

**The “Make or Take” Decision in an Electronic Market:
Evidence on the Evolution of Liquidity**

Robert Bloomfield, Maureen O’Hara, and Gideon Saar*

First Draft: March 2002
This Version: August 2002

***Robert Bloomfield (rjb9@cornell.edu) and Maureen O’Hara (mo19@cornell.edu) are from the Johnson Graduate School of Management, Cornell University. Gideon Saar (g Saar@stern.nyu.edu) is from the Stern School, New York University, and is currently a Visiting Research Economist at The New York Stock Exchange. Financial support for this project was obtained from New York University's Salomon Center for the Study of Financial Institutions.**

The “Make or Take” Decision in an Electronic Market: Evidence on the Evolution of Liquidity

Abstract

This paper uses experimental asset markets to investigate the evolution of liquidity in an electronic limit order market. Our market setting includes salient features of electronic markets, as well as informed traders and liquidity traders. We focus on the strategies of the traders, and how these are affected by trader type, characteristics of the market, and characteristics of the asset. We find that informed traders use more limit orders than do liquidity traders. We also find that liquidity provision shifts over time, with informed traders increasingly providing liquidity in markets. This evolution is consistent with the risk advantage informed traders have in placing limit orders. Thus, a market making role emerges endogenously in our electronic markets.

The “Make or Take” Decision in an Electronic Market: Evidence on the Evolution of Liquidity

Electronic markets have emerged as popular venues for the trading of a wide variety of financial assets. Stock exchanges in many countries including Canada, Germany, Israel, and the United Kingdom have adopted electronic structures to trade equities, as has Euronext, the new market combining eight former European stock exchanges. In the United States, Electronic Communications Networks (ECNs) such as Island, Instinet, and Archipelago use an electronic order book structure to trade as much as 45% of the volume on Nasdaq. There are now several electronic systems trading corporate bonds (e.g., eSpeed) and government bonds (Govpix), while, in foreign exchange, electronic systems such as EBS and Reuters dominate the trading of currencies. Eurex, the electronic Swiss-German exchange, is now the world’s largest futures market, and with the opening of the new International Securities Exchange, even options now trade in electronic markets.

Many such electronic markets are organized as electronic limit order books. In this structure, there is no designated liquidity provider such as a specialist or a dealer; instead, liquidity arises endogenously from the submitted orders of traders. Traders who submit orders to buy or sell the asset at a particular price are said to “make” liquidity, while traders who choose to hit existing orders are said to “take” liquidity. The spread and price behavior in such markets thus reflect the willingness of traders to supply and demand liquidity.

In this paper, we use an experimental market setting to investigate the evolution of liquidity in an electronic limit order market. Our market setting possesses the salient features of electronic markets: continuous trading, a visible “book” of orders, price-time order priority rules, instantaneous trade reporting rules, order cancellation capabilities, and both limit order and

market order functionality. While many experiments have used continuous double-auction market similar to the electronic markets we investigate (see the review by Sunder [1995]), our experiment is the first to focus primarily on the provision and use of liquidity in such markets. Our experimental market contains informed traders who have superior information and liquidity traders who face both large and small liquidity needs. We manipulate both the prior distribution and the realizations of security values. These manipulations allow us to analyze market behavior in ways unavailable in actual markets. In particular, we can analyze explicitly the strategies of informed and liquidity traders, and we can determine the factors that influence traders' make or take decisions.

Our particular focus in this paper is on three questions. First, how do informed and liquidity traders differ in their provision and use of market liquidity? Second, how do characteristics of the market, such as depth in the book or time left to trade, affect these strategies? And, third, how do characteristics of the underlying asset such as asset value volatility affect the provision of market liquidity? Addressing these questions allows us to provide insights not only into the functioning of electronic markets, but into the emergence of market liquidity as well.

Numerous authors in finance have examined aspects of these questions both theoretically and empirically, and there has also been related work in the experimental literature. Theoretical analyses of limit orders include Cohen, Maier, Schwartz, and Whitcomb [1981]; Rock [1990]; Angel [1994]; Glosten [1994]; Kumar and Seppi [1994]; Chakravarty and Holden [1995]; Parlour [1998]; Harris [1998]; Foucault [1999]; Parlour and Seppi [2001]; and Foucault, Kadan, and Kandel [2001]. Empirical studies of specific limit order markets include Biais, Hillion, and Spatt [1995]; Hollifield, Miller, and Sandas [1999]; Ahn, Bae and Chan [2001]; and Hasbrouck

and Saar [2001]. In general, these analyses have provided useful characterizations of limit order behavior, but the complexity of traders' decision problems has required selectivity in what aspects of trader or market behavior can be considered.

Our analysis provides a number of important new results. Of special significance, we find that informed traders actively submit limit orders. Indeed, both trader types use limit orders and market orders, but informed traders tend to use more limit orders than do liquidity traders. This behavior contrasts with the common assumption in the theoretical literature that informed traders only take liquidity, and do not provide it. One consequence of this behavior is that the book of orders has information content.

What we find particularly intriguing is that liquidity provision changes dramatically over time, and the key to this evolution is the behavior of the informed traders. When trading begins, informed traders are much more likely to take liquidity, hitting existing orders so as to profit from their private information. As prices move toward true values, the informed traders shift to submitting limit orders. This shift is so pronounced that towards the end of the trading period informed traders on average trade more often with limit orders than do liquidity traders. This has the intriguing implication that informed traders provide liquidity in various market conditions even as they speculate on their information. Liquidity traders who need to buy or sell a large number of shares, on the other hand, tend to use more limit orders early on, but as the end of the trading period approaches switch to market orders in order to meet their targets.

The informed traders also seem to change their strategies depending on the value of their information. When that value is high, informed traders tend to use more market orders in order to realize trading profit before prices adjust. When the value of their information is low, they move

very quickly to assume the role of dealers and trade predominantly by supplying limit orders to the market.

This dual role for the informed, acting as both traders and dealers, highlights the important ways that information influences markets. While it is the trading of the informed that ultimately moves prices to efficient levels, the superior information of the informed also makes these traders better able to provide liquidity to other traders in the market. Thus, unlike in theoretical models where the informed stop trading once their information is incorporated into prices, we find that the informed now profit further by taking on the role of liquidity providers and essentially earning the spread. In a symmetric information world, Stoll [1978] argued that the market maker would be a trader who was better diversified than the others and thus better able to bear risk. We show that in an asymmetric information setting, it is the informed traders who ultimately have the risk advantage because they know more about where the price should be. Thus, a market-making role arises endogenously in our electronic markets, adopted by traders for whom the risk of entering a limit order is lower than it is for other traders.

Our analysis may suggest why it is that electronic markets have been so successful in competing with more traditional market structures. Even in the presence of information asymmetry, the traders themselves will provide liquidity, eschewing the need for a formal, and typically more expensive, liquidity provider. While it is possible that such endogenous liquidity will dissipate in more uncertain market conditions, those same conditions make it difficult for designated liquidity providers to do much either.

The paper is organized as follows. In the next section we discuss the economic theory regarding limit order markets, with a particular focus on the factors affecting traders' order decisions. This section also sets out the questions we will address, and it provides a rationale for

why we use an experimental methodology in this research. Section 3 then describes our experimental markets and manipulations. Section 4 then presents our results. The paper's final section is a conclusion.

2. The Nature of Limit Order Markets

In an electronic market, traders face a number of choices in formulating their trading strategy. Certainly, a basic choice is whether to make or take an order. A trader makes an order by placing a limit order to buy or sell the asset at a specific price; a trader takes an order by agreeing to trade as the counter-party to an existing limit order. This latter trading strategy essentially corresponds to trading via a market order. While this decision can be thought of as “how” to trade, traders also must decide “when” to trade. A trader wishing to transact multiple shares can do so quickly, or she can spread her orders out. The trader can opt to trade early in the day, at the last minute, or at any point in between. Of course, in an electronic market deciding when to trade is also affected by the presence or absence of counter-parties wishing to trade. Finally, the trader faces the related decision of “what” to trade. Is she a buyer, a seller, or sometimes both? In an electronic market, each of these decisions affects not only the trader's individual profit and loss, but the behavior of the market as well. This latter linkage arises because liquidity is endogenous in an electronic market, arising solely from the trading strategies and collective behavior of the traders in the market.

