

The Effects of Market-Making on Price Dynamics

Sanmay Das

Dept. of Computer Science and Engineering
University of California, San Diego
La Jolla, CA 92093-0404
sanmay@cs.ucsd.edu

Abstract

This paper studies price properties in continuous double-auction markets in the presence of market-makers, agents with special responsibilities for maintaining liquidity and orderly price transitions. The effects of market-making are especially important immediately following exogenous price shocks. This paper analyzes the dynamics of prices around such events using simulation techniques and standard measures of market quality like the bid-ask spread. When the market-maker is a monopolistic price-setter, experiments show that myopic profit-maximization, apart from leading to poor market quality, is sub-optimal even for the market-maker in the sequential context. This observation leads to an interesting characterization of the monopolistic market-maker's exploration-exploitation dilemma as a tradeoff between price discovery and profit-taking. When traders can place limit orders, the presence of market-makers leads to more effective price-discovery and greater price stability under a simple model of limit-order placement.

1 Introduction

Most modern stock exchanges function through continuous double auctions. Traders can continually place orders to both buy and sell shares at different prices. There are two different types of orders, namely limit orders, which specify the price at which the trader is willing to buy or sell shares, but which may not get executed unless someone is willing to take the other side of the trade at that price, and market orders, which are orders to buy or sell immediately at the prevailing market price. Prices are decided based on the placement and execution of these orders.

Market microstructure is the study of how the specific rules governing trading affect price formation in markets. With the changing nature of trading in traditional stock markets, in which information is now more freely available and accessible and a large proportion of trades are initiated by automated trading systems, and the rapid increase in new kinds of electronic marketplaces which must decide on what kind of structure and rules to use, the study of microstructure has assumed increased practical importance.

Many aspects of exchanges and price dynamics have received a lot of attention recently from an algorithmic perspective. For example, Kakade *et al* [8] study algorithmic trading in terms of the problem of optimizing trade execution, and Even-Dar *et al* [3] examine the dynamics of order-books under different models of trading. This paper seeks to contribute to this growing literature by studying another of the key questions of market microstructure from an algorithmic and dynamic perspective: namely, what is the impact of *market-makers*, agents specially designated and/or employed by exchanges to provide liquidity and improve market quality?

While market-makers have traditionally been employed by large stock exchanges, many modern markets like prediction markets are now using market-makers to improve liquidity and market quality [14]. These markets are often novel and 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.

A particular focus of this paper is on price dynamics, rather than just equilibrium behavior. The efficient markets hypothesis, in its various forms, says that prices reflect all available information. I am concerned here with the question of *how* prices come to reflect the available information, and the processes that markets follow to incorporate new information into prices.

As an example of the type of question this paper seeks to answer within some specific market contexts, 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 informational signals about the value of a stock which the market-maker is not privy to. Typically, this happens when there is an informational shock that provides “inside information” to some traders.¹ The market-maker must set bid and ask prices to at least offset the adverse selection costs she incurs by trading with potentially better informed traders.

Sections 2 and 3 discuss research related to this problem in practice and in theory, respectively. Section 4 presents the stylized model that I use to study the problem, including the market-making algorithm, which is based on previous work [2]. Section 5 presents the main experimental results of the paper. In particular, I 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. I 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. Finally, Section 6 concludes with a discussion of the main results.

2 Microstructure Practicalities

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 [11, 10]. This is markedly different from much of economics and finance theory that explicitly abstracts away from the process of trading and assumes equilibrium pricing. Since 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 supply and demand of a stock. The insights provided by the study of microstructure can thus 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*. This section presents a brief overview of these types of orders and the role of the limit-order book. For the modeling purposes of this paper we are not concerned with the precise functioning of limit-order books in too much detail, but an excellent detailed description from a modern perspective can be found in the recent work of Kakade *et al* [8], and a more elaborate, albeit older, description is that of Schwartz [13].

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.”

How are these prevailing prices at which market orders trade determined? They come from the limit order book. Traders may also place orders of the type “Buy/Sell X shares at price Y.” These are known as

¹This does not have to be illegal. Paying for an analyst forecast can be thought of as a method of obtaining “inside information.”

