NEW YORK UNIVERSITY
STERN SCHOOL OF BUSINESS
FINANCE DEPARTMENT

Working Paper Series, 1994

Limit Order Trading

Puneet Handa and Robert A. Schwartz

FD-94-35

LIMIT ORDER TRADING

Puneet Handa*

Robert A. Schwartz**

First draft: February, 1991

Current draft: January, 1995

* Associate Professor of Finance, College of Business Administration, University of Iowa, Iowa City 52242.

** Professor of Finance and Economics, Stern School of Business, New York University, 90 Trinity Place, New York, New York 10006.

We thank Yakov Amihud, Marshall Blume, Ned Elton, Sandy Grossman, Jaques Hamon, Joel Hasbrouck, Bertrand Jacquillat, Ananth Madhavan, Maureen O'Hara, Jim Shapiro and participants at seminars held at Stern School of Business, Wharton School, New York Stock Exchange and Bankers Trust for their comments. Both authors are grateful for the funding provided by the NEC Corporation and the Center for Japan-U.S. Business and Economic Studies, Stern School of Business, New York University.

Abstract

The paper analyzes the rationale for and profitability of limit order trading. Although limit orders are essential to the functioning of order driven markets, their use has received relatively little attention in the literature. Trading via limit order is, in fact, sub-optimal when transaction prices change solely in response to new information. We suggest that in an order driven market a paucity of limit orders can result in order imbalances causing the transaction price to move temporarily away and then revert to the true price. Such short term changes in transaction price can offset losses incurred by limit order traders because of permanent changes in transaction price due to information. Further, we suggest that the market can be in ecological balance with liquidity driven price changes being just sufficient for the flow of market and limit orders to equilibrate. We use transactions data for the Dow Jones Industrial stocks for 1988 to compare a limit order strategy with a market order strategy, and find that limit order returns conditional on execution are greater than the unconditional market order returns, while limit order returns conditional on non-execution are lower. We also test the profitability of placing a network of buy and sell limit orders, and document the existence of a limit order spread that is appreciably greater than the posted bid-ask spread. Our findings suggest that trading via limit orders is desirable for participants who are willing to risk non-execution, and that trading via market orders is desirable for participants who are not willing to take the risk.

LIMIT ORDER TRADING

1. Introduction

An important design feature of a securities market is whether it is quote driven (such as NASDAQ in the United States and SEAQ in London), order driven (such as the Tokyo Stock Exchange), or both (such as the New York Stock Exchange, which has a specialist system, a public limit order book, and floor traders). A market is *quote driven* if dealers announce the prices at which other market participants can trade; it is *order driven* if some investors, by placing limit orders, establish the prices at which other participants can buy or sell shares. Although limit orders are routinely submitted to markets such as the New York and Tokyo Stock Exchanges, their desirability for investors has received little attention in the literature until recently.[1] This paper analyzes the rationale for, and profitability of, limit order trading. The issue is clearly important to market participants individually. It is also important from a market structure point of view: the profitability of limit order trading is essential to the viability of an order driven market. Moreover, an understanding of the order flow dynamics that make limit order trading viable enhances our knowledge of how a market's microstructure affects the return generation process.

In a world where transaction prices move solely in response to information, trading via limit order is costly because the individual who places a buy (sell) limit order has written a free put (call) option to the market. Consider an investor who wishes to buy shares of xyz common at $50 or better, and let that investor select between submitting a market order (that would execute, say, at $50) and submitting a limit order at $49. If news causes the share price to fall below $49, the option will be exercised and the individual can

---

[1] Glosten (1994) examines the role of limit orders in the context of an open electronic book. In a recent working paper, Harris and Hasbrouck (1993) empirically assess the profitability of limit orders versus market orders on the New York Stock Exchange.

lose from trading with a better informed investor. Alternatively, the price of shares on the market may not fall to $49, and the individual might miss the investment opportunity.

Copeland and Galai (1983) address the adverse selection problem from the viewpoint of a dealer. Building on Bagehot (1971), they show that the market maker's spread generates returns that can cover the cost of trading with informed participants. Limit order traders resemble dealers in that they provide liquidity and immediacy to the market. However, the primary objective of a limit order trader is to implement an investment decision and they do not continuously post two-sided quotes.

In a recent paper, Glosten (1994) derives an equilibrium where limit order traders implicitly gain from liquidity driven price changes but lose from information driven price changes. Glosten assumes two distinct classes of participants, those who trade by limit order and those who trade by market order, and does not model a participant's decision to trade via limit order or market order. In this paper we suggest that Glosten's analysis can be extended to consider the choice faced by an investor who wishes to buy or sell a share of the risky asset over a trading window.

Our investor can choose to trade via limit order and supply liquidity to the market or choose to trade via market order and demand liquidity from the market. The choice depends critically on the investor's beliefs about the probability of his or her limit order executing against an informed or a liquidity trader. Transaction price changes due to the arrival of a liquidity trader are temporary and reversible, and having a limit order execute against such price changes is desirable.[2] In contrast, transaction price changes due to the arrival of an informed trader are permanent and irreversible, and having a limit order execute against such price changes is undesirable. In addition, if a limit order fails to execute in a trading window, a

---

[2] That the transaction price can change without change in expected future payoffs is also found in Easley and O'Hara (1987), Grossman and Miller (1988), Admati and Pfleiderer (1989), Glosten (1989), and Leach and Madhavan (1993), among others.

decision has to be made at the end of the window of whether to trade at the prevailing transaction price or to forego trading and this decision has cost implications.

We endogenize the decision to trade via market or limit order and view the trading environment as an ecology where the supply of, and demand for, liquidity can be in natural balance. We conduct empirical experiments that use this conceptualization to assess the relative profitability of limit and market order trading.

Our empirical tests use 1988 transaction price data for the thirty Dow Jones Industrial stocks. The market order strategy involves buying or selling at market at the start of a trading session. The limit order strategy involves placing an experimental limit order which is subsequently converted to a market order if it does not execute within a specified period of time.[3] We find that returns to the limit order strategy conditional on the orders executing are greater than unconditional market order returns. We also find that returns to the limit order strategy conditional on the orders not executing are lower than the unconditional market order returns. Finally, we assess the profitability of placing a network of buy and sell limit orders around an asset's current price, and find evidence of what we call a *limit order spread* that is larger than the spreads posted on the market.[4]

The paper is organized as follows. In Section 2, we examine the rationale for trading via limit orders. Section 3 contains empirical evidence on the profitability of the limit order versus market order trading strategy. In Section 4, we examine a multiple limit order strategy (buy and sell) and empirically document the existence of a limit order spread. Our conclusions are presented in Section 5.