While there is a large literature in market microstructure analyzing the trading process, the specific literature looking at trader strategies in electronic limit order markets is still fairly small. This paucity reflects the difficulty of characterizing how, when, and what to trade when the market outcome attaching to individual strategies depends upon the collective strategies of all

other market participants as well. This trading problem is further complicated if some traders are better informed about the security's true value than others. The complexity of the trading environment, combined with the inter-dependence of traders' decisions, makes characterizing a trader's optimal order strategy quite difficult; adding in asymmetric information makes the problem generally intractable.

Most theoretical studies make their analyses tractable by imposing highly restrictive assumptions. These assumptions raise concerns about the robustness of their conclusions. We use experimental markets to test the robustness of predictions derived from restricted models, and to shed light on behavior in less restrictive settings. We impose rigorous experimental controls that allow us to attribute our experimental results unambiguously to variables that are important in theoretical work. For example, to investigate the effects of asset-value volatility on the submissions strategies of traders, we compare trading of high-volatility assets with trading of low-volatility assets. Because all other aspects of the markets are the same, comparing outcomes between the two markets characterizes the specific effects of volatility on market behavior. An obvious advantage of this approach is that traders are allowed to pursue whatever equilibrium strategies they prefer; what matters is simply how these strategies differ with the treatment variable. Perhaps equally important, experimental markets provide for multiple replications, allowing us to focus on the typical equilibrium outcome, and not merely on an outcome that is theoretically possible albeit highly unlikely.

The first stream of literature motivating our experiment achieves tractability by making restrictive assumptions about the behavior of informed traders, or by ignoring such traders completely. For example, the early literature looking at limit order markets focused on the trade-off between the immediate execution of taking the limit order versus the better price, and

uncertain execution, of making a limit order. Cohen, Maier, Schwartz and Whitcomb [1981] developed a “gravitational pull” model of limit orders to explain when a trader would submit a limit order as opposed to a market order (the functional equivalent of taking a limit order). These authors showed that as spreads narrow, the benefits of the better price available to limit order traders decreases, causing more traders to prefer the certain execution of the market order. As traders shift from limit orders to market orders, however, the spread widens, thereby increasing the attractiveness of the limit order price improvement potential. Thus, a trader’s decision regarding how to trade involves a dynamic balancing of the relative costs of price improvement and execution risk. However, Cohen, et al. ignore the role of informed traders in their market.

Rock [1990], Glosten [1994], and Seppi [1997] explicitly incorporate informed traders into their models, but assume that they always enter market orders instead of limit orders. This research allows a number of insights into the role of the “winner's curse” problem of limit order execution. If there is asymmetric information between traders, then limit order submitters may face an adverse execution risk: limit orders will more likely execute when they generate a loss to the limit order submitter.

Because the results and tractability of these models depend critically on the assumptions about informed traders, the first goal of our experiment is to examine behavior when these assumptions are relaxed. We therefore create a setting in which both liquidity and informed traders can choose between limit and market orders.

Another stream of literature examines how both liquidity and informed traders choose between limit and market orders, and makes the settings more tractable by exogenously imposing market characteristics (such as the state of the limit order book) affecting those decisions. The decisions are still quite complex. Consider, for example, the problem facing an informed trader.

The informed trader would like to profit from his information, and this suggests trading as frequently as possible. But rapidly taking limit orders will lead prices to quickly converge to full information levels. Alternatively, submitting a limit order or a series of limit orders might allow the trader to better hide his information, and to trade at better prices. But it does so by delaying trading, and exposes the trader to execution risk. If there are other informed traders, then this strategy may prove sub-optimal, in part because the clustering of orders on the book may signal the presence, and value, of new information. And if liquidity traders act strategically, they may delay trading to allow the competition of the informed to reveal these new prices.

Angel [1994] and Harris [1998] provide some predictions on how informed traders will behave. They argue that informed traders are less likely to use limit orders than are liquidity traders. Furthermore, informed traders are more likely to use market orders if the realized asset value is farther away from its expected value. This preference reflects the desire of informed traders to capitalize on their private information.

Harris [1998] also predicts that liquidity traders needing to meet a target will start by using limit orders, and then switch to market orders as the end of trading (their "deadline") approaches. A similar prediction applies to the informed traders: the likelihood of submitting a limit order decreases with time until the end of trading (when their information is revealed). In both cases, more time provides traders with flexibility to design a limit order strategy that avoids paying the spread.

To test these predictions, our experiment includes liquidity traders who are forced to buy or sell some number of shares before the market closes. We manipulate the extremity of realized security values relative to the prior expected value, as a way of manipulating the value of the

informed traders' information. We also examine trader behavior separately at different points during the trading period to test the predictions with respect to time.

A third stream of literature constructs more complete equilibria in which key market attributes (such as bid-ask spreads and book depth) arise endogenously. These dynamic equilibrium models allow traders' optimal strategies to depend on conjectures of other traders' strategies. To simplify the analysis, however, traders solve static problems in which they are allowed to take only one action (i.e., submitting a market or a limit order without the ability to return to the market and update their strategies).

Foucault [1999] uses such a model to predict a higher submission rate of limit orders by traders when true value volatility is greater. Since traders are unable to cancel their limit orders, higher volatility increases the likelihood that their limit orders will become mispriced. The greater risk of being picked off leads traders to price the limit orders less aggressively. This increases the spread in the market, which makes market orders more expensive and decreases their proportion in the order flow.

Parlour [1998] shows how traders' decisions are influenced by the (endogenously-determined) state of the limit order book. Her analysis focuses on "crowding out" that arises due to the time priority of orders already in the book. Thus, Parlour's model predicts that depth at the best price on the same side of the book decreases the likelihood of submitting a limit order, while depth at the best price on the opposite side of the book increases this probability. However, neither Foucault nor Parlour incorporate in their models traders with private information about the security. We test the implications of these two studies by manipulating the volatility of security value, and also by measuring the depth at the best prices in the limit order book.

3. Experimental Design

We now describe the nature of our experiment and the specific features of our markets. As a useful preliminary, we note the following definitions. A *cohort* is a group of six traders who always trade together. A *security* is a claim on a terminal dividend, and is identified by the value of the security and the traders' liquidity needs (described below). A *trading period* is a time interval during which traders can take trading actions. A *session* is a 75-minute period during which traders participate in a series of markets. Unless otherwise indicated, all prices, values and winnings are denominated in laboratory dollars (\$), an artificial currency that is converted into US currency at the end of the experiment.

3.1. Experimental Design

Basic Design. We seek to examine how the trading behavior of informed and uninformed traders differs with the volatility of security value, the extremity of realized value from the prior expected value, with elapsed time, and with the depth of the book. Our experiment includes eight cohorts of 6 traders, for a total of 48 participants.

To manipulate information, two informed traders were told the true security value before trading in the security began. Four liquidity traders were not told the security value until trading in the security has ended.

To manipulate volatility, we altered the distribution of security values. In a high-volatility setting, traders are told that values are distributed approximately uniformly over the interval from 0 to 50 laboratory dollars. In a low-volatility setting, traders are told that values are distributed according to a truncated bell-shaped distribution (over the interval 0 to 50) with a mean of 25

laboratory dollars and a standard deviation of 5. In both cases, the expected value is the same; only the variance about that expectation is different. Each cohort trades 10 securities in the high-volatility setting and 10 securities in the low-volatility setting.

To manipulate extremity, we presented traders with high-extremity realizations that were at least \$15 from expected value, and low-extremity realizations that were no more than \$7 from expected value.