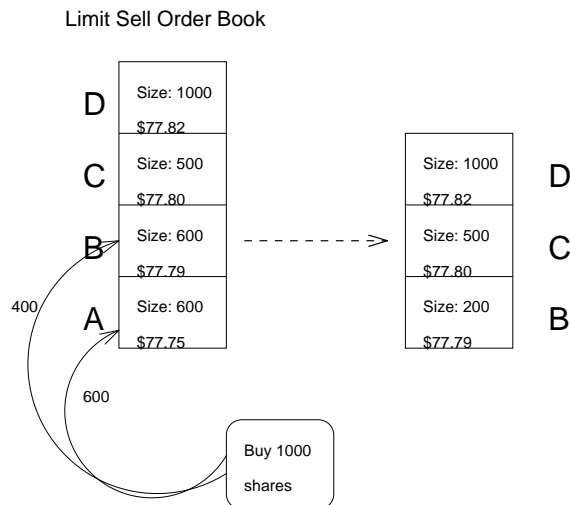


Figure 1: Example of a market buy order for 1000 shares executing against the limit sell order book. First, 600 shares execute at \$77.75 against Order A, then 400 shares execute against Order B at \$77.79. The volume weighted average price of this order execution is then \$77.766. Note that the remaining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

limit orders, and they are guaranteed price, but not execution. There are two limit order “book”s, the buy order book and the sell order book. The buy order book is prioritized in descending order of price and the sell order book is prioritized in ascending order of price. Orders at the same price are prioritized by arrival time. Assuming for the moment that the highest priced buy limit order is priced lower than the lowest priced sell limit order (this is justified below), the highest buy limit order and the lowest sell limit order constitute what are known as 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. If the size of the market order is greater than the size of the limit sell order, the limit order is removed from the book, and the remaining part of the buy is executed against the next order in the limit sell order book (with the process iterating if the next limit sell order is not sufficiently large to handle the entire remaining part of the market order). Any executed parts of the limit sell orders are taken off the books (if the buy order doesn’t take all of a limit sell order’s quantity, the remaining amount remains in the book). Figure 1 shows an example of how a market buy order for 1000 shares is executed and the states of the limit sell order book before and after the trade execution. The process for handling a market sell order through the limit buy order book is exactly analogous.

We can essentially ignore the cases when a limit buy order is at a price higher than (or equal to) the lowest limit sell order and when a limit sell order is at a price lower than (or equal to) the highest limit buy order by treating these as market orders. If only some portion of the new limit order crosses the other side of the order book, then that part is executed against the limit order on the other side, and the remaining portion enters the limit order book as it would have otherwise. Figure 2 shows an example of this.

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 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,

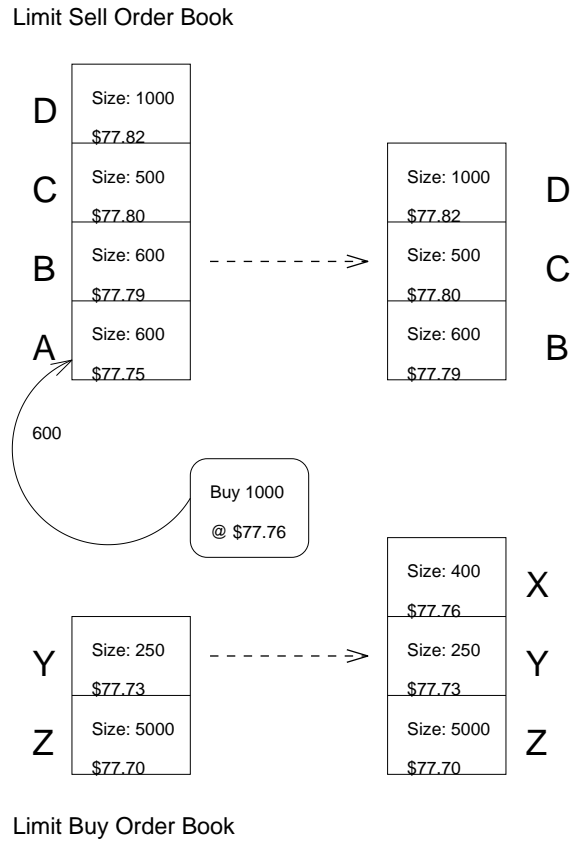


Figure 2: Example of a limit buy order for 1000 shares that crosses the current market ask price (lowest limit sell order price) executing against the limit sell order book. 600 shares execute at \$77.75 against Order A and the remaining 400 shares go into the limit buy order book because they no longer match or cross the ask price.