---

[3] In contrast, Harris and Hasbrouck (1993) base their tests of limit order trading on orders that have actually been submitted to the market.

[4] The documentation of the limit order spread is consistent with Grossman and Miller (1988) who postulate that a market maker can earn much more or less than the spread quoted at the time of the placement of the limit order.

2. A Rationale for Limit Order Trading

2.1 The Analytical Framework

In this section we present an analytical framework to assess an investor's decision to trade via limit order versus market order. Glosten (1994) provides a rationale for limit order trading. In his framework, there are two types of investors: patient traders, who may supply liquidity to the market, and other traders who wish to trade immediately. The former place limit orders, and the latter place market orders. A limit buy (sell) order trader can expect to lose if the order executes upon the arrival of an informed trader with a marginal valuation below (above) the limit price, and can expect to gain if the order executes upon the arrival of a liquidity trader. The trader will not choose to place a limit order unless the expected gain from transacting with a liquidity trader exceeds the expected loss from transacting with an informed trader.

Glosten (1994) does not explicitly model an investor's decision to trade via a limit order or a market order. In what follows, we suggest that his analysis can be extended to do this. We consider an investor who wishes to buy a share of the risky asset over a trading window (the analysis for a sell order is symmetric). Let the trading window be followed by an investment period of length t.[5]

Our investor faces the following microstructure. The opening price of the trading window is P_0. A limit order book exists at the start of the trading window and a trade is made during the window if a public market order arrives. For simplicity we assume that just one public trader, who is either informationally motivated or liquidity motivated, arrives during the trading window to buy or sell at market. That trader's order executes against the book and moves the transaction price to $P_x = S$ which holds at the

[5] Assuming that the stock is held over an investment period of a fixed length t facilitates comparing returns across different strategies.

end of the trading window. If the arriving trader is informed the true price of the security is P_x. If the arriving trader is liquidity motivated, the true price of the security remains at the opening price P_0. The following schema represents the sequencing of events:

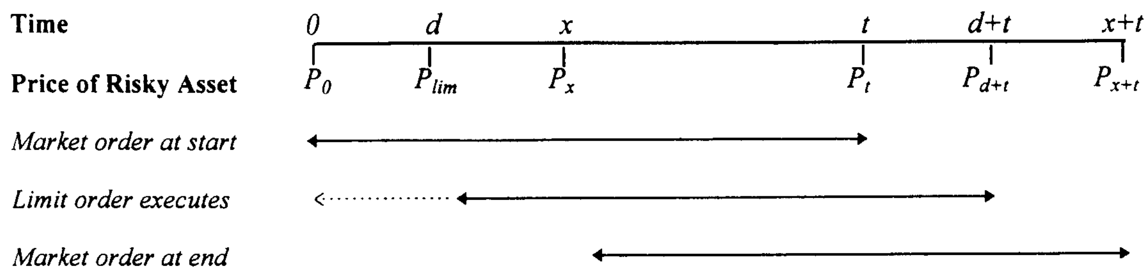

Consider the situation faced by the investor at the start of the trading window. If he or she uses a market order, the trade occurs immediately at the price P_0. The position is then sold at a price P_t after an investment period of length t. If the investor uses a limit order, the trade occurs within the trading window if a market order to sell clears the book down to a transaction price at or below the limit price P_{lim}. Otherwise a trade fails to occur within the trading window. If that happens, a second decision must be made: (1) withdraw the order and not trade, or (2) buy by market order at the last price of the trading window P_x, a price that we sometimes refer to as S for notational convenience.

We denote the probability density function of the last price S, conditional upon the arrival of an informed trader, by $f(S|I)$ and the probability density function of the last price S, conditional upon the arrival of a liquidity trader, by $f(S|L)$. We let F denote the probability that the last price of the trading window is equal to or below P_{lim} conditional upon a liquidity trader arriving. We assume that S is symmetrically distributed around P_0. Further, we assume that an informed trader arrives with an exogenously specified probability p, and that a liquidity trader arrives with a probability q (where $q = 1 - p$).

If the limit order executes within the trading window and the arriving trader is informed, the expected gain is given by:

$$p \cdot \int_{-\infty}^{P_{\lim}} \left(S - P_{\lim}\right) f\left(S \mid I\right) dS \le 0$$

If the limit order executes because of the arrival of a liquidity trader, our investor's expected gain is:

$$q \cdot \int_{-\infty}^{P_{\lim}} \left(P_0 - P_{\lim}\right) f\left(S \mid L\right) dS$$

$$= q \cdot F \cdot \left(P_0 - P_{\lim}\right) \geq 0$$

Thus, our investor's net expected gain from the execution of the limit order can be written as:

Expected Gain from Limit Order Execution:

$$p \cdot \int_{-\infty}^{P_{\lim}} \left(S - P_{\lim}\right) f\left(S \mid I\right) dS + q \cdot F \cdot \left(P_0 - P_{\lim}\right)$$

Our investor expects to gain from the execution of a limit order only if the expected loss from the arrival of an informed trader is more than offset by the expected gain from the arrival of a liquidity trader. For this to be the case, the probability of the arrival of a liquidity trader with a reservation value equal to or less than the limit price must be sufficiently large. If $q \cdot F = 0$, the transaction price fluctuates only due to information events. In this case, an investor must expect to lose from placing a limit order. On the other hand, if $q \cdot F > 0$, transaction price fluctuates additionally due to liquidity motivated trades. As, on expectation, the transaction price recovers from such fluctuations, it has a tendency to revert to the true price. As $q \cdot F$ increases, there can be a positive net gain from the execution of a limit order.

We define:

> **Bagging Cost**: When there are no liquidity driven changes in transaction price to cause the execution of a limit order (i.e., $q \cdot F = 0$), a limit order trader faces a negative expected gain from limit order execution. This negative expected gain is referred to as a *bagging cost*. Sufficient fluctuation in the transaction price due to market order trades by liquidity traders (i.e., $q \cdot F$ sufficiently large), can result in positive expected gain from limit order execution. This positive expected gain corresponds to a *negative bagging cost.*

We now turn to the case where the buy limit order fails to execute within the trading window. The limit order trader must then choose (1) to forego trading, or (2) to buy by market order at the last price of the trading window, S. If the arrival of an informed trader results in the last price, the closing purchase is at the true price and the expected gain is zero. If, on the other hand, the arrival of a liquidity trader results in the last price, the closing purchase is at a price other than the true price and results in a non-zero expected gain. If the arrival of a liquidity trader causes S to be above P_0, the closing purchase results in:

$$q \cdot \int_{P_0}^{\infty} (P_0 - S) f(S \mid L) dS \leq 0$$

If the arrival of a liquidity trader causes price S to be below P_0 but above the limit price P_{lim}, the closing purchase results in:

$$q \cdot \int_{P_{\text{lim}}}^{P_0} (P_0 - S) f(S \mid L) dS \geq 0$$

Of course, the price S cannot be below P_{lim} as it would have resulted in the limit order executing. Thus, recalling that S is distributed symmetrically around P_0, the limit order trader's net expected gain from the closing purchase is:

Expected Gain from a Closing Purchase:

$$q \cdot \int_{P_0 + (P_0 - P_{\text{lim}})}^{\infty} (P_0 - S) f(S \mid L) dS \leq 0$$

This means that an investor expects to lose from a closing purchase if the probability of arrival of a liquidity trader is non-zero. This also means that if $q \cdot F > 0$, an investor will choose to purchase at the end of the trading window only if the loss from foregoing a purchase exceeds the loss from a closing purchase.

