble, the stock trades at prices considerably higher than its true value). There are two principal kinds of traders in the market. *Informed traders* (sometimes referred to as fundamentalist traders) are those who know (or think they know) the true value of the stock and base their decisions on the assumption that the transaction price will revert to the true value. Informed traders will try to buy when they think a stock is

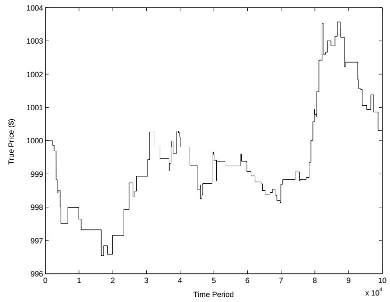

Figure 1: Example of the true value over time

undervalued by the market price, and will try to sell when they think a stock is overvalued by the market price. Sometimes it is useful to think of informed traders as those possessing inside information. *Uninformed traders* (also referred to as noise traders) trade for reasons exogenous to the market model. Usually they are modeled as buying or selling stock at random (one psychological model is traders who buy or sell for liquidity reasons). Other models of traders are often mentioned in the literature, such as *chartists* who attempt to predict the direction of stock price movement, but we have so far not incorporated such traders into our market model.

The true underlying value of the stock is V^i at time period i. All market makers are informed of V^0 at the beginning of a simulation, but do not receive any direct information about V after that[5]. At time period i, a single trader is selected from the trading crowd and allowed to place either a (market) buy or (market) sell order for one unit of the stock. An uninformed trader will place a buy or sell order for one unit at random if selected to trade. An informed trader who is selected to trade knows V^i and will place a buy order if $V^i > P_a^i$, a sell order if $V^i < P_b^i$ and no order if $P_b^i \leq V^i \leq P_a^i$.

In addition to perfectly informed traders, we also allow for the presence of noisy informed traders. A noisy informed trader receives a signal of the true price $W^i = V^i + \tilde{\eta}(0, \sigma_W)$ where $\tilde{\eta}(0, \sigma_W)$ represents a sample from a normal distribution with mean 0 and variance σ_W^2. The noisy informed trader believes this is the true value of the stock, and places a buy order if $W^i > P_a^i$, a sell order if $W^i < P_b^i$ and no order if $P_b^i \leq W^i \leq P_a^i$.

The true underlying value of the stock evolves according to a jump process. At time $i+1$, with probability p, a jump in the true value occurs[6]. When a jump occurs, the value changes according to the equation $V^{i+1} = V^i + \tilde{\omega}(0, \sigma)$ where $\tilde{\omega}(0, \sigma)$ represents a sample from a normal distribution with mean 0 and variance σ^2.

This model of the evolution of the true value corresponds to the notion of the true value evolving as a result of occasional news items. The periods immediately following jumps are the periods in which informed traders can trade most profitably, because the information they have on the true value has not been disseminated to the market yet, and the market maker is not informed of changes in the true value and must estimate these through orders placed by the trading crowd. The market-maker will not update prices to the neighborhood of the new true value for some period of time immediately following a jump in the true value, and informed traders can exploit the information asymmetry.

[5]That is, the only signals a market-maker receives about the true value of the stock are through the buy and sell orders placed by the trading crowd.

[6]p is typically small, of the order of 1 in 1000 in most of our simulations

2.3 An Algorithm for Market-Making

The market-maker attempts to track the true value over time by maintaining a probability distribution over possible true values and updating the distribution when it receives signals from the orders that traders place. The true value and the market-maker's prices together induce a probability distribution on the orders that arrive in the market. The market-maker's task is to maintain an online probabilistic estimate of the true value, which is itself a moving target.

Glosten and Milgrom [13] analyze the setting of bid and ask prices so that the market maker enforces a zero profit condition. The zero profit condition corresponds to the Nash equilibrium in a setting with competitive market-makers (or, more generally in any competitive price-setting framework [10]). Glosten and Milgrom suggest that the market maker should set $P_b = E[V|\text{Sell}]$ and $P_a = E[V|\text{Buy}]$. Our market-making algorithm approximately computes these expectations using the probability density function being estimated. It uses an explicit histogram representation of the density function and Bayesian updating based on observed orders from the trading crowd. A detailed description of the algorithm can be found in [9]. The price setting equations derived in that paper are as follows. In the case of perfectly informed traders, the market maker chooses P_b and P_a as:

$$P_b = 2\left(\sum_{V_i=V_{\min}}^{V_i=P_b}(\frac{1}{2}+\frac{1}{2}\alpha)V_i \Pr(V=V_i) + \sum_{V_i=P_b+1}^{V_i=V_{\max}}(\frac{1}{2}-\frac{1}{2}\alpha)V_i \Pr(V=V_i)\right) \quad (1)$$

and

$$P_a = 2\left(\sum_{V_i=V_{\min}}^{V_i=P_a}(\frac{1}{2}-\frac{1}{2}\alpha)V_i \Pr(V=V_i) + \sum_{V_i=P_a+1}^{V_i=V_{\max}}(\frac{1}{2}+\frac{1}{2}\alpha)V_i \Pr(V=V_i)\right) \quad (2)$$

Where $V_{\min}$ and $V_{\max}$ are practical upper and lower bounds on the true value and α is the proportion of informed traders. The probabilities $\Pr(V = V_i)$ are explicitly stored as part of the density estimation technique. In the case of noisy informed traders, the equations become:

$$P_b = 2\sum_{V_i=V_{\min}}^{V_i=P_b}(\frac{1}{2}-\frac{1}{2}\alpha+\alpha\Pr(\tilde{\eta}(0,\sigma_W)\leq(P_b-V_i)))V_i\Pr(V=V_i)+$$
$$2\sum_{V_i=P_b+1}^{V_i=V_{\max}}(\frac{1}{2}-\frac{1}{2}\alpha+\alpha\Pr(\tilde{\eta}(0,\sigma_W)\geq(V_i-P_b)))V_i\Pr(V=V_i) \quad (3)$$

and

$$P_a = 2\sum_{V_i=V_{\min}}^{V_i=P_a}(\frac{1}{2}-\frac{1}{2}\alpha+\alpha\Pr(\tilde{\eta}(0,\sigma_W)\geq(P_a-V_i)))V_i\Pr(V=V_i)+$$
$$2\sum_{V_i=P_a+1}^{V_i=V_{\max}}(\frac{1}{2}-\frac{1}{2}\alpha+\alpha\Pr(\tilde{\eta}(0,\sigma_W)\leq(V_i-P_a)))V_i\Pr(V=V_i) \quad (4)$$

Here, $\tilde{\eta}(0,\sigma_W)$ is a sample from a Gaussian distribution with mean 0 and variance σ_W^2, the variance of the noise process.

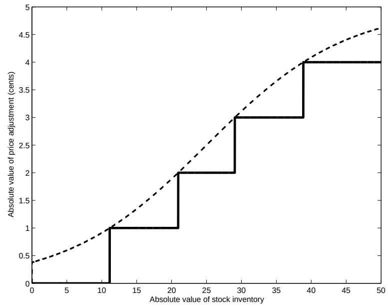

Figure 2: Step function for inventory control and the underlying sigmoid function

Various layers of complexity can be added on top of the basic algorithm. For example, minimum and maximum conditions can be imposed on the spread, and an inventory control mechanism could form another layer after the zero-profit condition prices are decided.

2.4 Inventory Control

Stoll analyzes dealer costs in conducting transactions and divides them into three categories [23]. These three categories are portfolio risk, transaction costs and the cost of asymmetric information. In the model we have presented so far, following Glosten and Milgrom [13], we have assumed that transactions have zero execution cost and developed a pricing mechanism that explicitly attempts to set the spread to account for the cost of asymmetric information. A realistic model for market-making necessitates taking portfolio risk into account and controlling inventory in setting bid and ask prices. In the absence of consideration of trade size and failure conditions, portfolio risk should affect the placement of the bid and ask prices, but not the size of the spread[7] [1, 23, 14]. If the market-maker has a long position in the stock, minimizing portfolio risk is achieved by lowering both bid and ask prices (effectively making it harder for the market-maker to buy stock and easier for it to sell stock), and if the market-maker has a short position, inventory is controlled by raising both bid and ask prices.