To manipulate elapsed time during trading for a security, we distinguished between decisions made at eight fifteen-second intervals during the 120 seconds of trading in each security.

Taken as a whole, our experiment uses a fully factorial repeated-measures design, with the following factors: trader type (informed, large liquidity trader, small liquidity trader), volatility (high, low) extremity (high, low), replication (there are three securities in each volatility x extremity combination), time (arbitrarily broken into eight 15-second periods), and cohort (eight cohorts of six traders each). Trader type and cohort membership are manipulated across traders, and all other factors are manipulated within traders.

Controls. The experiment also includes controls to ensure the treatment effects are not driven by differences that are not the focus of our study. To eliminate possible effects of minor differences in security, each cohort traded 12 securities that have identical deviations from the prior expected value of \$25 (see Table 1). Thus, tests of extremity and volatility allow us to compare outcomes across cohorts for securities that are identical in all key respects. To ensure that the total distribution of security values in each setting was distributed as indicated to traders, we also included eight additional securities with relatively extreme values (for the high-volatility

setting) or relatively central values (for the low-volatility setting). However, we did not include these securities in our analyses.

The experiment also controls the order of securities and treatments. Four of the cohorts traded first in the high-volatility setting and then in the low-volatility setting, while the other four traded in the opposite order. All cohorts traded the securities in exactly the same order. We also altered the sign of deviations from the prior expected value of \$25 across securities and across cohorts.

We count on the random assignment of participants to trader types to minimize the possibility that differences across trader types are driven by individual differences.

3.2 *Trading*

Market activity in each security takes place during two periods. During a “pretrading” period, traders have the opportunity to enter orders, but no trades are executed. During a main trading period, traders can continue to enter orders, and can also take other traders’ orders. We included the pretrading period to allow the order book to be full at the beginning of the main trading period.

Pretrading. Pretrading lasts 30 seconds, during which traders can enter bids (orders to buy one share at a chosen price) and asks (orders to sell one share at a chosen price). Traders can delete their orders at any time during the period. Traders can enter as many bids and asks as they wish, but cannot enter bids or asks that would result in one of their own outstanding bids having a price equal to or greater than the price of one of their own outstanding asks. All bids and asks must have integer prices between 0 and 50, inclusive. Traders are allowed to enter bids at prices above the lowest outstanding ask, and can enter asks at prices less than the highest bid. These

“crossing” orders are dealt with as discussed below. No trade takes place during the pre-trading period.

At the end of the pretrading period, the order book is purged of crossing orders in the following way. If the highest bid crosses with the lowest ask, the more recent of the two orders is deleted from the book. This process is repeated until the high bid is less than the low ask.

Main Trading. Main trading, which lasts 120 seconds, is exactly like pretrading, with two exceptions. First, traders are allowed to take other traders’ bids and asks. Traders take an ask by clicking a “buy 1” button, which allows them to buy one share at the lowest current asking price. Traders take a bid by clicking a “sell 1” button, which allows them to sell one share at the highest current bid price. Taking an ask is equivalent to entering a market buy order, while taking a bid is equivalent to entering a market sell order. Older limit orders are executed first. Second, traders are not allowed to enter limit orders that cross with existing limit orders from other traders. In other words, there are no "marketable limit orders," and immediate execution is achieved by submitting market orders (i.e., taking existing limit orders in the book).¹

Market Transparency. As soon as a trader enters an order, the order is shown on every trader’s computer screen, indicating that an unidentified trader is willing to buy or sell one more share at the posted price. As shown in Figure 1, the screen includes two graphs showing market activity. The left side of each graph shows every price at which an order has been posted (shown in green for the highest bid and lowest ask price, and yellow for other prices), and the number of shares posted at that price (shown by the number to the left of the graph). The right side of each graph shows every price at which the trader has personally posted an order, and the number of

¹ More specifically, a limit buy order can only be submitted at prices below the best ask price in the book, and a limit sell order can be submitted at prices above the best bid price in the book (so that the market cannot be "crossed" or "locked").

shares that the trader has posted at that price.² The center of each graph also includes a solid red line indicating the highest bid or lowest ask entered by any trader, and a solid green line indicating the highest bid or lowest ask entered by that particular trader.

Reporting of orders during trading is similar to reporting during pretrading. All trades are reported immediately to all traders, indicating the price and the trade direction (whether the trade involved a market buy and an ask or a market sell and a bid).

3.3 *Trader Types*

The market includes three types of traders. Two informed traders know the true value of the security, which they learn before trading begins. The remaining four traders have a trading “target” to meet before trading is complete. One trader’s target is to sell 20 shares; another’s is to buy 20 shares; another’s is to sell five shares, and another’s is to buy five shares. We refer to the first two as large liquidity traders and the last two as small liquidity traders. At the end of trading, liquidity traders are assessed a penalty equal to \$50 for each unfulfilled share. This penalty is large enough that it is worth trading at any price, no matter how unfavorable, to hit their target. The goal of a liquidity trader is to meet his or her target at the most favorable prices possible. Once they hit their targets, liquidity traders can buy or sell as many shares as they please without penalty.

3.4 *Subjects, Instructions and Incentives*

The experiments were conducted in the Business Simulation Laboratory (BSL) at the Johnson Graduate School of Management at Cornell University. The participants in the

² The screen does not provide traders with information about their time priority at a particular price. As is generally true in markets organized as electronic limit order books, traders infer their place in the queue from the number of shares in the book at a certain price when they submit an order.

experiments were Johnson MBA students. Each session involved twelve participants who were split randomly into two cohorts of six participants each. Upon arriving at the BSL, each subject received detailed written instructions, a copy of which is given in Appendix A. The instructions were reviewed in detail by the experiment administrator, who also answered any questions. The administrator then guided participants through the use and interpretation of the computer interface by trading two practice securities, which were exactly like the securities to be traded during the experiment, except that trading outcomes did not affect participants' cash winnings.

Traders started trading in each security with an endowment of \$0 in cash and zero shares. However, unlimited negative cash and share balances were permitted. Thus, traders could hold any inventory of shares they desired, including large short positions. Traders were told that at the end of trading, shares paid a liquidating dividend equal to their true value, so that their net trading gain or loss for a security would simply be equal to their ending share balance times the value of each share, plus their ending cash balance. Any penalties assessed to a liquidity trader for failing to hit a target are deducted from this trading gain or added to her trading loss.

Cash winnings for each session were determined by subtracting a "floor" from each trader's winnings in laboratory dollars, and then multiplying by an exchange rate that converts laboratory dollars into US dollars. The floor and exchange rate were derived from pilot experiments separately for each type of trader, and were designed so that each type would receive average winnings of approximately \$20/session. Traders were not told the floor or exchange rate, however, to minimize gaming behavior that might arise if traders knew they were unlikely to earn less than the minimum payment of \$5.

4. Results

The focus of our analysis is on the order strategies of traders: the choice between taking and making liquidity in a limit order market. We begin with market-wide summary statistics to provide a sense of how typical is the aggregate behavior that results from these experiments. We then examine differences in the use of market and limit orders by informed traders and liquidity traders, and we further investigate the differences between small and large liquidity traders. We next present results on how the submission rates of limit orders (relative to market orders) evolve through time, and on how the volatility of a security or the value of information held by informed traders affect trader strategies. Lastly, we examine how depth in the limit order book affects the traders' "Make or Take" decision.