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.

In recent work, Even-Dar *et al* have shown that there can be surprising instabilities in price dynamics arising from limit-order trading when agents have relative valuations of stocks [3]. This can happen even without considering one of the main problems that market-makers are intended to solve – ensuring a smooth transition of prices when there is a large external information shock. I should note that the model considered here does not directly relate to the question of order book instability as posed by Even-Dar *et al*, partly because the entire trading crowd in this model can be thought of as having absolute valuations in their model, and partly because this paper is focused on price dynamics following an external informational shock that is known to the trading crowd but not to the market-maker. However, there is a region of overlap in the problems, and it would be interesting to consider the problem of whether market-makers could theoretically eliminate possible “butterfly effects.”

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.

2.3 Market-Makers on the NYSE and the NASDAQ

In thinking about the role of market-makers, it is instructive to look at two major exchanges that use different models. The NYSE employs a *specialist*² for each stock, and the specialist’s role includes matching orders between other traders, always ensuring that trades execute at the best quoted prices. The specialist also conducts the call auction that determines the opening price of the stock at the beginning of a trading day. The specialist is obligated to step in and trade against her own inventory when there is an imbalance in supply and demand. Therefore, the specialist is essentially obligated to always maintain bid and ask quotes herself. The specialist is carefully regulated by the exchange, which evaluates her performance and the characteristics of stock price movement on a regular basis.

The NASDAQ, by contrast, relies on competition between market-makers to maintain low spreads and desirable properties. Multiple firms are allowed to make markets in the same stock, and they are all obligated to continuously post two-sided (bid and ask) quotes. Trading on the NASDAQ is completely electronic, so market-makers do not themselves need to engage in matching prices across other people’s orders.

While there are many other practical aspects to market-making, I will now abstract from the practicalities to examine the major aspects of market-making from a theoretical perspective and discuss what insights these can give us into market behavior.

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 price formation. While the entire limit-order book provides important information about market quality and price formation, we can gain plenty of 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.

The three main determinants of the bid-ask spread can be characterized as (1) transaction costs, (2) inventory holding costs, and (3) adverse selection costs [11]. It is easiest to think about these costs when all

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

transactions are taking place through a single market-maker. The market-maker, as a source of liquidity, is compensated through the existence of the spread, because she can buy for less than she sells for. 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 bears risk by holding inventory, and must be compensated for this (see, for example, Amihud and Mendelsohn [1], so a spread will arise even in the absence of transaction costs.

The third determinant 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 [6], 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 [2] 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 builds directly on that work by considering different possible behaviors on the part of the market-maker and the trading crowd.

4 The Model

This section formally defines the model of trading used for the remainder of this paper. In all cases there is one market-maker, and various traders who are sometimes referred to as the *trading crowd*. The basic market model closely follows that in [2], which extended the classic model of Glosten and Milgrom [6], although this paper allows for more types of behavior on the parts of both market-makers and the trading crowd.

4.1 Structure of trading:

Since we are interested in the dynamics of prices following a shock to the true value, we assume that one trading *episode* is divided into n units of time, or trading periods. There is just 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 determined by being sampled from a normal distribution with mean μ_V and standard deviation σ_V , where the parameters μ_V and σ_V are known to the market-maker. This can be thought of as representing 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 [9] is a canonical example), but the present work abstracts away from those considerations.

4.2 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 $W^i = V + x^i$, where x^i is a sample from a normal distribution with mean 0 and standard deviation σ_W . 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 [4], although in their case, traders can place limit orders in a wider range. Traders may not cancel limit orders in the model presented here.

MM Type	Avg. Profit	Avg. Spread	Avg. No. 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. δ is set to 0.10. All differences are statistically significant at a 0.05 confidence level. The plus-minus numbers in the “Avg. profit” column reflect 95% confidence intervals. Standard errors for the other two columns are trivial compared to the differences.