Inventory control can be incorporated into the architecture of our market-making algorithm by using it as an adjustment parameter applied after bid and ask prices have been determined by equations 3 and 4. An example of the kind of function we can use to determine the amount of the shift is a sigmoid function. The motivation is to allow for an initial gradual increase in the impact of inventory control on prices, followed by a steeper increase as inventory accumulates, while simultaneously bounding the upper limit by which inventory control can play a factor in price setting. Of course, the upper bound and slope of the sigmoid can be adjusted according to the qualities desired in the function. For our simulations, we use an inventory control function that uses the floor of a real valued sigmoid function with a ceiling of 5 cents as the integer price adjustment (in cents). The step function for the adjustment and the underlying sigmoid are shown in figure 2. Average market-maker inventories at the end of a simulation, and the negative correlation between market-maker profits and market volatility are both reduced by the use of such an inventory control module [9].

[7] One would expect spread to increase with the trade size.

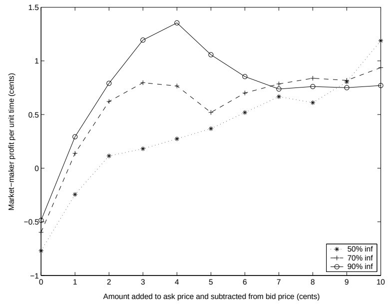

Figure 3: Market-maker profits as a function of increasing the spread

2.5 Profit Motive

The zero-profit condition of Glosten and Milgrom is expected from game theoretic considerations when multiple competitive dealers are making markets in the same stock. However, since our method is an approximation scheme, the zero profit method is unlikely to truly be zero-profit. Further, the market-maker is not always in a perfectly competitive scenario where it needs to restrict the spread as much as possible.

The simplest solution to the problem of making profit is to increase the spread by pushing the bid and ask prices apart after the zero-profit bid and ask prices have been computed using the density estimate obtained by the market-making algorithm. The major effect of this on the density estimation technique is that the signals the market-maker receives and uses to update its density estimate are determined by transaction prices, which are in turn determined by the bid and ask prices the market-maker has set. The precise values of the bid and ask prices are quite important to the sampling of the distribution on trades induced by the true value.

3 Empirical Results and Time Series Properties

3.1 Experimental Framework

Unless specified otherwise, it can be assumed that all simulations take place in a market populated by noisy informed traders and uninformed traders. The noisy informed traders receive a noisy signal of the true value of the stock with the noise term being drawn from a Gaussian distribution with mean 0 and standard deviation 5 cents. The standard deviation of the jump process for the stock is 50 cents, and the probability of a jump occurring at any time step is 0.005. The market-maker is informed of when a jump occurs, but not of the size or direction of the jump. The market-maker uses an inventory control function and increases the spread by lowering the bid price and raising the ask price by a fixed amount to ensure profitability. We report average results from 50 simulations, each lasting 50,000 time steps.

3.2 Profit Motive

Figure 3 shows the profit obtained by a single monopolistic market-maker in markets with different percentages of noisy informed traders. The numbers on the X axis show the amount (in cents) that is subtracted from (added to) the zero-profit bid (ask) price in order to push

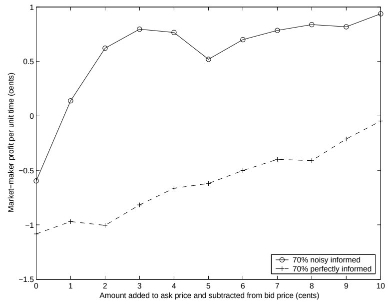

Figure 4: Market-maker profits with noisy informed traders and perfectly informed traders

the spread apart (we will refer to this number as the shift factor). It is important to note that market-makers can make reasonable profits with low average spreads – an example is given at the end of the section on competitive market-making. This is one test of whether the results of the model are reasonable and relevant to the real world.

With lower spreads, most of the market-maker's profits come from the noise factor of the informed traders, whereas with a higher spread, most of the market-maker's profits come from the trades of uninformed traders. Whenever the market-maker makes higher profits, this means that the trading crowd is losing money. Different percentages of informed traders lead to differently shaped curves. With only 50% of the traders being informed, the market-maker's profit keeps increasing with the size of the spread. However, increasing the spread beyond a point is counterproductive if there are enough noisy informed traders in the markets, because then the market-maker's prices are far enough away from the true value that even the noise factor cannot influence the informed traders to make trades. With 90% of the traders being informed, a global maximum (at least for reasonable spreads) is attained with a low spread, while with 70% of the traders being informed, a local maximum is attained with a fairly low spread, although the larger number of uninformed traders allows for larger profits with rather large spreads.