We define:

> **Non-execution Cost:** When there are no liquidity driven changes in transaction price (i.e., $q \cdot F = 0$), a limit order trader faces a zero expected gain from a closing purchase at the end of a trading window. On the other hand, fluctuation in transaction price due to liquidity motivated trades (i.e., $q \cdot F > 0$) results in a negative expected gain from a closing purchase, which we refer to as a positive *non-execution cost.*

We now write the expected returns from the limit order strategy, conditional on execution and non-execution, relative to the returns from a market order strategy. We use the notation R_m for the return to the market order strategy and the notation R_l for the return to the limit order strategy. When we condition a return on execution we use the notation e and when we condition on non-execution we use the notation n. Also $E(.)$ denotes the expectation operator at the start of the trading window.

When $q{\cdot}F$ is zero, our earlier result was a negative expected gain from a limit order execution, i.e. a positive *bagging cost*. The expected gain from a closing purchase was zero, i.e., a zero *non-execution cost*. Hence, for $q{\cdot}F = 0$, the relative return from the limit order strategy, conditional upon execution and non-execution is:

$$E(R_l | e) - E(R_m) \leq 0 ,$$

$$E(R_l | n) - E(R_m) = 0 .$$

When $q{\cdot}F$ is sufficiently large, we have a positive expected gain from limit order execution, i.e. a negative *bagging cost*, and a negative expected gain from a closing purchase, i.e., a positive *non-execution cost*. Hence, when the probability of arrival of a liquidity public trader is sufficiently large, the return from the limit order strategy, conditional upon execution and non-execution is:

$$E(R_l | e) - E(R_m) \geq 0 ,$$

$$E(R_l | n) - E(R_m) \leq 0 .$$

2.2 The Viability of An Order Driven Market

In this subsection, we consider the *viability* of an order driven market. In our context, an order driven market is viable if some participants place limit orders and other participants place market orders, so that trades and prices are established. We first define mean reversion in transaction price in the context of our model:

Mean Reversion in Transaction Price: In the context of our framework, a positive probability that a liquidity public trader arrives in a trading window and causes a limit order to execute (i.e., $q \cdot F > 0$), means that the transaction price reverts to the true price after the trade. We refer to this phenomenon as mean reversion in transaction price. *Sufficient* mean reversion, corresponding to a sufficiently high value of $q \cdot F$, gives a positive expected gain from limit order execution.

We show that, if the transaction price follows a mean reverting process, an order driven market is viable: relatively patient participants submit limit orders and relatively eager participants submit market orders. Conversely, a viable order driven market requires that the transaction price follow a mean reverting process. For the market ecology to be in balance, the mean reversion in transaction price must be just *sufficient* to compensate the marginal limit order trader adequately.

Viability of an Order Driven Market: An order driven market is viable if and only if transaction price is mean reverting. When an order driven market is viable, participants whose cost of foregoing trade is small will supply liquidity. Conversely, participants whose cost of foregoing trade is large will demand liquidity.

We present a verbal argument to preserve the simplicity of the result. First consider the case where $q \cdot F > 0$ and is sufficiently large so that the expected gain from limit order execution is positive. In this case, an investor whose cost of foregoing trade at P_0 or a higher price is small, can choose to place a limit order below P_0. For this investor, there is a positive expected gain from limit order execution; further, the negative expected gain from a closing purchase can be avoided by not trading at the end of the trading window. Hence, an investor whose cost of foregoing trade is small (who we may refer to as a *patient* trader) may choose to trade by limit order. This is not the case with an investor whose cost of foregoing trade at P_0 or a higher price is large, and who will not forego a trade at the end of the trading window. Such an investor (who we may refer to as an *eager* trader) will choose to trade by market order if the expected loss of a closing purchase is large and offsets the expected gain from limit order execution. In conclusion, with sufficiently large $q \cdot F$, some participants will place limit orders and others will place market orders making an order driven market viable.

Conversely, if $q \cdot F = 0$ an order driven market is not be viable. If $q \cdot F = 0$, we showed in the last section that the expected gains from limit order execution are negative and the expected gains from a closing trade are zero. This means that the unconditional expected gains from trading by limit order are negative for all investors. Hence, no investor will trade by limit order, making an order driven market not viable.

3. Empirical Evidence on Limit Order versus Market Order Trading

Limit orders are a way of capturing the higher level of short-run price volatility that is caused by liquidity driven price changes that are temporary, and our experiments may be interpreted as a *joint* test of (1) the limit order trading strategy and (2) the underlying price process. Our primary objective is, of course, to assess the profitability of limit order trading.

We conduct our experiments by replaying the transaction record and assessing the profitability of entering experimental, one-share market and limit orders. The one-share orders are certainly small enough so that, if they had in fact been entered, they would not have altered the transaction record. Our entries should be interpreted as marginal orders, and their profitability would suggest that the intra-marginal orders that did actually execute at the same price were indeed profitable.

Our experiments were run for the thirty Dow Jones Industrial firms which trade on the New York Stock Exchange (NYSE). The primary data source is the 1988 "Trades and Quotes" transaction file for NYSE stocks supplied by the ***Institute For The Study of Security Markets***. The importance of these firms and the frequency with which their shares trade make them particularly suitable for testing the limit order model.