The statistical tests in this section use a repeated-measures ANOVA. To judge statistical significance, we compute the average of the dependent variable within each cell (defined by the appropriate factors) for each of the eight cohorts. A repeated-measures analysis effectively treats each cohort as providing a single independent observation of the dependent variable. This design therefore reduces the problem, common in experimental economics, of overstating statistical significance by assuming that repetitions of the same actions by the same subject or group of subjects are independent events. When appropriate, we will use the ANOVA terminology of "main effect," "interaction," and "simple effect" to describe the statistical tests. A main effect examines the influence of one factor averaging over all the levels of the other factors. An interaction is when the effect of one factor is different at different levels of the other factors. A simple effect looks at the influence of one factor holding another factor at a specific level.³

³ For example, a significant Time main effect without a significant Type*Time interaction means that time exerts a similar influence on the behavior of all trader types. A significant Type*Time interaction implies that the different types of traders behave differently over time with respect to the dependent variable under investigation. In such a case, we would also examine the significance of simple effects that look at the influence of one factor holding

4.1. *Summary Statistics of Market-Wide Measures*

Figure 2 presents the evolution over time of three market-wide variables: volume, bid-ask spread, and absolute errors in trade price. Each Panel divides trading into eight 15-second intervals. While about 55 shares changed hands in a typical market, Panel A shows that volume exhibits the usual "U" shape observed in equity markets, with high volume at the open and the close of trading. Panel B shows that time-weighted spreads decline over time, and Panel C shows that price errors decline over time.⁴ These patterns suggest that markets behave reasonably well, in light of theoretical, archival, and experimental studies. Of particular importance is that our experimental markets appear to gradually incorporate information, a feature consistent with market efficiency.

4.2. *Summary Statistics of Traders' Strategies*

Panel A of Figure 3 presents summary statistics on the use of limit and market orders by informed traders, large liquidity traders, and small liquidity traders. The figure shows that informed traders submit more limit orders than liquidity traders do ($p = 0.0289$). That informed traders submit more limit orders stands in contrast to the prevailing wisdom in the theoretical literature. As Section 2 stresses, most theoretical models of limit order book markets assume that traders who provide liquidity through limit orders are uninformed about the true value of the security. Even the partial equilibrium models of Angel [1994] and Harris [1998], where informed

another factor at a specific level. So, we would look separately at the three types of traders to see whether time exerts significant influence on the behavior of each type.

⁴ The decline in the spread continues until the last interval, where it increases slightly. This is consistent with observed behavior in equity markets and can be explained by the desire of large liquidity traders to meet their trading targets and hence their use of more market orders than limit orders. The inelastic demand of the large liquidity traders creates profit opportunities for the other traders in the market who submit limit orders that are less aggressive, causing the spread to widen. This widening of the spread would also contribute to the small rise in ending price errors. Such behavior is consistent with the theoretical model of Brock and Kleidon [1992].

traders use limit orders under some circumstances, predict that informed traders would be less likely to submit limit orders than liquidity traders.⁵ We find this is not the case, revealing a complexity to informed behavior not captured by standard models.

The panel also provides information on what happens to the submitted limit orders. Interestingly, most limit orders submitted by the informed traders are left to expire in the book. This may reflect attempts by informed traders to "game" other market participants by submitting limit orders away from the market price, thereby creating an impression that the true value is different than otherwise believed. However, these orders also reflect genuine trading interest, as almost half of the trades of an informed trader (8.7 out of 17.8) occur when a limit order submitted by an informed trader is executed. Note that many more of the limit orders submitted by the liquidity traders are executed or cancelled, suggesting that these traders face (or fear they face) adverse selection.

Panel A also provides important information concerning the behavior of liquidity traders. Large liquidity traders trade an average of about 23 shares, slightly above their target of 20 shares. Small liquidity traders trade an average of about 14.4 shares, almost three times their target of five shares. The difference between the target and the actual number of shares traded points to a fundamental dissimilarity in the behavior of these two types of liquidity traders. The large liquidity traders are closer to the ideal definition of liquidity traders who trade for exogenous reasons. The average time it takes for a large liquidity trader to meet her target is 100 seconds, which means that she spends most of the trading period working towards completing this task. On the other hand, small liquidity traders meet their targets on average in 45 seconds (where the difference among the trader types is highly significant, $p < 0.0001$). After they hit their targets, small liquidity traders start speculating in the market. Since they do not possess any

⁵ This point is also made by Glosten [1994] for a special case of his model (see pp. 1150-1151).

special information about the true value of the security, their behavior is more analogous to that of day traders who attempt to predict short-term movements (or trends) in prices and trade to take advantage of them.

Panel B of Figure 3 sheds more light on this issue by examining the pattern over time in the average number of shares executed by trader types. Small liquidity traders start by intensely trading in the first two time intervals and thereafter maintain a relatively constant level of trading. Since they finish their targets on average by the end of the third time interval, the rest of the time they speculate on trends in prices or try to make money by providing liquidity and earning the spread. Large liquidity traders, on the other hand, trade much of their targets in the second half of the trading period when prices are closer to the true value of the security.

4.3 *Market versus Limit Orders Over Time*

Our primary analyses examine how traders choose between making and taking liquidity over time. To conduct these analyses, we define the (limit order) submission rate as the number of limit orders a trader submits divided by the sum of her limit and market orders.⁶ Figure 4 (Panel A) plots the submission rates of the three trader types separately for the eight time intervals. For all trader types, the use of market orders is more likely in the first interval (i.e., the submission rate is lower). At the end of pre-trading period, there is a book with limit orders waiting to begin trading. On average, there are 64.13 shares in the book at the beginning of a trading period, of which 13.78 are offered at the best bid and ask prices.⁷ It is natural for traders

⁶ We also replicated our analyses including only limit orders that improve, equal and come within a few ticks of the best existing order in the book. Our basic inferences are unchanged, so we report only the definition including all limit orders.

⁷ During the pretrading period, informed traders submit almost twice the number of limit orders submitted by liquidity traders (18.47 versus 9.33 for the small liquidity traders and 9.81 for the large liquidity traders). This difference is statistically significant ($p=0.0011$). Of the orders submitted by an informed during pretrading, 3.21 on average execute in the subsequent trading period, compared with 2.32 and 2.03 for the small and large liquidity

to start using more market orders in order to consume liquidity in the book. Informed traders pick-off mispriced limit orders in the book while liquidity traders use market orders to fill their targets. As time progresses, however, the trading strategies of the informed and the liquidity traders diverge.

The behavior of the large liquidity traders conforms tightly to the prediction of Harris [1998]. Harris argues that traders needing to meet a target by a certain deadline would start with a greater propensity for using limit orders, but as time progresses would switch to using more market orders. This strategy attempts to avoid paying the bid-ask spread, but it does expose the liquidity trader to execution risk. As the end of the trading period approaches, traders switch to market orders in order to meet their targets. Consistent with Harris' prediction, the submission rate of the large liquidity traders falls from about 70% in the first three intervals to under 35% in interval 8 (this pattern is statistically significant with $p < 0.0001$). An analysis of trading behavior provides similar results. Panel B of Figure 4 displays the "taking-rate" over time, defined as the percentage of trades completed using market orders (i.e., the number of market orders divided by the sum of market orders and executed limit orders). The large liquidity traders trade just under 40% of their shares by taking orders in the first interval, rising to about 65% in the final interval.

Informed traders behave in ways very different from theoretical predictions. Harris suggests that informed traders are less likely to use limit orders as time passes, but we find the opposite occurs. The submission rate of informed traders is just under 50% in the first interval, and *increases* to around 70% from the fourth interval on, a statistically significant change

traders, respectively. It is interesting to note that while the informed traders submit more limit orders during pretrading, they actively cancel fewer orders than do liquidity traders (1.20 for the informed compared with 1.24 for the small liquidity traders and 2.71 for the large liquidity traders). This difference is statistically significant with $p = 0.029$. Most of the orders submitted by the informed traders therefore are left to expire in the book at the end of the trading period (or disappear when "crossed" orders are eliminated at the end of pretrading). This may reflect the fact that since the informed know the true value of the security, leaving stale limit orders in the book is less risky for them than it is for the liquidity traders.

($p=0.0053$). Thus, informed traders are more likely to provide liquidity as time passes. The trading results are even more striking: Panel B of Figure 4 shows that informed traders execute almost 60% of their trades in the first interval by taking orders, but that proportion drops to just over 30% in the last interval (the pattern is significant with $p = 0.005$).