The market-maker's probability density estimates tend to be more concentrated with more noisy informed traders in the markets, because each trade provides more information. This leads to the empirical results being closer to theoretical predictions. For example, the prices leading to zero profit for the 70% informed and 90% informed cases fall between the 0 and 1 points on the X axis, which is close to what one would expect from the theory, whereas with perfectly informed traders zero profit is not obtained without using a large spread.

Figure 4 compares the profits obtained by the market-maker with 70% noisy informed traders as opposed to 70% perfectly informed traders. In the latter case, there is no advantage to be gained by having a smaller spread as there is with noisy informed traders. However, the market-maker's inability to make any profit even with a high spread seems surprising. This is partly attributable to the fact that the point at which the distribution is sampled is more important in the perfectly informed case because the signals only inform the market-maker of the probabilities that the true value is greater than or less than the last transaction price, instead of smoothly morphing points around the last transaction price by a mixture of a Gaussian distribution and a uniform distribution.

70% noisy informed traders				
	Competitive		Monopolistic	
Shift Factor	Profit	# Trades	Profit	# Trades
2	0.6039	38830	0.6216	39464
3	0.0157	594	0.8655	34873

Table 1: Market-maker profits (in cents per time period) and average number of trades in simulations lasting 50,000 time steps in monopolistic and competitive environments

3.2.1 Competitive Market-Making

In competitive environments, the highest bid price quoted by any market-maker and the lowest ask price quoted by any market-maker become the effective market bid-ask quotes, and the market-makers compete with each other for trades. A market-maker who does not make a sufficient number of trades will lose out to a market-maker who makes substantially more trades even if the latter makes less profit per trade.

This effect is particularly obvious in the experiment shown in table 1. In this experiment, two market-makers compete with each other, with the difference that one uses a shift factor of 2 and the other a shift factor of 3 for increasing the spread after the zero-profit bid and ask prices have been determined. If neither were using an inventory control mechanism, the market-maker using a shift factor of 3 would make no trades, because the market-maker using a shift factor of 2 would always have the inside quotes. The addition of inventory control allows the market-maker using a shift factor of 3 to make some trades, but this market-maker makes considerably less profit than the one using a shift factor of 2. In a monopolistic environment the market-maker using a shift factor of 3 outperforms the market-maker using a shift factor of 2 and the difference in magnitude of executed trades is not as large.

For a market with 70% of the trading crowd consisting of noisy informed traders and the remaining 30% consisting of uninformed traders, our algorithm, using inventory control and a shift factor of 1, achieves an average profit of 0.0074 ± 0.0369 cents per time period with an average spread of 2.2934 ± 0.0013 cents. These parameter settings in this environment yield a market-maker that is close to a Nash equilibrium player, and it is exceedingly unlikely that any algorithm would be able to outperform this one in direct competition in such an environment given the low spread. An interesting avenue to explore is the possibility of adaptively changing the shift factor depending on the level of competition in the market. Clearly, in a monopolistic setting, a market-maker is better off using a high shift factor, whereas in a competitive setting it is likely to be more successful using a smaller one. An algorithm for changing the shift factor based on the history of other market-makers' quotes would be useful.

3.3 The Effects of Volatility

Volatility of the underlying true value process is affected by two parameters. One is the standard deviation of the jump process, which affects the variability in the amount of each jump. The other is the probability with which a jump occurs. Table 2 shows the result of changing the standard deviation σ of the jump process and table 3 shows the result of changing the probability p of a jump occurring at any point in time. As expected, the

σ	Shift	Spread	Profit
100	1	2.7366	-0.7141
100	2	5.0601	-0.1410
50	1	2.2934	0.0074
50	2	4.4466	0.6411

Table 2: Market-maker average spreads (in cents) and profits (in cents per time period) as a function of the standard deviation of the jump process

p	Shift	Spread	Profit
0.005	1	2.2934	0.0074
0.005	2	4.4466	0.6411
0.0001	1	2.0086	0.8269
0.0001	2	4.0154	1.4988

Table 3: Average bid-ask spreads (in cents) and market-maker profits (in cents per time period) as a function of the probability of a jump occurring at any point in time

spread increases with increased volatility, and profit decreases. A higher average spread needs to be maintained to get the same profit in more volatile markets.