Hamon, Handa, Jacquillat and Schwartz (1993) report generally similar results for securities traded on the Paris Bourse. The NYSE, unlike the Paris bourse, however, is not a pure order driven market in that it includes a dealer (the specialist). But NYSE specialists participate in only thirty percent or so of all trades.[6] Specialists also have a negative obligation to give a public order priority when it is tied in price with the market maker's quote. Nonetheless, we caution that the specialist's ability to change his or her quotes when sensing that an informational event rather than a liquidity event has occurred, could put limit order traders at a relative disadvantage on the NYSE.[7]

3.1 Experimental Design

Our theoretical model considers the returns realized over an investment window after a purchase has been made using a market order or a limit order strategy. Our empirical analysis, which is based on this framework, requires the specification of three parameters: the length of the trading window (x days), the difference (expressed as a percent) between the current price (P_0) and the limit order price (l%), and the length of the investment window (t days).[8] We examine four limit order trading strategies corresponding to $l = 0.5\%$, 1%, 2% and 3%.[9] We choose a trading window of $x = 1$ day for $l = 0.5\%$ and 1%, $x = 2$ days for $l = 2\%$, and $x = 3$ days for $l = 3\%$. We avoid a window of less than one day so as to be consistent in

---

[6]We thank Jim Shapiro at the NYSE for having provided us with this number.

[7]See Rock (1995).

[8]One of these parameters, the length of the investment window, need not be specified in our multiple limit order tests reported in Section 4 below.

[9] Supplemental tests for $l < 0.5\%$ and $l > 3\%$, not reported here, yielded consistent results.

determining the opening and closing prices in a trading window. In addition, we use a larger trading window when the limit order is placed further from the opening price so as to allow more time for the liquidity and information events to occur that may result in an execution. We hold the investment window constant at t = 3 days to facilitate comparison of returns from the different trading strategies.

In our tests, a market buy order at a purchase price equal to P_0 on day 0 (the first day of the trading window) is compared with a limit buy order placed *l*% below P_0 (rounded to the nearest eighth). The limit order is followed until it executes or until the last price in the trading window is reached. If the limit order does not execute during the trading window, it is purchased at the opening price on the day following the trading window. For each trading window, all prices were standardized by setting the opening price on day 0 equal to 100 and re-scaling the ensuing price series.

The return to the market order is defined as the log of the opening price on day 3 minus the log of the opening price on day 0. The limit order return is defined as the log of the opening price on day $d+3$ minus the log of the purchase price. For shares bought *during* day $d-1$, $0 \le d-1 \le x-1$, the purchase price is the limit order price (P_{lim}) or the opening price of the day, if the opening price is below P_{lim}. For $d = x$, the purchase price is the opening price of day x. We ignore cash dividends on stocks and interest returns on cash balances; for interest rates greater than dividend yields, this results in the returns to limit order trading being understated.

Transactions data for the first 250 trading days of 1988 were used. The year was partitioned into ten 25-day sub-periods. Each sub-period was further sub-divided into windows. For the 0.5% and 1% strategies, there were eight windows per stock per sub-period; for the 2% strategy, there were five windows per stock per sub-period; and for the 3% strategy, there were four windows per stock per sub-period.

3.2 Transaction Prices

The average standardized purchase prices for the four limit order strategies are shown in Table 1. Panel A of the table reports overall and sub-period results for all (executed and un-executed) limit orders. For the 0.5% limit order strategy, the average standardized purchase price ranged from 99.728 to 100.522 for the sub-periods, and was 100.034 overall. We tested the null hypothesis that the average limit order purchase price was equal to the market order purchase price of 100, both for the individual sub-periods and overall. For each sub-period, we ran t-tests using the mean and variance of price across firms and across windows. One of the sub-period averages is significantly below 100, one is significantly above 100, and the remaining eight differ insignificantly from 100. We also tested the hypothesis that the limit order purchase price equals 100 using the grand mean and variance of price across all firms and sub-periods and could not reject it. Finally, we ran the non-parametric Wilcoxon signed rank test using sub-period means to test the null hypothesis that the overall median purchase price was 100. The p-value is 0.231, which again does not enable us to reject the null hypothesis.

For the 1% strategy, the overall purchase price is 99.987 and differs insignificantly from 100. Two sub-period average purchase prices are significantly below 100 and the rest are insignificantly different from 100. The results for the 2% limit order strategy are similar. One sub-period average is significantly below 100, one is significantly above 100, and the remaining seven are insignificantly different from 100. The overall purchase price is 100.019 and not significantly different from 100. The 3% limit order strategy yields one sub-period average significantly below 100, three significantly above 100, and the remaining six are insignificantly different from 100. The overall purchase price is 100.136 and not significantly different from 100. The Wilcoxon signed rank test gives a p-value of 0.094, and we cannot reject the null hypothesis at the 5% level of significance.

The purchase price for limit orders that executed was distributed around 100(1-l%) because of price rounding, and these values are not shown in the Table. Panel B of Table 1 gives sub-period and overall means for limit orders that did not execute. Ex-ante, the purchase price of an un-executed limit order, conditional on the fact that price did not drop *l*% or more during the trading window, is expected to be greater than 100. For the 0.5% limit order strategy, the purchase price ranged from 100.460 to 105.122, and had an average value of 101.299 which is significantly above 100. For the 1% limit order strategy, the purchase price ranged from 100.225 to 103.748, while the overall average was 100.885 which is significantly above 100. Similar results were obtained for un-executed limit orders based on the 2% and 3% limit order strategies; the overall purchase price was 100.854 and 101.037, respectively, and significantly above 100 in both cases.

In sum, the unconditional purchase price of all limit orders was not significantly different from the market order purchase price. On the other hand, the purchase price of un-executed limit orders was significantly higher than the overall market order purchase price of 100.

3.3 Returns

We define the following variables for each stock to measure and to compare returns for the limit and market order strategies over the 3-day investment windows:[10]

r_o^m : the unconditional return on the market order strategy.

r_o^l : the unconditional return on the limit order strategy.

r_e^l : the return on the limit order strategy conditional on execution.

r_n^l : the return on the limit order strategy conditional on non-execution.

Table 2 presents results on average limit order and market order returns, and on the percentage of limit orders that executed. The returns are aggregated over the ten sub-periods in 1988. The unconditional market order return is 0.039% for the 0.5% and 1% strategies, -0.081% for the 2% strategy, and 0.375% for

[10] The stock index *i* is suppressed to conserve notation.

the 3% strategy; it is significantly above zero for only the 3% strategy.[11] The comparable unconditional limit order return is 0.209% for the 0.5% strategy, 0.259% for the 1% strategy, 0.273% for the 2% strategy, and 0.295% for the 3% strategy, and it is always significantly above zero. Hence, the limit order strategy outperforms the market order strategy for the 0.5%, 1% and 2% tests, and under-performs in the 3% test. Percent executions partially account for this result: for the 0.5% test, 67.4% of the limit orders execute; for the 1% test, the percentage is 45.8%; for the 2% test, the execution percentage drops to 28.9%; and for the 3% test, only 22.7% of the limit orders execute.