This apparent inconsistency with the theory can be explained by the dynamic adjustment of prices to private information. Early on in trading, the security value differs from the true value, giving informed traders an incentive to take liquidity by trading via market orders. Over time, however, the security value adjusts, and at some point it is bracketed by the bid-ask spread. Now, market orders are unprofitable for the informed traders, but a new role for the informed traders emerges. Informed traders know the exact location of the true value. They can place limit orders around it and supply liquidity to the market, just like dealers, making a small profit by means of the bid-ask spread. In fact, informed traders in principle can provide liquidity in a much more efficient manner than liquidity traders due to their greater knowledge of the asset value.⁸

This changing role for the informed traders, from aggressive speculators to dealers, is a phenomenon that to a large extent was not recognized in the theoretical literature. We suspect this is because models must be simple enough to be solved, and the dynamic considerations that give rise to the change in the behavior of the informed traders were never fully modeled. It is the changing value of the private information due to the process of price adjustment that holds the key to this behavior.

The behavior of small liquidity traders seems to be in-between the behaviors exhibited by the large liquidity traders and the informed traders. They start with a relatively low submission

⁸ Since liquidity traders are faced with uncertainty over the true value, the "option value" argument implies that they will shade their limit order prices (i.e., make them less aggressive) to reflect that uncertainty. Since informed traders do not face such an uncertainty, they can price the limit orders more aggressively and be at the inside quote supplying liquidity to the other traders.

rate that increases through the third time interval (when on average they meet their targets) and declines steadily thereafter. Panel B of Figure 4 shows that small liquidity traders begin the trading period like large liquidity traders with a low taking rate (34.45%) as they attempt to use limit orders to save on trading costs. They do not exhibit a clear pattern, however, and end the trading period with approximately the same taking rate.

4.4. Volatility and Extremity

Theoretical models suggest that informed traders would be more likely to use market orders when the value of their information is greater (Angel [1994], Harris [1998]). The risk of being picked off in Foucault [1999] implies that greater security value volatility would result in a higher submission rate of limit orders. We now turn to analyses that reveal the importance of information value by highlighting the effects of the extremity of the security realization, and also look at the effects of fundamental value volatility. Table 2 presents the mean submission rates for each trader type, within each of the four cells created by crossing extremity of the realized security value (high or low) with security value volatility (high or low). The results are aggregated over time intervals.⁹

The data clearly show that informed traders behave differently when the realized value of the security is extreme. Consistent with the predictions of Angel [1994] and Harris [1998], informed traders are more likely to use market orders when value extremity is high (meaning that their information is more valuable). The ANOVA analysis of the submission rate documents a

⁹ Since there are no significant interactions of Time and Extremity (or Volatility), we simplify the exposition by analyzing the effects of extremity or volatility with a less complex factorial design that does not consider time effects.

significant Type*Extremity interaction ($p=0.003$) and a marginally significant Type*Extremity*Volatility interaction ($p=0.0543$).

The lowest submission rate of limit orders by informed traders is in the Volatility_{Low}/Extremity^{High} cell, 61.64%. This cell can be thought of as a "surprise" cell in that the liquidity traders do not expect the realized value to be too far from the expected value because the volatility of the security is relatively low (a bell-shaped distribution). Therefore, the informed traders can use market orders to pick-off the liquidity traders (making a profit) for a longer period of time before prices adjust. When the realized value is extreme but the perceived volatility is higher, aggressive trading by market orders would result in greater revision in beliefs about the true value of the security and therefore will be less profitable. Consistent with this argument, the use of limit orders by informed traders in the Volatility^{High}/Extremity^{High} cell is higher, 67.8%.

In the Volatility_{Low}/Extremity_{Low} cell, which we can view as a "non-surprise" cell, we would expect the informed traders to move very quickly to play the role of dealers. There is little to gain by using market orders since the realized value is close to the expected value, and so prices are likely to adjust very quickly. Indeed, informed traders use more limit orders in that cell, with a submission rate of 79.27%, to make money from the bid-ask spread. The analysis of simple effects confirms that the Extremity factor influences the trading strategies of informed traders significantly ($p=0.0101$), with a weaker Volatility*Extremity effect ($p=0.0486$).¹⁰

We do not find support for the prediction of Foucault [1999] that higher volatility would result in a higher submission rate of limit orders. While the submission rate of the large liquidity

¹⁰ Since liquidity traders do not know the realization of the security value, the Extremity manipulation cannot affect their strategies directly. In principle, differences in the strategies of the informed traders may affect market conditions and thus indirectly cause the liquidity traders to change their trading strategies. The simple effect tests, however, cannot reject that the behavior of the liquidity traders is the same across the two Extremity regimes.

traders, our most "typical" liquidity traders, is lower in the $\text{Volatility}_{\text{Low}}/\text{Extremity}_{\text{Low}}$ cell, consistent with Foucault, the differences among the cells are not statistically significant.¹¹

Our results are essentially the same when we examine taking rates: informed traders take much less liquidity when extremity is low than when it is high, particularly in the "surprise cell" where volatility is low. Extremity and volatility have no statistically significant effects on the taking rates of the liquidity traders.

Theory predicts that liquidity traders should lose when trading against the informed traders, and this is exactly what we find. An informed trader makes on average \$40.66 in trading profit per market, while a small (large) liquidity trader loses on average \$22.89 (\$17.77). The smaller losses of the large liquidity traders can be due to the fact that they tend to trade late in the trading period when prices are closer to the true value of the security.

Panel C of Table 2 shows the trading profit and losses by volatility/extremity cells. It is clear that the informed (liquidity) traders make (lose) the most money in the $\text{Volatility}_{\text{Low}}/\text{Extremity}_{\text{High}}$ "surprise" cell, where the value of information is greater and it is easier to pick-off the liquidity traders. In the $\text{Volatility}_{\text{Low}}/\text{Extremity}_{\text{Low}}$ "non-surprise" cell, the trading profit of the informed is lowest, in line with their diminished role as speculators and the increase in their activity as dealers (profiting from the bid-ask spread). The ANOVA analysis of profit shows a significant $\text{Type} * \text{Volatility} * \text{Extremity}$ interaction ($p=0.0054$).

¹¹ One advantage of conducting an experiment rather than using actual market data is our ability to separate ex-ante volatility, which we manipulate using the distribution of security values that is known to all traders, and ex-post volatility or price volatility, which is also affected by the realized value of the security. The Extremity manipulation affects ex-post volatility since prices move more as they adjust to information when the realized value is extreme. The submission rates of both types of liquidity traders are slightly greater in the $\text{Extremity}_{\text{High}}$ cells. While these differences are not statistically significant, they may signal a possible volatility effect. It may be the case that instead of reacting to the ex-ante volatility of the security, traders react to actual price volatility. Such behavior makes sense because informed traders are present in the market. Liquidity traders could interpret high transaction price volatility as evidence of an extreme realized value and a greater likelihood of being picked-off as prices adjust. While there are no informed traders in Foucault's model, he generates a similar risk of being picked-off by combining true value volatility with the inability to cancel limit orders. Therefore, the weak pattern we observe with respect to extremity may reflect the intuition of Foucault's model but is generated by a different mechanism (the presence of informed traders in the market rather than frictions on the responses of traders to public information).

4.5. *Trading Strategies and the State of the Book*

The state of the limit order book can influence order submission strategies. Our main findings are that all traders react similarly to depth on the opposite side of the book, but that informed and liquidity traders behavior diverges when we consider same side depth. We interpret this latter result as indicative of the quasi-dealer function played by informed traders.

Our tests are motivated by Parlour [1998], who provides predictions on how depth at the best bid or offer (BBO) affects the likelihood of submitting a limit or a market order. She predicts that traders are less likely to use a limit order if the book on the same side as the order, say the ask side for a sell order, is thicker. This so called "crowding out" effect arises due to the time priority of orders already in the book lowering the probability of execution of a new limit order at the ask. Parlour also predicts that traders are more likely to use limit orders if the book on the other side, say the bid side for a sell order, is thicker. (We note that there are no informed traders in Parlour's model; it is unclear whether these predictions will hold for either informed or liquidity traders in an environment in which some traders possess private information.)