4.3 The market-maker model:

The market-maker uses an algorithm developed in detail in [2], from which most of this section is derived (the myopically optimizing market-maker presented here is novel). The key aspect of the algorithm is that the market-maker uses the information conveyed in trades to update her beliefs about the “true” price 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 $[\mu_V - 4\sigma_V, \mu_V + 4\sigma_V]$. The density estimate is initialized by taking values of the normal pdf at all points in the range and normalizing the vector.

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

4.3.1 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. Here I 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

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

For each possible bid and ask price, the market-maker can thus compute an expected profit. Figure 3 shows the form this vector typically takes for bid and ask prices. Glosten 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, $E[V]$ [6]. 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 $P_a = E[V|\text{Buy}]$ [6]. Previous research has focused on explicitly solving these equations (which is more efficient than computing the entire vector of expected profits) and then changing the zero-profit prices to take other considerations such as inventory

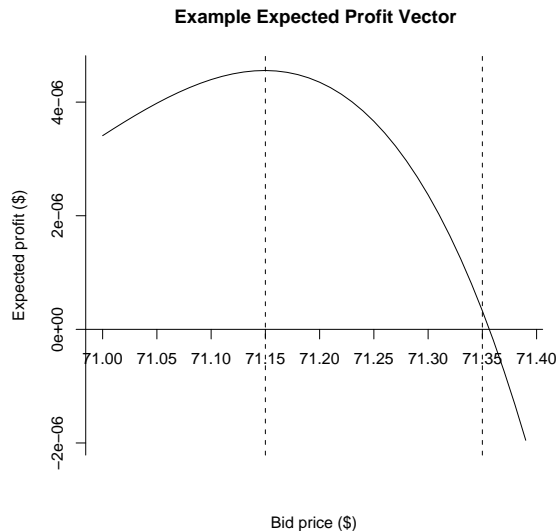


Figure 3: An example of the expected profit vector for the bid price. This is computed using the initial density estimate of the market-maker before any trades have occurred. The vertical lines show the myopically optimal (left) and zero-profit (perfectly competitive, and actually yielding some ϵ expected profit) prices (right). The expected profit keeps going down as the bid price increases further.

control or profit-making into account [2]. However, it is also easy to derive the zero-profit prices by taking the first non-negative expected-profit price as we move downwards and upwards from $E[V]$ as P_b and P_a respectively. Note that “zero-profit” is actually somewhat of a misnomer, because the market-maker does make some small expected profit from each trade, a point I return to below.

Another point worth noting here is that 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 myopically, selecting the prices with highest expected profit at each trading period. 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 the 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. In fact, I will show later that the myopic optimization technique 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.

4.3.2 Updating the Density Estimate

The market-maker uses the Bayesian updating method described in [2]. 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.

5 Simulation Results

5.1 Experimental Design

The experiments reported here closely 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. 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.

5.2 Interpreting the Results

The key measure of market quality used here is the bid-ask spread. If the spread is large, this means that the average trader is incurring a much higher cost of trading. This can be particularly costly for traders who 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. When bid and ask prices are set competitively, the spread also serves as a good measure of the heterogeneity of information in the market, because it is not artificially elevated for profit-making considerations.

From the perspective of the exchange itself, it would be preferable to have low spreads and a high volume of trading. However, it is natural for the spread to be significantly larger immediately following an informational shock, since information tends to be very heterogeneous in the market so the question becomes one of how quickly the spread returns to a reasonable level.

Another point to note is that prices follow a two-regime behavior in the simulations reported here, which is one of the key findings reported in [2]. A price jump (which for the purposes of this paper occurs at the beginning of each episode) is followed by a period of high spreads, heterogeneous 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. Some of the numbers reported below are critically dependent on the interaction of these two regimes, and changing the length of each episode would affect the numbers. Shorter episodes would be dominated by the statistics of the heterogeneous information regime with higher spreads and less trading, and longer episodes would be dominated by those of the homogeneous information regime. Nevertheless, these numbers will be important for understanding the algorithmic trading problem of the market-maker – it is better, however, to think of market properties in terms of the graphs below that show price behavior over time.