The average limit order return conditional on execution increases in the percentage parameter, and is always significantly higher than both the unconditional limit order and the unconditional market order return. It is 0.288% for the 0.5% strategy, 0.429% for the 1% strategy, 0.532% for the 2% strategy, and 1.820% for the 3% strategy, and always significantly different from zero at the 1% level. The average limit order return conditional on non-execution is always lower than the overall limit order return, but not necessarily lower than the unconditional market order return. It is 0.046% for the 0.5% strategy, 0.115% for the 1% strategy, 0.167% for the 2% strategy, and -0.152% for the 3% strategy, and always insignificantly different from zero.

Overall, the results are indicative of short-run negative auto-correlation in returns. Stock price changes appear to be noisy signals of informational change, and have a tendency to revert to previous levels. The exaggerated price movements during a trading window are presumably liquidity driven. Our theoretical model suggests that the noise may, in part, be endogenous. That is, if an insufficient percentage of limit orders execute, fewer will be placed and short-run price volatility will increase, until the marginal limit order trader is adequately compensated.

---

[11] Average market order returns can differ for the four limit order tests because the specific calendar days included in these tests are not identical.

3.4 Bagging

As discussed in Section 2, if prices follow random walk, limit order traders suffer from the winner's curse: a buy limit order executes if and only if the equilibrium price jumps to or below the limit order price, a phenomenon referred to as bagging. With bagging, if the limit price is reached in a trading window, the return on the limit order strategy is expected to be lower than the *unconditional* market order return. Table 3 presents the difference between the limit order return *conditional on execution*, and the *unconditional* market order return (where, for each stock, the unconditional market order return is the average market order return in a given sub-period). Overall, the differential limit order return is 0.332% for the 0.5% strategy, and is significantly greater than zero at the 1% level and with a Wilcoxon test p-value of 0.004. It is positive for all but one sub-period and significantly greater than zero for two sub-periods.

The differential limit order return for the 1% strategy is 0.561% which is significantly above zero at the 1% level and has a Wilcoxon p-value of 0.003. Again, it is positive for all but one sub-period and significantly greater than zero for four sub-periods. For the 2% strategy, the differential limit order return increases to 0.892%, and is significantly greater than zero at the 1% level with a Wilcoxon p-value of 0.009. As with the other tests, it is positive for all but one sub-period and is significantly greater than zero for seven sub-periods. Finally the 3% strategy yields a differential limit order return of 1.786% which is significantly greater than zero at the 1% level and has a Wilcoxon p-value of 0.001. It is positive for all but one sub-period and is significantly greater than zero for six sub-periods.

The differential limit order returns conditional upon execution, which are consistently positive and steadily increase from the 0.5% test to the 3% test, are inconsistent with the hypothesis that a positive bagging cost dominates the gains from liquidity events. There is no evidence that executed limit orders, on average, earn lower returns than the benchmark return. On the contrary, it appears that limit order executions occur due to liquidity driven price changes with sufficient frequency, and that prices tend to rebound over relatively short investment horizons. This being the case, why do not all investors find it profitable to place limit orders? The answer, provided in the next section, involves the non-execution cost.

3.5 Non-Execution Cost

We capture the non-execution cost by forcing the investor to buy at the market price prevailing at the end of a trading window. We compare returns to the limit order strategy *conditional on non-execution* to the *unconditional* market order return (where, for each stock, the unconditional market order return is the average market order return in a given sub-period). Table 4 presents the differential returns over sub-periods and overall for the four strategies. Overall, the differential limit order return is -0.164% for the 0.5% strategy. This is insignificant based on the t-test and with a Wilcoxon test p-value of 0.344. The differential return is negative for six sub-periods, including one sub-period where it is significantly negative.

For the 1% strategy, the differential limit order return is -0.067%, which is insignificant based on the t-test and with a Wilcoxon test p-value of 0.212. It is negative for four sub-periods, and significantly so for one sub-period. For the 2% strategy, the differential limit order return is 0.132%, and insignificant with a Wilcoxon test p-value of 0.156. It is negative for four sub-periods, significantly greater than zero for three sub-periods, and significantly less than zero for two sub-periods. Finally, for the 3% strategy, the differential limit order return of -0.626 is significantly less than zero based on the t-test and with a Wilcoxon test p-value of 0.452. It is negative for eight sub-periods, significantly so for five, and significantly positive for two sub-periods.

In conclusion, the evidence suggests that there is a cost of non-execution, although it is not consistently significant. The cost is negative for the 2% strategy (as the differential return is positive) but close to zero. The cost is significantly positive for the 3% strategy. Thus, an eager investor (with a high intensity to trade) is unlikely to profit from the limit order strategy since such an investor faces the cost of non-execution. A patient investor (with a well-balanced portfolio), however, can avoid the cost of non-execution by not trading at all if the limit order does not execute.

3.6 *Overall Differential Returns*

The unconditional differential limit order returns $r_o^l - r_o^m$ are reported in Table 5 for the four strategies. The differential limit order return to all limit orders is 0.171% for the 0.5% test, and positive for 8 sub-periods (being positive and significant for two). It is significantly greater than zero at the 1% level with a Wilcoxon test p-value of 0.012. The corresponding return for the 1% test is 0.220%, being positive for 9 out of the 10 sub-periods. Again the return is significant at the 1% level and has a Wilcoxon test p-value of 0.007. For the 2% test, the overall differential limit order return is 0.355%, being positive in six sub-periods and significantly greater than zero at the 1% level but with a Wilcoxon test p-value of 0.069 which is not significant at the 5% level. For the 3% test, the overall differential limit order return is a negative -0.08%, being negative in seven sub-periods. The return is insignificantly different from zero based on its t-statistic and has a Wilcoxon test p-value of 0.351.

Overall, the limit order strategy performs as well or better than the market order strategy. High non-execution costs and low probability of execution can lead to lower limit order returns as in the 3% test. On the other hand, positive differential returns appear to be driven by limit order executions that are followed by a rebound in prices as opposed to a bagging cost. The 0.5%, 1% and 2% tests fall in this category. The superior performance of limit order returns for some traders (those who are relatively patient) may be deduced from the behavior of the unconditional differential returns.

