

The Effects of Market-Making on Price Dynamics

Sanmay Das
Dept. of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180-3590
sanmay@cs.rpi.edu

ABSTRACT

This paper studies market-makers, agents responsible for maintaining liquidity and orderly price transitions in markets. Market-makers include major firms making markets on global stock exchanges, as well as software agents that run behind the scenes on novel electronic markets like prediction markets. We use a sophisticated model of market-making to build richer agent-based models of markets and show how these models can be useful both in understanding properties of existing markets and in predicting the impacts of structural changes. For example, we show how competition among market-makers can lead to significantly faster price discovery following a jump in the true value of an asset. We also show that myopic profit-maximization, apart from leading to poor market quality, is sub-optimal even for a monopolistic market-maker. This observation leads to an interesting characterization of the market-maker's exploration-exploitation dilemma as a tradeoff between price discovery and profit-taking.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Algorithms, Economics

Keywords

Agent-based simulations :: emergent behavior, Economic paradigms :: electronic markets and institutions

1. INTRODUCTION

With the dramatic increase in electronic exchanges and automated trading in recent years, it has become important to develop new computational and algorithmic tools for analyzing market properties and designing software agents that participate in market activities. Many aspects of exchanges and price dynamics have recently received attention from an algorithmic perspective. For example, Kakade *et al* [10] study algorithmic trading in terms of the problem of opti-

mizing trade execution, and Even-Dar *et al* [5] examine the dynamics of order-books under different models of trading.

This paper contributes to this growing literature by studying the role of *market-makers*. Most modern exchanges, ranging from major equity markets like the NYSE and NASDAQ to electronic betting and prediction markets like tradesports.com, employ or designate agents with special responsibilities for maintaining liquidity and orderly price transitions. These market-makers are obligated to continuously post and honor two-sided (buying and selling) prices.

While market-makers have traditionally been employed by large stock exchanges, novel modern markets like prediction markets are now using market-makers to improve market quality [18]. These markets are often illiquid, and the presence of market-makers can bootstrap them into a sufficiently liquid phase to attract trading. We need to better understand the impact of market-makers in such markets.

The market-making problem can be analyzed from two perspectives. First, how does one design an effective market-making algorithm, given some knowledge of the market structure? Second, what are the implications of the presence of market-makers using these algorithms on price dynamics? This paper builds on previous theoretical and simulation studies on market-making to analyze price properties in stylized market models when market-makers are present with different constraints on their behavior.

We focus particularly on price dynamics, rather than just equilibrium behavior. The efficient markets hypothesis, in its various forms, says that prices reflect all available information. But *how* do prices come to reflect the available information? What processes do markets follow to incorporate new information into prices? As an example, suppose new positive information about a stock is relayed to all the participants in a stock market. We know that the traded price should go up to reflect the new information. However, what process will it follow in its rise? Will the increase be orderly and in small increments or will there be a sudden jump? How will the price process be affected by different possible market structures? Computational modeling is an ideal tool for studying these problems.

Specifically, this paper considers the market-maker's pricing problem in markets populated by traders who receive better information about the value of a stock than the market-maker. Typically, this happens when there is an informational shock that provides information to some traders (for example, the release of an analyst report to subscribers). The market-maker must set bid and ask prices to at least offset the adverse selection costs she incurs by trading with

Cite as: The Effects of Market-Making on Price Dynamics, Sanmay Das, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16, 2008, Estoril, Portugal, pp. 887-894.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

potentially better informed traders.

We start by summarizing some relevant background in market microstructure, and then present the model and simulation results. We show that a myopically optimizing market-maker does not achieve maximal long-run profit, and pose the problem of how to optimally balance exploration and exploitation in this setting. We also demonstrate that the presence of market-makers can speed up the process of price discovery and lead to better market quality even when other traders are allowed to place limit orders.