5.3 A Price-Setting Market-Maker

The first set of simulations are in a market in which the market-maker is a monopolistic price-setter. In addition to the two market-maker types described above, we consider a third type of market-maker, who

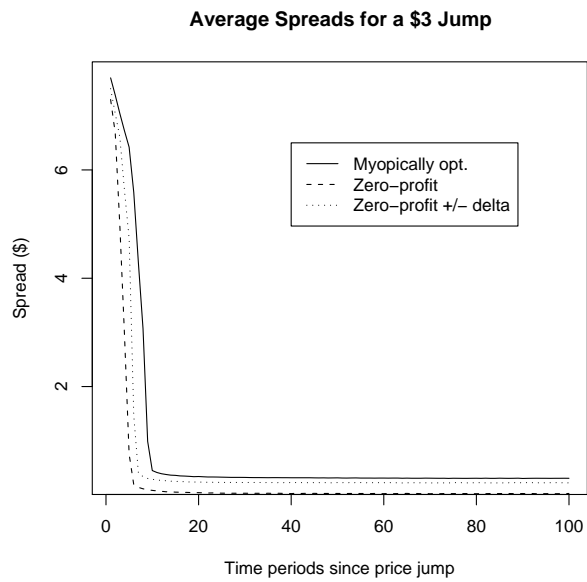
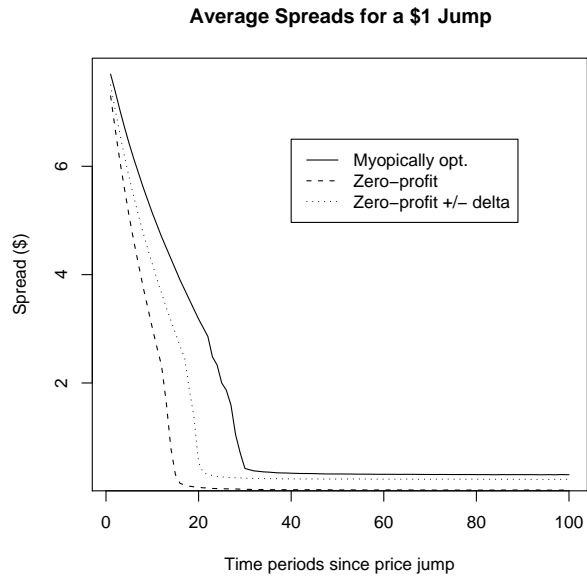


Figure 4: Dynamics of the spread for the three different market-maker types when there is an external (positive) shock of \$1 (left) or \$3 (right) to the “true” stock price. Spreads are averaged over 10 episodes.

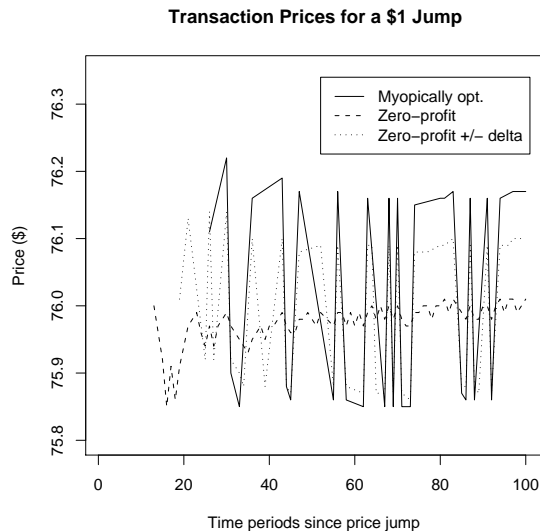


Figure 5: Transaction price behavior for the three different market-maker types in one example paired episode when there is an external (positive) shock of \$1 to the “true” stock price. The episode is paired in the sense that the trader arriving at trading period i has the same valuation across the three different market-maker types.