3.7 *Market-adjusted Returns*

In this section, we compare market-adjusted returns to the limit order strategy with market-adjusted returns to the market order strategy. The objective is to determine whether the differential performance of limit orders and market orders for individual stocks is attributable to order flow dynamics that are unique to individual stocks, or to the behavior of the aggregate order flow for the broad market. That is, we examine the extent to which the execution of a limit order for a specific stock at a particular point in time is the outcome of a liquidity event specific to that stock, or to a more pervasive, mean-reversion producing order flow imbalance that causes limit orders on the same side of the market to execute for a wide cross-section of stocks. We start by defining the differential return for stock *i* in window *t* as

$$R_{it} = r_o^{\,l} - r_o^{\,m}$$

where $r_o^{\,l}$ is the overall limit order return of stock *i* in window *t*, and $r_o^{\,m}$ is the average market order return of stock *i* in the sub-period that contains window *t*. Similarly, for an equally-weighted portfolio we define

$$R_{pt} = r_{po}^{\,l} - r_{po}^{\,m}$$

where $r_{po}^{\,l}$ is the overall limit order return of the portfolio in window *t*, and $r_{po}^{\,m}$ is the average market order return of the portfolio in the sub-period that contains window *t*. While R_{it} can be interpreted as the excess return for stock *i*, R_{pt} can be interpreted as the excess return for the portfolio. We test the relationship between R_{it} and R_{pt} by running the regression,

$$R_{it} = \alpha_i + \beta_i R_{pt} + \eta_{it}$$

where η_{it} is the differential return for a stock that is un-correlated with the differential return on the market portfolio. Because our portfolio comprises only 30 stocks, we eliminate the effect of each stock's own return on the portfolio return by constructing a customized index for each stock, using only the returns for the other 29 stocks in the sample.[12] For stock *i*, a positive covariance term, β_i would suggest that the profitability of a limit order as compared to a market order in a trading window is in part attributable to

[12] This procedure, known as the "Lachenbruch method," is described in Lachenbruch (1967).

limit orders in aggregate having outperformed market orders for that period. The residual term, η_{it} captures the component of limit order profitability that is specific to stock *i*.

An analysis of η_{it} is interesting for two reasons. First, removing the common market element for each stock allows us to view the observations for the 30 companies as independent samples, which increases our confidence in the significance of the findings. Second, one might expect that liquidity events are unique to individual stocks, and that they cause limit order execution to be profitable on a stock-by-stock basis. To test this hypothesis, we examine the behavior of the market-adjusted component of differential returns by running the following panel regression:

$$\eta_{it} = \gamma_1 + \gamma_2 L_{it} + \varepsilon_{it}$$

where *i* represents the stock, *t* represents the window, and

η_{it} is the market-adjusted differential return,

$L_{it} = 1$ when the limit price is reached in a window and $L_{it} = 0$ when the limit price is not reached,

ε_{it} is an error term.

Hence, a stacked vector of the market-adjusted differential returns *for all stocks* is regressed on a stacked vector of indicator values L_{it}, to obtain estimates of the coefficients γ_1 and γ_2, and their standard errors.

As in our previous analysis, $\gamma_1 = 0$ and $\gamma_2 < 0$ is consistent with random walk and with limit order execution, on expectation, being costly. On the other hand, $\gamma_1 < 0$ and $\gamma_2 > 0$ is consistent with mean reversion. Finally, $\gamma_1 + \gamma_2 > 0$ suggests that mean reversion is sufficiently positive for the limit order strategy to be profitable even when closing transactions are forced.

The overall and sub-period results after we control for the systematic market component in limit order trading, are shown in Table 6. These results provide clear and strong support for the hypothesis that trading via limit orders provides gains from execution and has costs when a limit order fails to execute. For the 0.5% test, the market-adjusted differential return for limit orders that execute (given by $\gamma_1 + \gamma_2$) is 0.164%, while the market-adjusted differential return for limit orders that do not execute (given by γ_1) is

-0.341%. The parameters γ_1 and γ_2 are well-behaved: γ_1 is significantly less than zero overall and takes a negative value in nine of the ten sub-periods, while γ_2 is significantly greater than zero overall, is greater than γ_1 in nine sub-periods, and takes a positive value in the same nine sub-periods.

For the 1% test, the market-adjusted differential return for limit orders that execute is 0.252% and for limit orders that fail to execute it is -0.213%. The parameter γ_1 is significantly less than zero overall and is negative in eight sub-periods, while γ_2 is significantly greater than zero overall, is positive in nine sub-periods, and exceeds γ_1 in eight sub-periods. For the 2% test, the market-adjusted differential return for limit orders that execute is 0.509%, and the market-adjusted differential return for limit orders that do not execute is -0.209%. The parameter γ_1 is significantly less than zero overall and is negative in all but one sub-period, while γ_2 is significantly greater than zero overall, is positive in nine sub-periods and greater than γ_1 in the same sub-periods. Finally, for the 3% test, the market-adjusted differential return for limit orders that execute is 1.155% and the market-adjusted differential return for limit orders that do not execute is -0.339%. The parameter γ_1 is significantly less than zero overall and is negative in all ten sub-periods, while γ_2 is significantly greater than zero overall, and is greater than γ_1 and positive in all sub-periods.

The results on adjusted differential returns support the hypothesis that differential limit order returns conditional on execution are positive, and that differential limit order returns conditional on non-execution are negative. We find, after controlling for the common market effect in individual stock returns, that bagging is not a cost to limit order traders, but that non-execution is. Again we conclude that investors who are particularly eager to transact may prefer the market order strategy, while investors who gain relatively little by trading at current prices (and who are thus willing to risk not executing) may prefer the limit order strategy.

3.8 Robustness Tests

We have thus far tested the profitability of buying by limit orders using transaction data for the year 1988. This particular year was marked by generally rising stock prices and the impact of trend in prices is not clear. A possibility is that an upward trend could result in a reduction in bagging costs and an increase in non-execution costs for buy orders.

In one simple test of robustness, we examine the profitability of a multiple limit order strategy where a hypothetical trader places both buy and sell limit orders for an extended period of time. This test is discussed in detail in the next section. In this section, we focus on robustness tests that use a very similar experimental design to the results reported in the earlier sections.

In our first test of robustness, we invert the price series, treating the last price observation as the first, the second to last price observation as the second, etc. We examine the profitability of a buy limit order strategy on this inverted price series using a test design identical to the one reported in the earlier sections. One advantage of using the inverted price series is that we are able to preserve both the experimental design and the correlation structure of the price series, while altering the trend in prices. In an alternate test that we perform, the profitability of short-selling by limit order is also examined during the trading period, with the short position subsequently covered at the end of the investment window. In this formulation of the robustness test, the experimental design assumes that there are no constraints or requirements to short-selling, and that short-selling can proceed in a manner similar to purchasing stock.

Table 7 compares the relative returns we obtained from a regular (non-inverted) buy test, with the relative returns from a buy limit order strategy using the inverted price series, and with the relative returns from a short-selling limit order strategy. The results in Table 7 should be interpreted in two dimensions. First, in equilibrium, we expect the overall return to trading via limit order to be similar to the return to trading via market order. Second, we expect that mean reversion in prices will cause executed limit orders to perform better than market orders, and un-executed limit orders to perform worse than the market orders.