3.4 Time Series Properties of Transaction Prices

Liu *et al* present a detailed analysis of the time series properties of returns in a real equity market (they focus on the S&P 500 and component stocks) [16]. Their major findings are that return distributions are leptokurtic and fat-tailed, volatility clustering occurs (that is, big price changes are more likely to be followed by big price changes and small price changes are more likely to be followed by small price changes)[8] and that the autocorrelation of absolute values of returns decays according to a power law, and is persistent over large time scales, as opposed to the autocorrelation of raw returns, which disappears rapidly.

Raberto *et al* are able to replicate the fat tailed nature of the distribution of returns and the clustered volatility observed in real markets [21]. However, the Genoa Artificial Stock Market explicitly models opinion propagation and herd behavior among trading agents in a way that we do not[9]. Nevertheless, our model is also able to replicate the important stylized facts of real financial time series, including the leptokurtic distribution of returns, clustered volatility and persistence of the autocorrelation of absolute returns.

A return over a particular time period is defined as the ratio of the prices at which two transactions occur which are separated by that period in time. In our model, a one step return is the ratio of the prices at which two successive transactions occur. We record all transaction prices and assume that the intervals between transactions are the same. All experiments in this section are in a market with 70% noisy informed traders, 30% uninformed traders, and a market-maker using a shift factor of 1 (which results in the

[8] Liu *et al* are certainly not the first to discover these properties of financial time series. However, they summarize much of the work in an appropriate fashion and provide detailed references, and they present novel results on the power law distribution of volatility correlation.

[9] A jump in the true value will lead informed traders in our model to make the same decisions on whether to buy or sell, but not because of imitative behavior among the agents themselves.

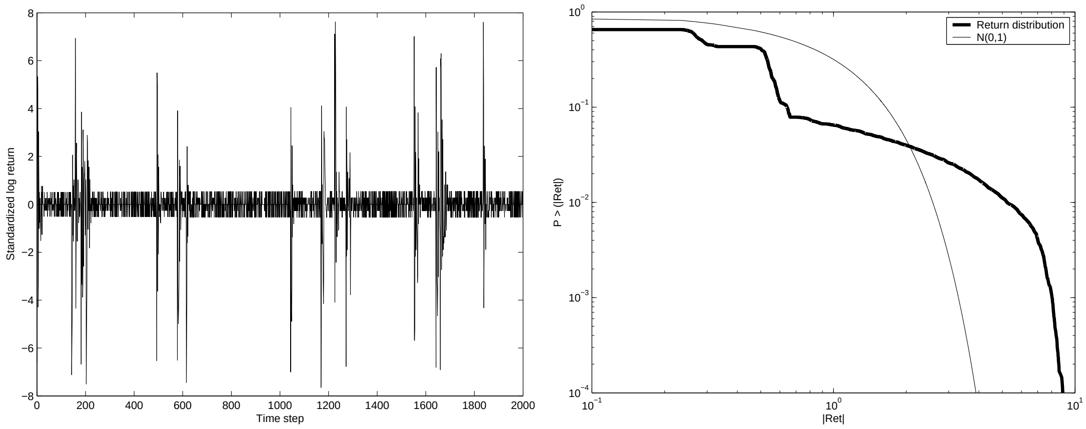

Figure 5: Standardized log returns over time (left) and leptokurtic distribution of absolute standardized log returns (right)

market-maker's profit being close to zero). We work with log returns in this model, where the log return is $\log P_{n+1} - \log P_n$. Figure 5 shows the standardized log returns over time and their distribution. Standardized log returns are log returns detrended by the mean and rescaled by the standard deviation. The clustering of volatility, sharp tails, and leptokurtic distribution of returns are evident from the figure. The fat tail is evident from the right half of the graph, where the area being covered by the distribution of returns "pokes out" from the area covered by a normal distribution. The sample kurtosis for this experiment was 28.7237.

The other important feature of real financial time series that our market also shows are the long-range persistence of the autocorrelation of absolute returns and the rapid decay of autocorrelation of raw returns (figure 6). Interestingly, the decay of autocorrelation appears to be linear, which is in contrast with the power law decay observed by Liu *et al.* If we look more closely, the decay is linear on a log-log scale for the first 25 lags (indicative of a power law decay) (figure 7). The long range persistence of autocorrelation is an important feature of real financial markets [16], but in comparison, Raberto *et al* fail to observe persistence of autocorrelation beyond 80 lags in the Genoa market, and they do not see decay consistent with a power law at any scale. Real markets, the Genoa market and our artificial market all show quick decay of the autocorrelation of raw returns.