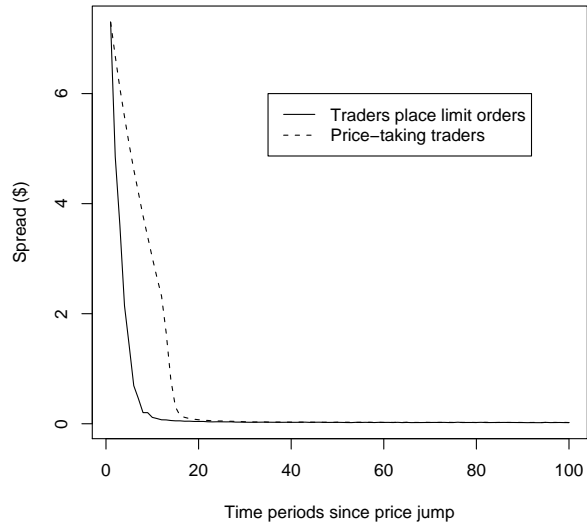
simply adds and subtracts a certain δ (in this case we set $\delta = 0.10$, or ten cents) to the zero-profit ask and bid prices respectively. Table 1 shows that this algorithm outperforms the myopically optimizing market-maker in terms of profit while at the same time providing more liquidity to the market and maintaining a lower average spread. This is discussed in more detail below in the context of the market-maker’s exploration-exploitation tradeoff.

The zero-profit market-maker, who can be thought of as the sum of 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 4, which shows example behavior for two different price jumps (note that the model is symmetric, so the direction of the jump does not matter). It is also interesting to note that the market-maker does make some profit over the course of the simulation as we would expect, 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 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 5 shows the actual price process in an example episode where the true value jumped up one dollar. There are various interesting things that help explain the results in Table 1 and Figure 4. The first transaction does not occur for quite some time after the start of the episode. This is because 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.

The “bounce” between trades is a phenomenon that occurs partly because of the spread and partly because of the market-maker adjusting her beliefs in response to past trades. Once the market-maker’s beliefs have become quite concentrated, the former effect dominates the latter. It is clear that the market is much more orderly and prices do not fluctuate as much in the competitive case (the zero-profit market-maker).

Average Spreads for a \$1 Jump (Zero-profit MM)



Average Spreads for a \$1 Jump (Myop. Opt. MM)

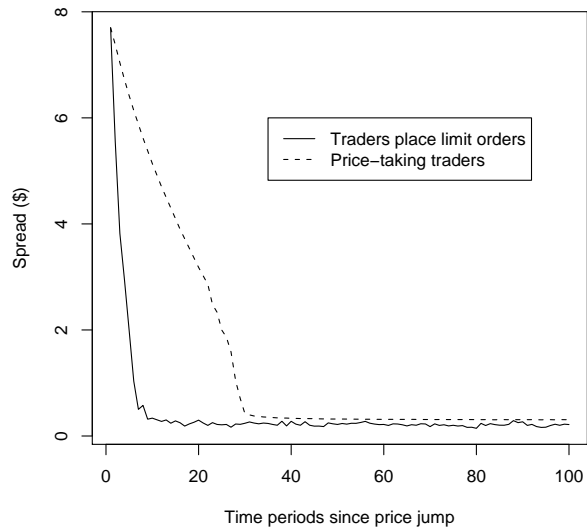


Figure 6: Dynamics of the spread for markets with a perfectly competitive market-maker (left) and a myopically optimizing market-maker (right) with and without traders who can compete with the market-maker by placing limit orders. There is an external (positive) shock of \$1 to the “true” stock price. Spreads are averaged over 10 episodes.

5.3.1 The Exploration-Exploitation Tradeoff

Let us consider the problem strictly from the perspective of a profit-maximizing market-maker, without worrying about market properties for the moment. The fact that a simple algorithm like the “zero-profit plus-minus δ ” method outperforms the myopically optimizing market-maker shows that there is something problematic in the myopic algorithm in terms of optimizing the outcome 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 certainly do better than the myopically optimizing method.

Why does myopic optimization fail to achieve sequential optimality? The main reason arises from the fact that prices serve two functions – they both allow the market-maker to make profits, and the execution prices enable her to learn more about the information out there in the world, i.e. 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.

Note that there might be an individual 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. What is not taken into account in that step is the tradeoff between the information gained from that trade and how it might help in future pricing decisions, and the profit from that trade. If we characterize the first of these as the market-maker’s “price discovery” role, and the second as her “profit-taking” desire (which is her desire to cash in right now on her beliefs), we can see that the market-maker’s exploration-exploitation tradeoff is also the 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. It is possible that the market-maker should try to learn the new true price early by maintaining low spreads, and then exploit her information later, or the strategy may be something altogether different. Understanding and characterizing the optimal profit-making strategy is an important direction for research. Of course, the simple model does not consider strategic traders, who may themselves learn from prices instead of having fixed valuations. This would make the problem even harder, and would be similar in flavor to the model of Kyle [9] and its descendants [7, 5, inter alia], although they typically hold market-maker behavior fixed by competitive considerations and optimize for the traders.