2. MARKET MICROSTRUCTURE

The field of market microstructure is concerned with the specific mechanisms and rules that govern trading in a market and how these mechanisms impact price formation and the trading process (two comprehensive surveys are [15, 13]). This is markedly different from much of economics and finance theory that abstracts away from the process of trading and assumes equilibrium pricing. Traders in real markets have to interact with some kind of realistic pricing mechanism. We cannot just assume that all trading takes place at an equilibrium price conveniently determined by a Walrasian auctioneer who sets the price to clear the market with knowledge of every potential trader's demand function. The insights provided by the study of microstructure can have great impact in the design and regulation of markets.

2.1 Limit Orders and Market Orders

Most modern markets function as continuous double auctions with two types of orders, known as *limit orders* and *market orders*.¹ Market orders are guaranteed immediate execution, but not price. That is to say, if an agent places a market order to buy or sell a certain number of shares, those shares will be bought or sold at the prevailing market prices. This is what is commonly thought of as a buy or sell order in the market. A typical market order will be of the form "Buy/Sell X shares."

The prices at which market orders execute are determined by the limit order book. Traders may also place *limit orders* of the type "Buy/Sell X shares at price Y." the highest buy limit order and the lowest sell limit order constitute the market *bid* and *ask* prices, and the difference between them is known as the *bid-ask spread* or just the spread. When a market buy order arrives, it is executed against the lowest limit sell order, and, similarly, a market sell order is executed against the highest limit buy order. Figure 1 shows an example.

2.2 Liquidity and Market-Making

Different exchanges may employ one or multiple *market-makers*. The key function of a market-maker is the provision of liquidity – ensuring that there is enough interest in a stock to maintain a reasonable amount of trading without steep price changes. An exchange would want to employ firms to make markets in securities in order to ensure the smooth functioning of the market. As a trader, one wants to be assured that (1) market orders will get executed in a

¹For the modeling purposes of this paper we are not concerned with the precise functioning of the order book, but an excellent description from a modern perspective can be found in the recent work of Kakade *et al* [10], and a more elaborate, albeit older, description is that of Schwartz [17].




Figure 1: Example of a market buy order for 1000 shares executing against the limit sell order book. First, 600 shares execute at \$77.75 against Order A, then 400 shares execute against Order B at \$77.79. The volume weighted average price of this order execution is then \$77.766. Note that the remaining 200 shares of Order B remain on the book.

reasonable amount of time, and (2) market orders of a reasonable size will execute at prices close to the quoted bid and ask prices. If an exchange cannot make these guarantees, traders may not be willing to trade on the exchange, and firms would not want to be listed on the exchange.

Of course, liquidity in markets is partly a function of traders themselves placing limit orders. The depth of the limit order book, and the differences in price between adjacent orders usually provide a good measure of liquidity. However, most traders are never obligated to place limit orders, and limit orders can also be cancelled after being placed. Therefore, one can never guarantee execution of a market order in the absence of an institutional structure that ensures the existence of a trader always willing to take the other side of a trade.

A large shock, say from a surprising earnings report, can quickly lead to major changes in the valuations of those trading in a stock, and it is likely that one side of the market (the buy side if the shock is negative and the sell side if it is positive) will become thin. The market-makers in a stock are supposed to ensure that there are no large sudden jumps in the price by stepping in and absorbing the other side of orders no one else wants to absorb in this situation. This will guarantee that no trader feels like they got a bad deal because the trade immediately before theirs executed at a significantly better price.

Major exchanges have adopted different models for liquidity provision through market-making. For example, the NYSE employs a single, regulated *specialist*² for each stock, while the NASDAQ relies on competition between multiple market-makers.