4 Conclusions and Future Work

4.1 Summary

This paper describes the key components of an agent-based model of dealership markets. The market-makers in our model use Glosten and Milgrom's basic concept of how to set bid and ask prices and generalize this to a more realistic environment than the toy problems originally considered in that framework [13]. Our market model allows us to study issues of asymmetric information at a new level of computational detail, with informed and uninformed traders.

Simulations to validate our models of market-making and trading agents reveal that our markets display realistic properties; for example, bid-ask spreads are higher in more volatile

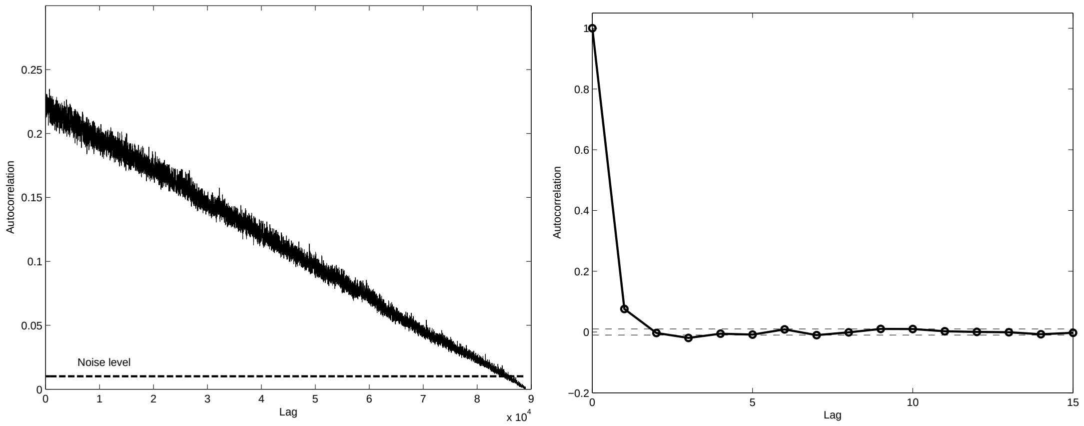

Figure 6: Autocorrelation of absolute (left) and raw (right) returns

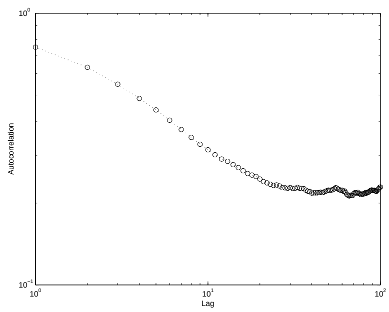

Figure 7: Autocorrelation of absolute returns on a log-log scale for lags of 1-100

markets and market-makers with lower spreads perform better in competitive situations.

Artificial markets populated by the kinds of trading crowds and market-makers we describe replicate some of the important time series phenomena of real financial markets. For example, the leptokurtic distribution of returns and the persistence of the autocorrelation of absolute returns along with the rapid decay of the autocorrelation of raw returns are important phenomena in financial time series [16]. These phenomena are replicated to some extent in the artificial markets described by Lux [17] and Raberto *et al* [21] among others, but only with explicit models of opinion propagation and evolutionary behavior in the trading crowd. The fact that our model does not need to postulate such behavior, instead relying on the simple interaction between informed and uninformed traders, may point to an important underlying regularity of such time series phenomena.

4.2 Future Directions

Our market model will allow us to study a number of interesting issues in both market microstructure and the dynamics of the trading crowd. In terms of microstructure, especially market-making, our agent-based model allows us to study issues like market quality in regulated monopolistic vs. unregulated competitive dealer markets (like NYSE vs. NASDAQ), and the differences between theoretically predicted and empirically observed effects of decimalization.

In the context of multi-agent learning, it would be interesting to have competitive market-makers using the same basic algorithm try to learn the optimal shift factor to use when competing against each other. It is conceivable that this could give rise to cooperative (or collusive) behavior without the need for explicit communication.