To test the association between depth and submission rates, we record for each market and limit order the depth level at the best prices on both the same side and the opposite side of the book. To simplify the analysis, we consider only four levels of depth: (1) depth less than or equal to two shares, (2) depth greater than two shares and less than or equal to four shares, (3) depth greater than four shares and less than or equal to six shares, and (4) depth greater than six shares.¹² We calculate a separate submission rate for each trader in each level by dividing the

¹² We examined a couple of different bounds for the categories to make sure that our results are not driven by the arbitrary allocation into categories. The results we discuss in the text seem to be robust to perturbations of the bounds.

limit orders submitted in a certain level by the sum of limit and market orders in that level. We also repeat the analysis considering only limit orders submitted at the best existing bid or offer.¹³

Panel A of Figure 5 presents the submission rates for all orders conditional on four levels of depth on the same side as the order. The behavior of informed and large liquidity traders differ when there is a lot of depth in the book. The statistical test for differences in behavior (interaction Type*Depth), borders on significance ($p=0.0627$). The large decline in category 4 for the large liquidity traders is consistent with Parlour's prediction that traders are less likely to submit limit orders when depth on the same side of the BBO is greater. The simple effects analysis shows that this pattern in the behavior of the large liquidity traders is statistically significant ($p=0.0355$).

In contrast, the informed traders show higher submission rates when the book is deep on the same side as the order. This behavior may be due to the liquidity supply role taken on by the informed traders. In the theoretical model, existing depth in the book is not the making of the individual trader who decides between a limit and a market order. Here, both informed traders are more likely to act as dealers when such a strategy is profitable. Competition among the informed traders may cause them to submit additional limit orders even as depth on the same side increases.

Submission rates for limit orders submitted only at the best existing bid or offer provide even stronger evidence of the same behavior. For these orders the interaction of trader type and depth level is strongly significant ($p=0.0016$), as is the simple effect testing the pattern in the informed traders' submission rates across depth levels ($p=0.0017$). While in the figure it looks as

¹³ Note that in both cases, we condition only on depth in the book at the best bid or offer (not the entire book). The difference between the regular submission rate to the one that considers only the best existing bid or offer has to do with the limit orders submitted, not with the book. The first measure is the number of all limit order submitted (for a given depth level) divided by the sum of the limit and market orders submitted (at that depth level). The second measure has in the numerator only limit orders whose prices match the existing best bid or offer in the book.

if small liquidity traders also submit more limit orders as the depth on the same side increases, this pattern is not statistically significant and is not robust to perturbations in the bounds of the categories.

Panel A of Figure 6 presents the submission rate of limit orders conditional on BBO depth on the other side of the book (say the ask side for a buy order). We find that all trader types increase their submission rates as the depth on the other side increases. The statistical tests confirm that the main effect of the Depth factor is highly significant ($p < 0.0001$), but the interaction of Type*Depth is not. Thus, we cannot say unequivocally that the behavior of the informed traders is very different from the behavior of the liquidity traders. The steep increase in the submission rate of the informed traders, however, is highly significant ($p = 0.0001$). The results are essentially identical when we consider only orders submitted at the best bid or offer (with $p = 0.0318$ for the main effect of Depth).

Parlour's prediction is that the greater the ask (bid) depth, the more likely are buyers (sellers) to submit limit orders. The intuition behind her result is that, focusing on the ask side, buyers rationally anticipate the crowding out of limit sell orders when the ask depth is greater, and so limit buy orders become more attractive (in anticipation of more market sell orders). Our results strongly support this prediction of her model.

5. Conclusions

In this study we have investigated the evolution of liquidity in an electronic limit order market. Our particular focus has been on the order behavior of informed and liquidity traders. As we have shown, informed traders and liquidity traders act quite differently in term of order submission strategies. Large liquidity traders prefer limit orders early in the trading period, but

their preference shifts over time toward using more market orders. Their behavior corresponds relatively well to the predictions from theoretical models.

This is not the case for informed trader behavior. We find that informed traders play a particularly intriguing, and heretofore unappreciated, role in electronic markets. As predicted by theory, the informed do trade to capture the value of their private information. However, this trading behavior is complex, involving the use of both limit orders and market orders. Indeed, informed traders actually submit more limit orders than they do market orders. This reflects the ability of informed traders to know better the true value of the asset, and so a source of profit for the informed is earning the bid-ask spread via limit order submission. This behavior results in the informed providing liquidity to the market. Such a “making liquidity” role explains why electronic markets can endogenously create liquidity even in the presence of information asymmetry.

Electronic markets are now a ubiquitous feature of securities markets. As trading increasingly shifts to electronic platforms, the nature of liquidity provision takes on greater importance. An important contribution of our research is to make clearer how liquidity is provided in such a setting, and the factors that influence it.

6. References

- Ahn, H.-J., Bae, K.-H., Chan, K., 2001. Limit orders, depth and volatility: Evidence from the Stock Exchange of Hong Kong. *Journal of Finance* 56, 767-788.
- Angel, J. J., 1994. Limit versus market orders. Unpublished working paper. School of Business Administration, Georgetown University.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655-1689.
- Brock, W.-A., Kleidon, A.-W., 1992. Periodic market closure and trading volume: a model of intraday bids and asks. *Journal of Financial Dynamics and Control* 16, 451-489.
- Chakravarty, S., Holden, C. W., 1995. An integrated model of market and limit orders. *Journal of Financial Intermediation* 4, 213-241.
- Cohen, K. J., Maier, S. F., Schwartz, R. A., Whitcomb, D. K., 1981. Transaction costs, order placement strategy, and existence of the bid-ask spread. *Journal of Political Economy* 89, 287-305.
- Foucault, T., 1999. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets* 2, 99-134.
- Foucault, T., Kadan, O., Kandel, E., 2001. Limit order book as a market for liquidity. Unpublished working paper. HEC School of Management.
- Glosten, L. R., 1994. Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127-61.
- Harris, L., 1998. Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions and Instruments* 7.
- Hasbrouck, J., Saar, G., 2001. Limit orders and volatility in a hybrid market: The Island ECN. Unpublished working paper. Stern School of Business, New York University.
- Hollifield, B., Miller, R. A., Sandas, P., 1999, Empirical Analysis of Limit Order Markets, Unpublished working paper, Carnegie-Mellon University.
- Kumar, P., Seppi, D., 1994. Limit and market orders with optimizing traders. Unpublished working paper. Graduate School of Industrial Administration, Carnegie Mellon University.

- Parlour, C., 1998. Price dynamics in limit order markets. *Review of Financial Studies* 11, 789-816.
- Parlour, C. A., Seppi, D. J., 2001. Liquidity-based competition for order flow. Unpublished working paper. Graduate School of Industrial Administration, Carnegie Mellon University.
- Rock, K., 1990. The specialist's order book and price anomalies. Unpublished working paper. Graduate School of Business, Harvard University.
- Seppi, D. J., 1997. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies* 10, 103-150.
- Stoll, H., 1978, The Supply of Dealer Services in Securities Markets, *Journal of Finance*, 33, 1133-1151.
- Sunder, S. (1995), Experimental Asset Markets: A Survey. In *The Handbook of Experimental Economics*. Edited by J.H. Hagel and A.E. Roth. Princeton, NJ. Princeton University Press.

Table 1
Absolute Deviation of Security Value from Prior Expected Mean Value of \$25.

Actual values of securities traded equaled 25 plus or minus the indicated number. The sign of the deviation from 25 varied across securities and across cohorts. Half of the cohorts traded securities in order 1, while the other half traded securities in order 2. Only the securities indicated in bold were used in the analysis; other securities were included only to allow security values to be distributed as indicated to traders. Low-volatility securities have deviations of 7 or less; high-volatility securities have deviations of 12 or more.