2.3 Microstructure Theory

Market microstructure theory typically relies on stylized market models to gain insights into the functioning of the market, and how different structural changes can impact

²In practice market-makers tend to be large banking firms.

price formation. There is an extensive literature in economics and finance on information-based market modeling. While it is impossible to provide a comprehensive bibliography here, O’Hara’s book [15] is an excellent starting point. Besides the canonical models of Glosten and Milgrom [7] and Kyle [11], a number of papers by Easley and O’Hara (for example [4]) consider the process of price adjustment as a reaction to information in securities markets. Grossman and Miller [8] consider the effects of market structure on liquidity. There is also a large empirical literature that is beyond the scope of this paper. We turn instead to a discussion of how to think about market quality and liquidity.

While the entire limit-order book provides important information about market quality and price formation, we can gain insight just by examining the bid and ask prices over time. The bid-ask spread serves as an indicator of market quality and liquidity, and, if it is assumed that the market functions reasonably efficiently, it can also proxy for hidden variables like the heterogeneity of information or beliefs about the valuation of a stock.

If we model all transactions as going through a single market-maker, then the spread compensates the market-maker for three different kinds of costs: (1) transaction costs, (2) inventory holding costs, and (3) adverse selection costs [15]. The market-maker must be compensated for the cost of doing business, which falls under the category of transaction costs. Even in the absence of transaction costs, the market-maker must be compensated for the risk she bears by holding inventory [1], so a spread will still arise.

The third type of cost mentioned above, and the focus of this paper, is the adverse selection cost borne by the market-maker in interacting with traders who potentially have better information available to them. This was first studied in detail by Glosten and Milgrom [7], who showed that transaction prices form a martingale under the assumption that a zero-profit market-maker sets bid and ask prices to the expected value of the stock given that a sell or buy order, respectively, is received. Das [3] extended this model by providing a practical algorithm for setting dollar-and-cent prices under some assumptions about the nature of the trading crowd. This paper considers different possible behaviors on the part of the market-maker and the trading crowd.

2.4 Agent-Based Financial Modeling

Agent-based modeling is assuming an increasingly important role in the study of microstructure. For example, Darley *et al* [2] used agent-based models to predict the impact of decimalization on the NASDAQ. LeBaron provides an excellent summary of much of the early work in agent-based computational finance [12]. At the same time, much computational modeling has been from the econophysics perspective. Econophysicists typically use low-intelligence models of agents and focus more on the dynamics of the collective behavior of these agents, trying to show how complex phenomena can arise from simple agents and simple interactions, using the tools of statistical mechanics [14; 9, *inter alia*]. This paper takes a middle ground – we assume low intelligence behavior on the parts of the trading crowd, but use this to formulate an interesting optimization problem for the market-maker, who is more likely to spend the resources on solving difficult problems.

3. THE MODEL

The market model is based on those in [3] and [7]. In all cases there is one market-maker, and a *trading crowd*.

3.1 Structure of trading:

Each trading *episode* is divided into n units of time, or trading periods. There is one security or stock in the market. For convenience, assume the existence of a true, or fundamental, value of the stock. In each episode, at time 0, this true value V is sampled from a normal distribution with mean μ_V and standard deviation σ_V (μ_V and σ_V are known to the market-maker). This represents an informational shock occurring before each trading episode begins. The true value remains constant for the rest of the episode.

At each trading period i , the market-maker sets bid and ask prices for one unit of stock, P_b^i and P_a^i . The market bid price P_b^i is then the maximum of the market-maker’s bid price and the prices in the limit buy order book, and the market ask price P_a^i is the minimum of the market-maker’s ask price and the prices in the limit sell order book. Only unit trade sizes are considered at each time period. Information can also be conveyed by the patterns and size of trades (Kyle’s model [11] is a canonical example), but the present work abstracts away from those considerations.

3.2 The trader model:

At each time period i , one trader is randomly selected from the (assumed infinite) trading crowd. This trader values the stock at $W^i \sim N(V, \sigma_W^2)$. If $W^i > P_a^i$ then the trader buys one unit of the stock. If $W^i < P_b^i$ then the trader sells one unit of the stock. When neither of these conditions hold, two different possibilities are considered. First, if the trading crowd is not permitted to place limit orders, then the trader places no trade. If the trading crowd is allowed to effectively compete with the market-maker by placing limit orders, then the trader selects a price P_l^i uniformly at random between P_b^i and P_a^i . Then a limit order is placed for one unit at the price P_l^i . The order is a buy order if $P_l^i < W^i$ and a sell order otherwise. This is a modification of the “zero-intelligence” model of Farmer *et al* [6], although in their case, traders can place limit orders in a wider range. Traders may not cancel limit orders.