An examination of all limit orders that execute, irrespective of the strategy being tested, reveals that such orders either perform significantly better or are not significantly different from a market order. Second, an examination of limit orders that fail to execute, irrespective of the strategy tested, reveals that such orders either perform significantly worse or are not significantly different from a market order. Third, the results on the overall performance of limit orders are as expected mixed: in the inverted buy tests, limit orders out-perform market orders for the 1% and 2% strategies, but under-perform in the 3% case; in the short-sell tests, limit orders out-perform market orders for the 3% strategy, but under-perform in the 1% and 2% tests. Overall, we conclude that trending does not influence the results reported in the earlier sections. This interpretation is further confirmed by the results reported in Section 4 which follows.

4. Profitability of a Multiple Limit Order Trading Strategy

Our tests have thus far assessed the profitability of either buying or selling by limit order, and of holding the position for a relatively short period before closing it out. We now assess the profitability of placing a network of buy and sell limit orders, and of allowing a sequence of purchases and sales to occur over an extended calendar period before closing out any accumulated position that may have developed. The multiple limit order test we use to assess the profitability of limit order trading is structured as follows.

For each Dow stock, our hypothetical trader casts a network of limit orders around the opening price for the first trading day of a test period. The orders are placed a fixed distance apart where the distance is measured as a percentage (s = 1%, 2%, 3%, 4%, 5%) of the opening price. That is, if the opening price is P_0, limit orders to buy are entered at $(1-s)P_0$, $(1-2s)P_0$, $(1-3s)P_0$, etc., and limit orders to sell are entered at $(1+s)P_0$, $(1+2s)P_0$, $(1+3s)P_0$, etc. All entered orders are rounded to the nearest eighth. The dollar spread implied by the parameter s depends, of course, on the price level at which the stock is trading. For the 1% strategy, the unrounded dollar spreads average \$0.48 and range from \$0.05 to \$1.58. As s increases to 5%, the average, minimum and maximum rise monotonically to, respectively, \$2.42, \$0.26 and \$7.08.

As executions are realized, limit orders are appropriately reinstated to keep the network intact and centered on the price of the last limit order that executed. We restrict our tests to a naive limit order strategy, without allowing for a revision in the quotes in response to current market conditions, and without any inventory control. We follow the trades, inventory levels, and profits for three 80-day periods during the year 1988.[13] The total return to the multiple limit order strategy can be written as the sum of all proceeds from sales less the sum of all payments for purchases. Purchases and sales are made as price fluctuates during a test period and, because price can drift, the trader generally holds an unbalanced inventory at the end of each period. Consistent with the design of our single (buy or sell) limit order tests, returns for the multiple limit order strategy are measured in two ways: (1) as the gain per round trip (referred to as RTGAIN later in this section) that is realized as prices fluctuate during a test period, not including the closing (inventory re-balancing) transaction, and (2) as a limit order spread (referred to as LOS later in this section) which includes the closing trade.[14] The limit order spread is obtained by dividing the total dollar return including the closing transaction, by the total number of roundtrips including the closing transaction. A positive bagging cost is, thus, manifest in a negative RTGAIN, and a positive non-execution cost is manifest in an LOS less than RTGAIN.

More formally, we use the following notation to assess the returns to the limit order strategy:

---

[13] Three periods of eighty trading days each were used as data were available for all stocks in our sample for a period of 242 days in 1988.

[14] This is similar to the standard bid-ask spread which gives the per share profits of one round trip, i.e., buying at the bid and selling at the ask. The limit order spread can be interpreted as the returns per share by buying as limit bids execute and selling as limit asks execute, followed by a re-balancing to restore the opening inventory position.

Na = number of shares sold during the test period (excluding sales at the closing trade),

Nb = number of shares bought during the test period (excluding purchases at the closing trade),

Pa = average selling price per share over all trades accept the closing trade,

Pb = average buying price per share over all trades accept the closing trade, and

Pc = closing price per share.

Letting Π be the total return, we have:

$$\Pi = (P_a \cdot N_a - P_b \cdot N_b) - P_c \cdot (N_a - N_b)$$

Rearranging the equation, we can write the total return as:

$$\Pi = (P_a - P_b) \cdot Min(N_a, N_b) + (\overline{P} - P_c) \cdot (N_a - N_b)$$

$$\text{where } \overline{P} = P_a \text{ if } N_a > N_b$$
$$\text{and } \overline{P} = P_b \text{ if } N_a < N_b.$$

and the limit order spread (*LOS*) as

$$LOS = \frac{\Pi}{Max(N_a, N_b)}$$

The first term on the right hand side of the expansion for Π, that is, $(P_a - P_b) \cdot Min(N_a, N_b)$ captures the *returns* from the round trips.[15] The *bagging cost* is given by $-(P_a - P_b) \cdot Min(N_a, N_b)$. Similarly, the second term, $(\overline{P} - P_c) \cdot (N_a - N_b)$, gives the *returns* from closing the position at the end of the period. Therefore, the *non-execution cost* is given by $-(\overline{P} - P_c) \cdot (N_a - N_b)$.

In the case of bagging, the difference between the average ask price and the average bid price, $(P_a - P_b)$ establishes whether bagging is, in fact, a positive cost. If bagging is a positive cost, $(P_a - P_b)$ is, on expectation, negative, and if bagging is a negative cost (because of mean reversion in prices) this

[15] We are grateful to the referee for suggesting this line of analysis.

component is, on expectation, positive. To see this more clearly, consider the following. If during a test period, news is on net bearish, limit buy orders will execute as price falls and P_b will, on average, be greater than P_a. Alternatively, if news is on net bullish, limit sell orders will execute as price rises and P_a will, on average, be less than P_b. Hence, when prices follow a random walk $(P_a - P_b)$ will, on average, be zero. On the other hand, with mean reversion, prices will tend to rise after limit buy orders execute and to fall after limit sell orders execute, leading to P_a greater than P_b. For the empirical results that follow, we represent the bagging cost as

$$BC = (-1) \cdot RTGAIN \cdot RTRIPS$$

where BC = Bagging Cost
$RTGAIN$ = Gain per Round Trip = $(P_a - P_b)$
and $RTRIPS$ = Number of Round Trips = $Min(N_a, N_b)$