5.4 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 6 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-maker having to actually make those trades.

In the zero-profit case there is very little difference in the spreads once the market has converged to the homogenous information regime, since it would be very hard to go lower than the 2-3 cent spreads which the zero-profit market-maker maintains even without competition. However, in a market with a myopically optimizing market-maker, the limit-order trading induces a lower spread in the homogenous information regime, as can be seen in Figure 7. The average spread fluctuates more when traders can place limit orders because there is more of a difference in behavior across episodes, due to the large range of limit orders that can be placed by traders.

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.

5.5 The Absence of a Market-Maker

Another question that is important from the perspective of market quality is that of how much the market-maker really helps. In fact, many studies of market microstructure ignore the role of the market-maker

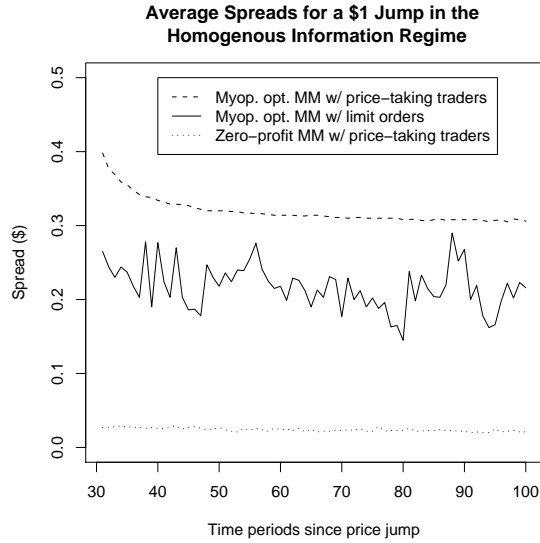


Figure 7: A closer look at average spreads in the homogenous information regime for the myopically optimizing market-maker with and without competition, as compared to the perfectly competitive (zero-profit) market-maker.

MM Type	Avg. Profit	Avg. Spread	Avg. No. 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 for the three different market-maker types in the setting where traders can also place limit orders. These results are averages from 100 episodes, each consisting of 100 trading periods, and each 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. The plus-minus numbers in the “Avg. profit” column reflect 95% confidence intervals. Standard errors for the other two columns are trivial compared to the differences.

completely, assuming prices will form competitively through traders placing limit orders.

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. This section compares spreads in these markets with spreads in a market with a myopically-optimizing market-maker who competes with the traders who place limit orders. Markets with zero-profit market-makers are in some senses a limiting case with pricing as close to competitive as feasible, but Figure 8 shows that price discovery is faster and spreads are smaller when there is a myopically optimizing market-maker present. Table 2 also shows that the presence of the market-maker increases liquidity.

This shows that even a market-maker who is just trying to optimize her own immediate profit in a competitive setting can improve the quality of the market while at the same time earning profit. 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 can help the process of price discovery, especially following informational shocks. In efficient markets, we would expect such traders to come into existence on their own in light of the available opportunity for profit.

Of course, in practice actual market-makers will fall somewhere between the myopic optimizers and the ϵ -profit makers described in this paper. Where exactly on the spectrum they fall or should fall can 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 over time) or left up to the effects of competition.

6 Discussion

From a market behavior perspective, this paper shows that the presence of 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 the sequential context? Myopic optimization can be outperformed by relatively simple algorithms, so it will be interesting to try and devise a better algorithm for balancing the market-maker’s exploration-exploitation tradeoff and understanding what this algorithm would imply for market properties.

These conclusions are, of course, subject to the caveat that this paper presents a stylized market model, and in particular the trading model is quite simple. However, simple trading models have been shown to produce rich and interesting market behavior in many cases (for example, [4, 2, 12]) and there are valuable insights to be gained from studying market properties in these models. In particular, the role and importance of market-makers will hopefully garner more attention in the algorithmic literature and in studies of novel electronic markets.