3.3 The market-maker model:

The market-maker uses an algorithm which extends that developed by Das [3]. This section explains that algorithm and novel extensions to myopically optimizing market-makers and markets in which traders place limit orders. The key aspect of the algorithm is that the market-maker uses the information conveyed in trades to update her beliefs about the “true” value of the stock, and sets prices based on these beliefs. The market-maker maintains a probability density estimate over the true price of the stock. This estimate is maintained by assigning positive probabilities to discrete points that correspond to dollar-and-cent values in the range $[\mu_V - 4\sigma_V, \mu_V + 4\sigma_V]$.³ The density estimate is initialized by taking values of the normal pdf at all points in the range and normalizing the vector. While the initial prior is Gaussian, the MM’s beliefs do not remain Gaussian after the first update.

There are two key steps involved in the market-making algorithm. The first is the computation of bid and ask prices

³The actual number of standard deviations used can affect pricing near the ends of the distribution.

given a density estimate of the kind described above, and the second is the updating of the density estimate given the information implied in trades.

3.3.1 Calculating Prices

Assuming she has access to a density estimate of the form specified above, the market-maker can compute the expected profit she would make from any particular bid or ask price. The expected profit computations assume that the stock can be liquidated at the “true” value at the end of an episode. Here we explain the process for the bid price; the ask price computation is analogous. Let π_S denote profit from a market sell order being received. That is, π_S is the expected profit given that if any order is received, it is a market sell order. Equivalently, it is the expected profit at that time step if the market-maker’s ask price is infinite. Then expected profit at a time step will be the sum of π_S and an equivalently computed π_B . Dropping the superscripts i , because we are only considering one trading period:

$$E[\pi_S | P_b = x] = \sum_{y=V_{\min}}^{V_{\max}} \Pr(V = y) \Pr(\text{Sell} | P_b = x, V = y)(y - x)$$

Now, $\Pr(V = y)$ is known from the density estimate. The term remaining to be computed is $\Pr(\text{Sell} | P_b = x, V = y)$. A trader will only sell if she thinks the stock is overvalued, i.e the trader’s valuation is lower than the bid price, so it must be the case that $y + \mathcal{N}(0, \sigma_W^2) < x$. Therefore

$$\Pr(\text{Sell} | P_b = x, V = y) = \Pr(\mathcal{N}(0, \sigma_W^2) < x - y)$$




Figure 2: An example of the expected profit vector for the bid price. This is computed using the initial density estimate of the market-maker before any trades have occurred. The vertical lines show the myopically optimal (left) and ϵ -profit (perfectly competitive, right) prices.

For each possible bid and ask price, the market-maker can thus compute an expected profit. Figure 2 shows the form this vector typically takes for bid and ask prices. Glosten

MM Type	Profit	Spread	# Trades
Zero-profit	0.507 ± .076	0.61	80.3
Myopically opt.	4.438 ± .164	1.46	30.6
Zero-profit ±δ	4.674 ± .127	1.03	45.6

Table 1: Simulation results for the three different market-maker types. These results are averages from 100 episodes, each consisting of 100 trading periods, and each with an independently sampled “true” value. All differences are statistically significant at a 0.05 confidence level. Plus-minus numbers reflect 95% confidence intervals. Standard errors for spread and number of trades are trivial compared to the differences.

and Milgrom have shown that to at least break even in expectation, the market-maker must set the bid price lower than, and the ask price higher than, $E[V]$ [7]. Two types of market-maker considered in this paper are zero-profit market-makers and myopically optimizing market-makers.