Security Number	Order 1	Order 2
	Low-Volatility Setting	High-Volatility Setting
1	3	21
2 (Low-Extremity)	2	2
3 (High Extremity)	15	15
4	4	13
5 (Low-Extremity)	4	4
6	7	24
7 (High-Extremity)	20	20
8 (High-Extremity)	16	16
9 (Low Extremity)	6	6
10	8	9
	Low-Volatility Setting	High Volatility Setting
11	21	3
12 (Low-Extremity)	7	7
13 (High Extremity)	16	16
14	13	4
15 (Low-Extremity)	19	19
16	24	7
17 (High-Extremity)	3	3
18 (High-Extremity)	17	17
19 (Low Extremity)	1	1
20	9	8

Table 2
Submission Rate and Taking Rate by Volatility and Extremity Factors

This table presents the submission rates, taking rates, and average trading profit per market for the different types of traders in each of the Volatility/Extremity cells. Submission Rate is defined as the number of limit orders a trader submits divided by the sum of her limit and market orders. Taking Rate is defined as the percentage of trades completed using market orders (the number of market orders a trader submits divided by the sum of her market orders and executed limit orders). Trading profit per market is defined as the sum, across all limit orders and market orders a trader submits, of the difference between the trading price and the true value of the security. The Volatility factor of the security's true value has two levels: high (a uniform distribution) and low (a bell-shaped distribution). The Extremity factor (how different the realized value of the security from its unconditional mean) has two levels: high (realized values that are at least \$12 from expected value) and low (realized values that are no more than \$7 from expected value). The two informed traders know the true value of the security before trading begins. One small liquidity trader needs to sell 5 shares and another needs to buy 5 shares. One large liquidity trader needs to sell 20 shares and another needs to buy 20 shares. We first compute the variable under investigation for an individual trader and then the average for a trade type within each of the eight cohorts. The numbers in the table represent the averages across the cohorts.

Panel A: Submission Rate of Limit Orders

	Informed	Small Liquidity	Large Liquidity
Volatility _{Low} /Extremity _{Low}	79.3%	65.3%	52.5%
Volatility _{Low} /Extremity _{High}	61.6%	66.4%	59.3%
Volatility _{High} /Extremity _{Low}	71.3%	62.2%	59.0%
Volatility _{High} /Extremity _{High}	67.8%	64.2%	60.7%

Panel B: "Taking Rate" or the Percentage of Trades Completed using Market Orders


	Informed	Small Liquidity	Large Liquidity
Volatility _{Low} /Extremity _{Low}	34.7%	43.7%	61.3%
Volatility _{Low} /Extremity _{High}	55.0%	40.0%	51.5%
Volatility _{High} /Extremity _{Low}	48.3%	45.1%	54.2%
Volatility _{High} /Extremity _{High}	50.7%	43.0%	51.1%

Panel C: Average Trading Profit per Market in Laboratory Dollars

	Informed	Small Liquidity	Large Liquidity
Volatility _{Low} /Extremity _{Low}	17.15	-8.81	-8.33
Volatility _{Low} /Extremity _{High}	83.48	-51.35	-32.13
Volatility _{High} /Extremity _{Low}	25.10	-8.35	-16.75
Volatility _{High} /Extremity _{High}	36.90	-23.04	-13.85

Figure 1
Example of a Trading Screen

This figure presents a screen snapshot for a practice security. The screen includes two graphs showing market activity. The left side of each graph shows every price at which an order has been posted (shown in green for the highest bid and lowest ask price, and yellow for other prices), and the number of shares posted at that price (shown by the number to the left of the graph). The right side of each graph shows every price at which the trader has personally posted an order, and the number of shares that the trader has posted at that price. The center of each graph also includes a solid red line indicating the highest bid or lowest ask entered by any trader, and a solid green line indicating the highest bid or lowest ask entered by that particular trader.




The main trading period has begun. Use the Order Entry Graph to enter your orders. Right click to remove orders.


Figure 2
Market-Wide Summary Statistics

This figure presents summary statistics for volume, bid-ask spread, and price errors over time. Volume is defined as the number of shares traded. Bid-ask spread is the difference between the best price in the book for buying a share and the best price in the book for selling a share. Absolute errors in trade price are defined as the absolute value of the difference between the trade price and the true value of the security. The variables are computed separately for each 15-second interval in the trading period.

Panel A: Volume



Panel B: Bid-Ask Spread



Panel C: Absolute Errors in Trade Price


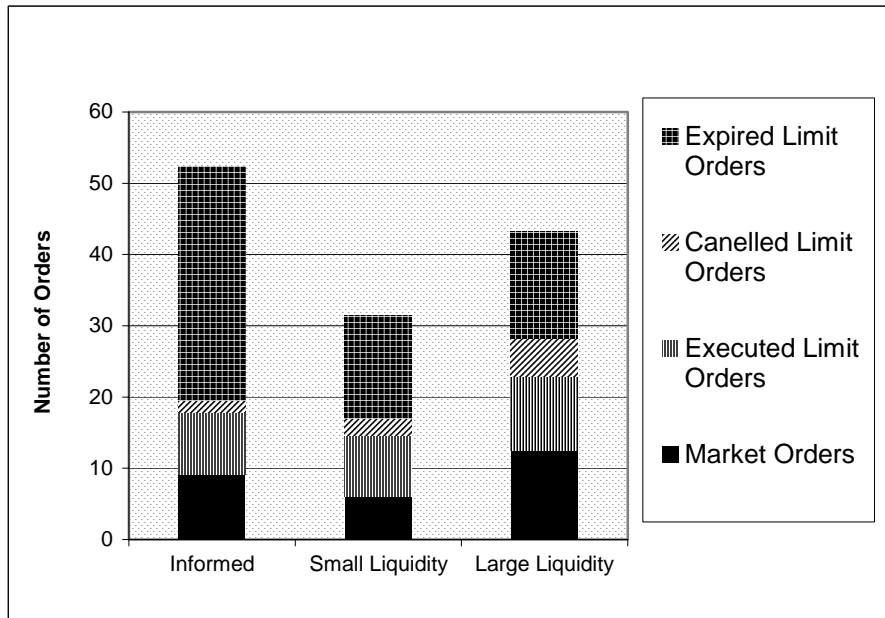


Figure 3
Summary Statistics of Traders' Strategies

This figure presents summary statistics on orders and trades for the different types of traders. The two informed traders know the true value of the security before trading begins. One small liquidity trader needs to sell 5 shares and another needs to buy 5 shares. One large liquidity trader needs to sell 20 shares and another needs to buy 20 shares. Each limit order submitted by a trader is for one share. Market orders are defined as the taking of limit orders at the best prices in the book. Expired limit orders are orders that are left in the book at the end of the trading period while cancelled limit orders are those that traders actively remove from the book. Panel B plots the average number of shares executed by a trader who belongs to one of the three types. This variable is computed separately for each 15-second interval in the trading period.

Panel A: Market and Limit Orders



Panel B: Shares Executed by Trader Type





Figure 4
Limit and Market Order Strategies

This figure presents the submission rates and taking rates for the different types of traders over time. The two informed traders know the true value of the security before trading begins. One small liquidity trader needs to sell 5 shares and another needs to buy 5 shares. One large liquidity trader needs to sell 20 shares and another needs to buy 20 shares. Submission Rate is defined as the number of limit orders a trader submits divided by the sum of her limit and market orders. Taking Rate is defined as the percentage of trades completed using market orders (the number of market orders a trader submits divided by the sum of her market orders and executed limit orders). The variables are computed separately for each 15-second interval in the trading period.