References

- [1] Y. Amihud and H. Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8:31–53, 1980.
- [2] Sanmay Das. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, 5(2):169–180, April 2005.
- [3] Eyal Even-Dar, Sham M. Kakade, Michael Kearns, and Yishay Mansour. (In)Stability properties of limit order dynamics. In *Proceedings of the ACM Conference on Electronic Commerce*, 2006.
- [4] J. Doyne Farmer, Paolo Patelli, and Ilija I. Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences, USA*, 102(11):2254–2259, 2005.
- [5] F.D. Foster and S. Viswanathan. Strategic trading when agents forecast the forecasts of others. *The Journal of Finance*, 51:1437–1478, 1996.

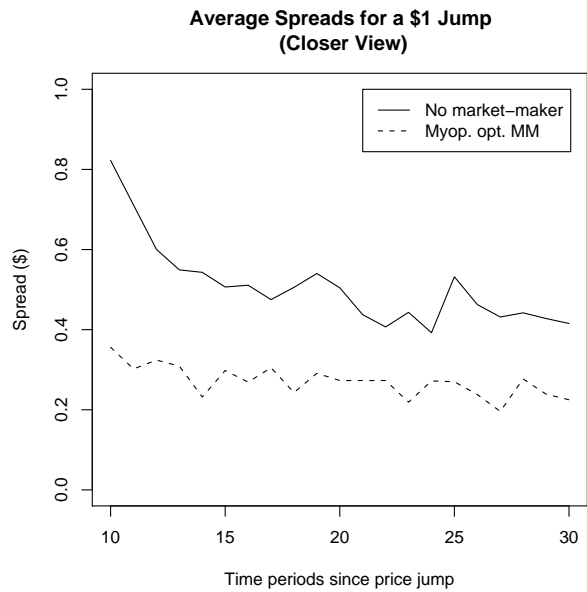
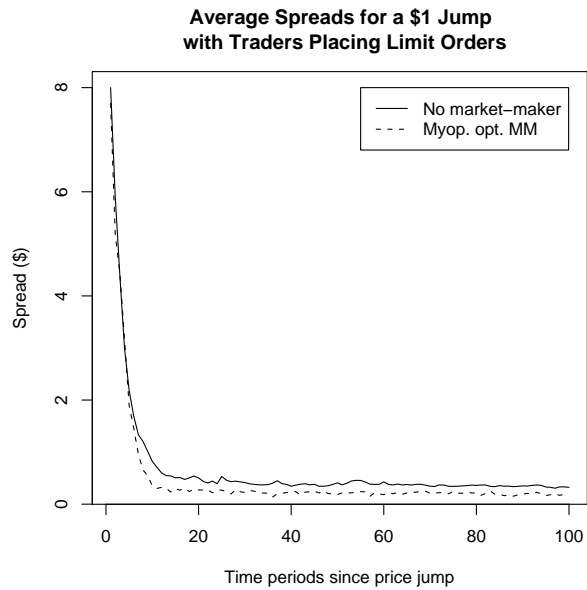


Figure 8: Dynamics of the spread for markets in which traders can place limit orders without a market-maker and with a myopically optimizing market-maker. There is an external (positive) shock of \$1 to the “true” stock price. Spreads are averaged over 100 episodes for the case without a market-maker and 10 episodes for the case with a market-maker.

- [6] 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.
- [7] C.W. Holden and A. Subrahmanyam. Long-lived private information and imperfect competition. *The Journal of Finance*, 47:247–270, 1992.
- [8] Sham M. Kakade, Michael Kearns, Yishay Mansour, and Luis Ortiz. Competitive algorithms for VWAP and limit-order trading. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 189–198, 2004.
- [9] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1336, 1985.
- [10] A. Madhavan. Market microstructure: A survey. *Journal of Financial Markets*, pages 205–258, 2000.
- [11] M. O’Hara. *Market Microstructure Theory*. Blackwell, Malden, MA, 1995.
- [12] M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299:319–327, 2001.
- [13] Robert A. Schwartz. *Reshaping the Equity Markets: A Guide for the 1990s*. Harper Business, New York, NY, 1991.
- [14] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.