A Zero-Profit Market-Maker.

In a perfectly competitive, frictionless environment, the market-maker’s equilibrium strategy is to set prices so as to obtain zero profit in expectation. This condition leads to a nice characterization of the bid and ask prices as $P_b = E[V|\text{Sell}]$ and $P_a = E[V|\text{Buy}]$ [7]. Of course, given that prices are actually quoted in some increments (like cents), the zero-profit strategy effectively becomes an ϵ -profit strategy and the market-maker quotes the closest integral value with non-negative expected profit.⁴

The zero-profit equations can be solved explicitly [3]. In the context of the model in this paper, the equations reduce to (for the ask price):

$$P_a = E[V|\text{Buy}] = \frac{1}{\Pr(\text{Buy})} \sum_{V=V_{\min}}^{V_{\max}} \Pr(N(0, \sigma_W^2) + V > P_a) V \Pr(V)$$

Then P_a is the (attractive) fixed point of the operator defined by the right-hand side.⁵ This fixed point can be computed very efficiently by starting to iterate from $E[V]$. The bid price can be computed in the same way.

The concept of the zero-profit (or ϵ -profit) market-maker can stand in as a proxy for efficient prices in the market as a whole. The fiction of a price-setting market-maker is convenient, but the same prices could be achieved by a variety of other means. This way of looking at price dynamics is useful from the systems perspective, although the algorithmic problems of market-making are interesting in and of themselves.

A Myopically Optimizing Market-Maker.

If the market-maker is a monopolistic price-setter, then she can set prices so as to make positive profits. The simplest such model is a market-maker who optimizes myopi-

⁴In terms of the expected profit vector described above, the zero-profit prices are the first non-negative expected-profit prices as we move downwards and upwards from $E[V]$ as P_b and P_a respectively.

⁵Technically P_a will be the ceiling of the fixed point.

cally, selecting the prices with highest expected profit at each trading period. The short-term optimal bid price is then the one that maximizes the expression $E[\pi_S|P_b = x]$ derived above. The maximum is well-behaved in the model presented here, with a single local and global maximum always occurring in practice for the bid and ask prices.

The selection of prices that attain these maxima guarantees that for any given trading period, the expected profit is maximized, but does not guarantee overall profit maximization in the sequential context. We show later that myopic optimization does not perform optimally over time by demonstrating a method that performs better. This leaves open the important algorithmic question of how to design improved, or perhaps optimal, algorithms for this problem.

When the market-maker can face competition through other traders placing limit orders, all bid prices lower than the best bid in the order book, and all ask prices higher than the best ask in the order book will yield zero expected profit for that trading period, since the probability of executing a trade at that price will be zero. In this case the myopically optimal strategy for the market-maker reduces to placing bids and asks that are just inside the spread (best bid plus one cent and/or best ask minus one cent) if they have positive expected profit. This is similar to parasitic strategies often employed by so-called “day traders.”

Note that it is even more difficult to think about an optimal sequential market-making strategy in the competitive framework because the market-maker is no longer guaranteed a monopoly over all trade executions.

3.3.2 A “widening the spread” heuristic

A heuristic strategy for making profit suggested in [3] is to compute the zero-profit bid and ask prices and then add (subtract) some value δ to the ask (bid) price. Some intermediate value of δ will optimize profits within this class of strategies, because widening the spread too much would lead to insufficient trading to make profits.

3.3.3 Updating the Density Estimate

The market-maker uses the Bayesian updating method described in [3]. All the points in the density estimate can be updated based on whether a buy order, sell order, or no order (equivalently, a limit order) was received. As an example, suppose a buy order was received.

$$\Pr(V = x|\text{Buy}) = \frac{\Pr(\text{Buy}|V = x) \Pr(V = x)}{\Pr(\text{Buy})}$$

The denominator is the same for all x and can thus be ignored and the updated vector renormalized. The second term in the denominator is just the prior (from the existing density estimate), and the conditional probability of a buy order can be computed as above from the normal distribution of noise in trader valuations. The density estimate is initialized at the beginning of each episode to be normal with mean μ_V and standard deviation σ_V .