More detailed examination of the time series properties of returns and an analysis of why they differ from real markets in the characteristics in which they do differ is an important step. We are also interested in calibrating the artificial market parameters to real markets. For example, the probability of a jump or the standard deviation of the jump could be usefully linked to occurrences in real markets. Finally, it is important to investigate richer, more complex market-models. The first step in this direction is to incorporate consideration of different trade sizes, which we have ignored in the existing model.

In terms of trading dynamics, we are actively inverstigating different distributions of information among traders. For example, suppose their are multiple signals of the true price, some or all of which are observable by particular traders. How do traders decide which signals to use (especially if there is a cost to switching between signals), and what effect does this have on market properties? We believe that our artificial market framework is an exciting testbed for exploring these and other questions.

Acknowledgments

I would like to thank Tommy Poggio, Andrew Lo and Adlar Kim for useful comments. This research was sponsored by grants from Merrill-Lynch, National Science Foundation (ITR/IM) Contract No. IIS-0085836 and National Science Foundation (ITR/SYS) Contract No. IIS-0112991. Additional support was provided by the Center for e-Business (MIT), DaimlerChrysler AG, Eastman Kodak Company, Honda R&D Co., Ltd., The Eugene McDermott Foundation, and The Whitaker Foundation.

References

[1] Y. Amihud and H. Mendelson. Dealership Market: Market-Making with Inventory. *Journal of Financial Economics*, 8:31–53, 1980.

[2] Abhijit V. Banerjee. A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.

[3] M. Brown and V. Darley. The Future of Trading: Biology-based Market Modeling at Nasdaq. *Perspectives in Business Innovation*, 4, 1999.

[4] N. T. Chan and C. R. Shelton. An Electronic Market Maker. Technical Report 2001-005, MIT Artificial Intelligence Laboratory, 2001.

[5] T. Chan. *Artificial Markets and Intelligent Agents*. PhD thesis, Massachusetts Institute of Technology, 2001.

[6] S. Cincotti, S.M. Focardi, M. Marchesi, and M. Raberto. Who Wins? Study of Long-Run Trader Survival in an Artificial Stock Market. *Physica A*, 2002. In press.

[7] Rama Cont and Jean-Philipe Bouchaud. Herd Behavior and Aggregate Fluctuations in Financial Markets. *Macroeconomic Dynamics*, 4:170–196, 2000.

[8] V. Darley, A. Outkin, T. Plate, and F. Gao. Sixteenths or Pennies? Observations from a Simulation of the Nasdaq Stock Market. In *IEEE/IAFE/INFORMS Conference on Computational Intelligence for Financial Engineering*, 2000.

[9] Sanmay Das. Intelligent market-making in artificial financial markets. Technical Report 2003-005, MIT Artificial Intelligence Laboratory, 2003.

[10] D.Fudenberg and J.Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.

[11] D. Friedman and J. Rust. The Double Auction Market Institutions, Theories, and Evidence. Technical report, Santa Fe Institute Studies in the Sciences of Complexity, 1991.

[12] M. Garman. Market Microstructure. *Journal of Financial Economics*, 3:257–275, 1976.

[13] L.R. Glosten and P.R. Milgrom. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics*, 14:71–100, 1985.

[14] S.J. Grossman and M.H. Miller. Liquidity and Market Structure. *Journal of Finance*, 43:617–633, 1988.

[15] B. LeBaron. Agent-based Computational Finance: Suggested Readings and Early Research. *Journal of Economic Dynamics and Control*, 24:679–702, 2000.

[16] Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C. Peng, and H.E. Stanley. Statistical Properties of the Volatility of Price Fluctuations. *Physical Review E*, 60(2):1390–1400, 1999.

[17] T. Lux. 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, 1998.

[18] A. Madhavan. Market Microstructure: A Survey. *Journal of Financial Markets*, pages 205–258, 2000.

[19] M. O'Hara. *Market Microstructure Theory*. Blackwell, Malden, MA, 1995.

[20] André Orléan. Bayesian Interactions and Collective Dynamics of Opinion: Herd Behavior and Mimetic Contagion. *Journal of Economic Behavior and Organization*, 28:257–274, 1995.

[21] M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Agent-based Simulation of a Financial Market. *Physica A*, 299:319–327, 2001.

[22] Robert A. Schwartz. *Reshaping the Equity Markets: A Guide for the 1990s*. Harper Business, New York, NY, 1991.

[23] H. Stoll. The Supply of Dealer Services in Securities Markets. *Journal of Finance*, 33:1133–1151, 1978.