Panel A: Submission Rate of Limit Orders



Panel B: "Taking Rate" or the Percentage of Trades Completed using Market Orders





Figure 5
Submission Rate Conditional on BBO Depth on the Same Side as Order

This figure presents the submission rates of the different types of traders conditional on depth in the book at the best bid or offer (BBO) on the same side as the order. The two informed traders know the true value of the security before trading begins. One small liquidity trader needs to sell 5 shares and another needs to buy 5 shares. One large liquidity trader needs to sell 20 shares and another needs to buy 20 shares. Panel A reports the submission rate of limit orders (the number of limit orders a trader submits divided by the sum of her limit and market orders) conditional on four BBO depth levels: (1) depth less than or equal to two shares, (2) depth greater than two shares and less than or equal to four shares, (3) depth greater than four shares and less than or equal to six shares, and (4) depth greater than six shares. Panel B reports the number of limit orders submitted at the BBO for each level divided by the total number of limit and market orders in that depth level.

Panel A: All Limit Orders



Panel B: Limit Orders at the Best Bid or Offer






Figure 6
Submission Rate Conditional on BBO Depth on the Other Side of the Book

This figure presents the submission rates of the different types of traders conditional on depth in the book at the best bid or offer (BBO) on the other side of the book (i.e., the ask side for a buy order). The two informed traders know the true value of the security before trading begins. One small liquidity trader needs to sell 5 shares and another needs to buy 5 shares. One large liquidity trader needs to sell 20 shares and another needs to buy 20 shares. Panel A reports the submission rate of limit orders (the number of limit orders a trader submits divided by the sum of her limit and market orders) conditional on four BBO depth levels: (1) depth less than or equal to two shares, (2) depth greater than two shares and less than or equal to four shares, (3) depth greater than four shares and less than or equal to six shares, and (4) depth greater than six shares. Panel B reports the number of limit orders submitted at the BBO for each level divided by the total number of limit and market orders in that depth level.

Panel A: All Limit Orders



Panel B: Limit Orders at the Best Bid or Offer





APPENDIX: INSTRUCTIONS TO SUBJECTS

Overview

During this session, you will trade shares of 20 securities that have values denominated in “laboratory dollars.” We convert your trading gains into U.S. dollars to determine your payment. At the end of trading, we will ask a series of questions about your experience.

Value of a Security

The true value of each security is distributed over the interval [0,50], in one of the ways shown below. For 10 securities, values will follow the distribution on the left, in which having a



true value of 25 is as likely as having a true value of 2 or 48. For the other 10 securities, values will follow the distribution on the right, in which values near 25 are much more likely than more extreme values. The experiment administrator will make sure that everyone in your market knows which distribution is in force at all times during trading.

How to Trade

You trade shares by entering orders that others can “take” or by “taking” orders that others have entered. All orders are for one share, but you can enter and take multiple orders at each price. Here are your options:

- ◆ **Entering a Bid** A bid is an order to **buy** a share at a stated price. You will buy at that price if someone else chooses to take your bid, and sells a share to you at the price you indicated.
- ◆ **Entering an Ask.** An ask is an order to **sell** a share at a stated price. You will sell at that price if someone else chooses to take your ask, and buys a share from you at the price you indicated.
- ◆ **Taking a Bid or Ask.** If you click on the “SELL 1” button on the BID column, you will sell a share at the highest bid. If you click on the “BUY 1” button on the ASK column, you will buy a share at lowest ask.
- ◆ **Removing a bid or ask.** You can remove (cancel) a bid or ask that you entered, simply by *right-clicking* on it. Doing so removes all orders you entered at that price.

The Trading Period

Trading for each security is split into two parts: pre-trading and main trading.

- ◆ **Pre-Trading (30 seconds).** During the pre-trading period, traders can enter bids and asks, but no one can take them. At the end of pre-trading, the highest bid and the lowest ask will be paired up. If they “cross” (the bid is priced higher than or equal to the ask) the more recent order will be deleted. This will be repeated until there are no crossing orders remaining.
- ◆ **Main Trading (120 seconds).** During the main trading period, all traders can enter bids and asks, and can also take the bids and asks entered by other traders.

Some Trading Rules

The following rules keep you from entering or taking any orders you please.

- ◆ You can't trade with yourself. If you try to take an order you entered, your request will be rejected.
- ◆ You can never enter a bid at a price greater than your own ask, or an ask at a price less than your own bid. Doing so would be like trying to trade with yourself.
- ◆ During Main Trading, you can't enter a crossing order (a bid higher than an existing ask or an ask lower than an existing bid). If you are willing to buy at the lowest ask, simply click the “BUY 1” button. If you are willing to sell at the highest bid, simply click the “SELL 1” button.

Trader Types

The session includes two types of traders.

- ◆ **Informed Traders** know the true value of the security, which they learn before trading begins.
- ◆ **Target Traders** don't know the value of the security. However, each has a trading “target” to meet before trading is complete. One trader's target is to sell 20 shares; another's is to buy 20 shares; another's is to sell 5 shares, and another's is to buy 5 shares. If you have a target, it will be stated clearly on your screen. At the end of trading, you will be assessed a penalty equal to \$50 for each unfulfilled share. This penalty is large enough that it is worth trading at any price, no matter how unfavorable, to hit your target. The goal of a target trader is to meet his or her goal at the most favorable prices possible. Once they hit their targets, target traders can buy or sell as many shares as they please without penalty.

Making Money

You start each security with \$0 and 0 shares. However, negative cash and share balances are permitted. Thus, you can buy shares even if you don't have money to pay for them, and you can sell shares you don't own ("short selling").

At the end of trading for each security, the shares you own pay a liquidating dividend equal to their true value. If you have a positive balance of shares, the dividend is added to your cash balance for each share you own. If you have a negative balance of shares, the dividend is subtracted from your cash balance for each share you own. The resulting number is your trading gain (if positive) or trading loss (if negative). Any penalties assessed for failing to hit a trading target are deducted from your trading gain or added to your trading loss.

You make money every time you buy a share for less than true value or sell a share for more than true value. For example, buying a share worth \$30 at a price of \$23 creates a gain of \$7. Selling that share at that price creates a loss of \$7.

Converting Laboratory Dollars into US Dollars.

Laboratory winnings, as described above, will be converted into US\$ according to the formula

US\$ Payment = Exchange Rate x (Net Gain/Loss in Laboratory \$ + "Adjustment").

You are guaranteed a minimum payment of US\$10.

You will not learn the exact adjustment or exchange rates. However, we will tell you two key facts. First, the exchange rate is positive, meaning that the more laboratory dollars you win, or the fewer you lose, the more \$US you take home. Second, different types of traders have different adjustments. Target traders have a positive adjustment, meaning that they can lose laboratory dollars but still take home a substantial payment. Informed traders have negative adjustments, meaning that they need to earn some number of laboratory dollars in order to receive more than the minimum payment. The parameters are set so that the average winnings will be approximately US\$25 for each person for the session.

Other Rules

Please do not talk with other subjects or look at their computer screens without explicit permission from the experiment administrator. Please ask the administrator before leaving the room for any reason.

Consent Form

This is a research experiment intended to help regulators and researchers understand the functioning of financial markets and business decision-making. If you have any questions about the administration of this experiment, please contact Professor Robert Bloomfield (rjb9@cornell.edu, 255-9407, 450 Sage Hall, Cornell University). Participants can also contact the University Committee on Human Subjects (uchs-mailbox@cornell.edu, 255-2943, 115 Day Hall, Cornell University).

In order to participate, you must sign the consent form below:

I consent to participate in this trading session, and agree to abide by all of the rules previously described to me by the administrator throughout my participation. I recognize that:

- (1) If I breach any of the rules governing my behavior, I forfeit my right to any money I might have earned by participating;
- (2) I have the right to leave the experiment at any time, without penalty, but that in doing so I forfeit my right to any money I might have earned by trading;
- (3) This experiment has been approved by the Cornell University Committee on Human Subjects as research that uses no deception of any kind.

I understand that my performance will be kept private, and that the research is intended to help clarify the nature of decisions made in business settings like the one arising in this session.

Signed _____

Date _____