4. SIMULATION RESULTS

4.1 Experimental Design

The experiments reported here follow the model described above. For the “true value” distribution, $\mu_V = 75$, $\sigma_V = 1$. The standard deviation of the distribution from which individual trader valuations are drawn, $\sigma_W = 0.2$. Each

episode consists of $n = 100$ trading periods. When market-makers use the spread-widening heuristic, they use $\delta = 0.10$. Results for those simulations in which the jump amounts are the same across episodes are averaged across 10 episodes, while results for simulations in which the jump amounts are themselves random are averaged over 100 episodes.

4.2 Interpreting the Results

We will see below that prices follow a two-regime behavior, which is one of the key findings reported in [3]. A price jump is followed by a period of high spreads, heterogenous information, and low volume of trading. Once this heterogeneity of information is resolved by the market-maker, trading settles into a regime of lower spreads and higher volume trading. We call these regimes the *price discovery* regime and the *efficient markets* regime respectively. Some of the numbers reported below are dependent on the interaction of these two regimes, and changing the length of each episode would affect the numbers. For example, shorter episodes may be dominated by the statistics of the price discovery regime with higher spreads and less trading.

We can examine market properties from the perspectives of at least three different players: (1) The average trader incurs a higher cost of trading when the spread is high. This is particularly so for traders who are not trading on information but instead need to change their positions for exogenous reasons such as liquidity constraints or hedging considerations. It can also be costly for traders who attempt to learn from prices, since trading is restricted in periods of higher spread. (2) The exchange itself would prefer low spreads and a high volume of trading. However, it is natural for the spread to be significantly larger immediately following an informational shock, and the exchange is interested in having the spread return quickly to a reasonable level quickly. (3) Self-interested market-makers want to optimize their profit. While the exact behavior of the spread that achieves this is not obvious, the market-makers would definitely want spreads significantly wider than the competitive spreads during the efficient market regime.

4.3 A Price-Setting Market-Maker

The first set of simulations are in a market in which the market-maker is a monopolistic price-setter. Table 1 shows that the spread-widening heuristic outperforms the myopically optimizing market-maker in terms of profit while at the same time providing more liquidity to the market and maintaining a lower average spread.

The zero-profit market-maker, who can be thought of as the aggregate outcome of a market with multiple competitive market-makers, of course provides the most liquidity to the market and quickly converges to a new regime of homogenous information, allowing for a greater volume of trading. The lower average spread and greater volume of trading can be seen from Table 1, and the fact that the market-maker converges to the homogenous information regime faster can be seen from Figure 3, which shows example behavior for two different price jumps (the model is symmetric, so the direction of the jump does not matter). The market-maker does make some profit over the course of the simulation, even though she is setting prices as competitively as possible. This is because trades can only occur at integral prices, whereas the truly zero-profit prices may fall between two dollar-and-cent values. This might be part of the solution




Figure 3: Dynamics of the spread for the three different market-maker types when there is an external (positive) shock of \$1 (left) or \$3 (right) to the “true” stock price. Spreads are averaged over 10 episodes.




Figure 4: Transaction price behavior for the three different market-maker types in one example paired episode when there is an external (positive) shock of \$1 to the “true” stock price. The episode is paired in the sense that the trader arriving at trading period i has the same valuation across the three different market-maker types.

to one of the questions of microstructure theory – where do the small profits come from that must exist to persuade market-makers to enter the game?

Figure 4 shows the actual price process in an example episode where the true value jumped up one dollar. Note that: (1) The first transaction does not occur for quite some time after the start of the episode. The market-maker sets the spread very wide, and then slowly narrows the spread as she learns that no one is willing to trade at such a high spread. If the amount of the jump had been greater, trading would have started earlier. The first few trades then collapse the market-maker’s density estimate into a much more concentrated region, allowing for smaller spreads and then considerably more trading. (2) The “bounce” between trades occurs partly because of the spread and partly be-

cause of the market-maker adjusting her beliefs in response to past trades. Once the market-maker’s beliefs have become concentrated, the former effect dominates the latter. The market is much more orderly and prices do not fluctuate as much in the competitive case.

4.3.1 The Exploration-Exploitation Tradeoff

The fact that a simple algorithm like the spread-widening heuristic outperforms myopic optimization shows that the myopic algorithm does not optimize over the sequential game. This may not matter in a competitive setting because there the market-maker has to compete for every trade, but when she has a monopoly on price-setting, she can do better than the myopically optimizing method.

Myopic optimization fails to achieve sequential optimality because prices serve two functions. In addition to being the profit-making mechanism, prices also convey information about the “true” value of the stock. A market-maker with narrower spreads can concentrate her density estimate more quickly than one with larger spreads. This will allow her to make more trades, and potentially to make more profits from future trades.⁶ The market-maker’s exploration-exploitation tradeoff can be thought of as a tradeoff between “price discovery” and “profit-taking.” The optimal strategy for a market-maker in this setting is uncharacterized – this is an interesting open problem.

4.4 Competition in the Limit Order Book

This section explores the effects on price dynamics of allowing traders to place limit orders using the model specified above. Figure 5 shows that the competition induced by limit-order placement on the part of traders leads to faster discovery of the new price for both zero-profit and myopically optimizing market-makers. This is because traders will come in and place limit orders early in the process, leading to an earlier start to the trading process, and therefore, more information becoming available earlier without the market-

⁶There is also a tradeoff between the probability of a trade occurring and the expected profit given that the trade does occur, but this is solved for in the profit maximization step.




Figure 5: Dynamics of the spread for markets with a perfectly competitive market-maker (left) and a myopically optimizing market-maker (right) with and without traders who can compete with the market-maker by placing limit orders. There is an external (positive) shock of \$1 to the “true” stock price. Spreads are averaged over 10 episodes.

maker having to actually make those trades.

In the zero-profit case there is little difference in the spreads once the market has converged to the homogenous information regime, since it would be hard to go lower than the 2-3 cent spreads which the zero-profit market-maker maintains even without competition. In a market with a myopically optimizing market-maker, the limit-order trading induces a lower spread in the homogenous information regime.

Table 2 shows that market quality is significantly improved in terms of both the average spread and number of trades per episode for the case of the myopically optimizing market-maker. For this experiment the “zero-profit plus-minus δ ” market-maker used $\delta = 0.05$ so that the profits could be comparable with that of the myopically optimizing market-maker, since the profits when using $\delta = 0.10$ were significantly lower.

MM Type	Profit	Spread	# Trades
Zero-profit	$0.480 \pm .110$	0.24	88.52
Myopically opt.	$2.910 \pm .184$	0.43	56.68
Zero-profit $\pm \delta$	$2.821 \pm .113$	0.32	73.57
None	N/A	0.64	43.36

Table 2: Simulation results when traders can also place limit orders. Results are averages from 100 episodes, each consisting of 100 trading periods with an independently sampled “true” value. δ is set to 0.05. The difference between the profits received by the myopically optimizing and zero-profit $\pm \delta$ market-maker is not statistically significant at the 0.1 confidence level, while all other differences are significant at the 0.05 confidence level. Plus-minus numbers reflect 95% confidence intervals. Standard errors for spread and number of trades are trivial compared to the differences.

4.5 The Absence of a Market-Maker

Many studies of market microstructure ignore the role of the market-maker, assuming prices form competitively through limit order placement. To evaluate the influence of market-making, one can simulate markets with only the trading crowd placing limit orders and market orders using the model described above. Markets with zero-profit market-makers are a limiting case with pricing as close to competitive as feasible, but Figure 6 shows that price discovery is faster and spreads are smaller even when there is a *myopically optimizing* market-maker. Table 2 also shows that the presence of the market-maker increases liquidity.

This shows that even a market-maker who is solely trying to optimize her own immediate profit in a competitive setting can improve the quality of the market. Thus, market-making can serve as an effective trading strategy for individual agents who do not possess superior information but are willing to learn from prices, and their presence in turn helps the process of price discovery. In efficient markets, we would expect such traders to come into existence on their own in light of the available opportunity for profit.

In practice, market-makers may be regulated by the market (by, for example, setting rules on how large the spread may be at any time, or penalizing market-makers for falling short on measures of market quality) or their behavior may be left up to the effects of competition.

5. DISCUSSION

From a market behavior perspective, this paper shows that market-makers can speed up the process of price discovery and lead to better market quality even when the market-makers are not heavily regulated. In practice we would expect market-makers to be more competitive, and therefore perform even better along these dimensions, than the myopically optimizing market-maker considered here.

From the algorithmic standpoint, this paper poses an important open problem – what is the optimal market-making algorithm for a monopolistic, price-setting market maker in




Figure 6: Dynamics of the spread for markets in which traders can place limit orders without a market-maker and with a myopically optimizing market-maker. There is an external (positive) shock of \$1 to the “true” stock price. Spreads are averaged over 100 episodes for the case without a market-maker and 10 episodes for the case with a market-maker.

the sequential context? Myopic optimization can be outperformed by simple algorithms, so it will be interesting to try and devise better algorithms for balancing the market-maker’s exploration-exploitation tradeoff. Initial experiments suggest that a good strategy may be to maintain artificially low spreads soon after a jump, and make up the losses incurred in the price-discovery regime by exploiting the quicker price discovery process in the homogenous information regime.

While the trading model of this paper is stylized, simple models have been shown to produce rich and interesting market behavior in many cases (for example, [6, 3, 16]). These models have qualitatively reproduced many real market phenomena. We hope the role and importance of market-makers will garner more attention in the algorithmic literature and in studies of novel electronic markets.

6. REFERENCES

- [1] Y. Amihud and H. Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8:31–53, 1980.
- [2] 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.
- [3] S. Das. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, 5(2):169–180, April 2005.
- [4] D. Easley and M. O’Hara. Time and the process of security price adjustment. *Journal of Finance*, 47:577–606, 1992.
- [5] E. Even-Dar, S. M. Kakade, M. Kearns, and Y. Mansour. (in)stability properties of limit order dynamics. In *Proceedings of the ACM Conference on Electronic Commerce*, 2006.
- [6] J. D. Farmer, P. Patelli, and I. I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences, USA*, 102(11):2254–2259, 2005.
- [7] L. Glosten and P. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.
- [8] S. Grossman and M. Miller. Liquidity and market structure. *Journal of Finance*, 43:617–633, 1988.
- [9] N. F. Johnson, P. Jefferies, and P. M. Hui. *Financial Market Complexity: What physics can tell us about market behavior*. Oxford University Press, UK, 2003.
- [10] S. M. Kakade, M. Kearns, Y. Mansour, and L. Ortiz. Competitive algorithms for VWAP and limit-order trading. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 189–198, 2004.
- [11] A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1336, 1985.
- [12] B. LeBaron. Agent-based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, 24:679–702, 2000.
- [13] A. Madhavan. Market microstructure: A survey. *Journal of Financial Markets*, pages 205–258, 2000.
- [14] R. N. Mantegna and H. E. Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge University Press, UK, 2000.
- [15] M. O’Hara. *Market Microstructure Theory*. Blackwell, Malden, MA, 1995.
- [16] M. Raberto, S. Cincotti, S. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:319–327, 2001.
- [17] R. A. Schwartz. *Reshaping the Equity Markets: A Guide for the 1990s*. Harper Business, New York, NY, 1991.
- [18] J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.