Now consider the non-execution cost. Any buy imbalance that develops during the test period can be eliminated either by sell limit orders executing during the test period, or by a closing sale at the end of the test period. Similarly, any sell imbalance cab be eliminated either by buy orders executing during the test period, or by a closing purchase at the end of the test period. The component $(N_a - N_b)$ is the size of the closing trade at the end of a test period. Rising prices cause $(N_a - N_b) > 0$ and $(\overline{P} - P_b) < 0$, and falling prices have the reverse effect; hence $sign[(N_a - N_b)] = -sign[(\overline{P} - P_c)]$ and the non-execution cost is, on expectation, positive. We use the following notation to represent the non-execution cost:

$$NEC = (-1) \cdot DC \cdot IMB$$

where NEC = Non - execution Cost

DC = Differential Per Share at Closure = $(\bar{P} - P_c)$

and IMB = Share Imbalance = $(N_a - N_b)$

Note that, if prices do not follow a mean reverting process, the hypothetical limit order trader is doubly cursed: a bagging cost is incurred when bullish or bearish news causes orders predominantly on one side or the other of the market to execute, and a non-execution cost is incurred when the imbalance is eventually restored. On the other hand, with mean reversion, the limit order strategy may be profitable (that is, a negative bagging cost can outweigh the positive non-execution cost). If so, limit orders may be used as a *volatility capture* trading strategy.

Table 8 reports the total return from the multiple limit order test as well as the bagging and non-execution costs associated with it, for the five limit order strategies (s = 1%,...,5%). It also reports the (implied) limit order spread for each of the five strategies. Consistent with the results reported in the previous sections, the bagging cost (BC) is negative, i.e., there is a positive return rather than a loss from executing limit orders. Moreover, this gain is statistically significant for all five strategies. Examining the components of this gain, the round trip gain (RTGAIN) is always positive and increasing in the spread parameter. The number of round trips (RTRIPS) declines rapidly in the spread parameter but is always significantly greater than zero, i.e., our hypothetical trader does indeed get an opportunity to reverse some of the trades.

The non-execution cost (NEC) captures the closing trade at the end of a test period and this cost turns out to be always positive. The first component of the non-execution cost is the differential per share at closure (DC), a number which, on average for our sample, is negative though not significantly different from zero. The second component of the non-execution cost is share imbalance (IMB) and this number is always positive, indicating a tendency to sell more than to buy in a period. Overall, because prices were generally rising in 1988, the hypothetical trader oversells during the test periods and has to regain a zero inventory position by purchasing at the higher closing price. This results in a loss at closure, i.e. in a positive non-execution cost. Both of the components of non-execution cost have highly skewed distributions (with a few heavy losses and several smaller gains as well as losses at closure).

The net profitability of the multiple limit order strategy is clearly brought out by the *limit order spread* numbers. As one moves from a 1% to a 5% strategy, the *limit order spread* (LOS) increases steadily from $0.23 to $1.24. This result is our strongest confirmation that mean reversion exists in prices in an order driven market, and that it makes limit order trading profitable for some public participants.

The *limit order spread* is a function of two countervailing forces: a negative *bagging cost* which results in the hypothetical trader actually gaining from the execution of his or her limit orders, and a positive *non-execution cost* which results in a loss to the trader if he or she closes a position to return to zero inventory. In order to better understand how the limit order spread is determined, we studied the correlation between its various components. The bagging cost arises when a trader, on average, sells low and buys high (as captured by the term RTGAIN), and on the number of round trips made within a period, RTRIPS. In our sample, a negative bagging cost arose when the trader, on average, bought low and sold high, i.e., RTGAIN was positive, and, moreover, some round trips were actually achieved, i.e., RTRIPS was significantly greater than zero. The correlation between RTGAIN and RTRIPS was -0.336 and not significantly different across strategies, which indicated that a lower per share gain is associated with a higher number of round trips. The limit order spread (LOS) was only weakly negatively related to RTRIPS (correlation of -0.174) but strongly positively related to RTGAIN (correlation of 0.619) and RTGAIN seemed to drive the gains for the trading strategy.

The non-execution cost arises when the closing price is high (i.e., DC is negative), and the trader does buy at close (i.e., IMB is positive), or vice versa. The evidence is that DC and IMB are, indeed, highly negatively correlated (correlation of -0.558 for all strategies pooled together, and not significantly different across strategies) which is consistent with the significant positive non-execution cost we have documented. Further, the limit order spread is positively correlated with DC (correlation of 0.405) and negatively correlated with IMB (correlation of -0.336), indicating that rising prices which caused a trader to oversell, caused lower limit order spreads in 1988.

We regressed the limit order spread (LOS) on its four components, RTGAIN, RTRIPS, DC and IMB, and on four indicator variables (IND_1, IND_2, IND_4 and IND_5) to capture any systematic differences between the strategies that are not accounted for by the four component variables.[16] The indicator variable IND_1 took the value 1 when the 1% strategy was used and the value zero otherwise, and so on. The regression used data pooled across time periods, across stocks and across the five trading strategies. Our objectives were (a) to understand the partial effect of each component variable on LOS after controlling for the other component variables and after controlling for the specific trading strategy employed, and (b) to assess whether LOS had any systematic component that depended on the specific trading strategy employed which was not explained by the four component variables. We expected the four component variables to explain a significant portion of LOS even across strategies, and we expected the indicator variables not to capture any systematic effects beyond those captured by the four component variables. Our regression results are presented below. The dependent variable LOS, and the independent variables RTGAIN and DC are measured in cents, and the independent variables RTRIPS and IMB in number of shares.

$$LOS = \underset{(-1.98)}{-25.68} + \underset{(16.90)}{0.64} \cdot RTGAIN + \underset{(1.24)}{0.24} \cdot RTRIPS + \underset{(9.91)}{0.24} \cdot DC + \underset{(0.82)}{0.93} \cdot IMB$$
$$- \underset{(-0.10)}{2.12} \cdot IND_1 + \underset{(0.43)}{6.81} \cdot IND_2 + \underset{(0.56)}{8.96} \cdot IND_4 + \underset{(0.46)}{7.37} \cdot IND_5$$
$$R^2 = 0.529$$

[16] As our primary interest was to get some insights into partial correlations between LOS and its component variables, for simplicity we restricted the regression test to a linear fit.

The regression clearly indicates that the key components driving the limit order spread in 1988 were RTGAIN and DC. The number of round trips within a period (RTRIPS) and the share imbalance at closure (IMB) had insignificant effect on the limit order spread. Further, the indicator variables were all insignificant as expected, indicating that no systematic differences exist beyond those captured by RTGAIN and DC, in the determination of the limit order spread across the five different multiple limit order strategies employed. Hence, the significant difference that exists in the limit order spread (LOS) across strategies ($0.23 for the 1% test monotonically increasing to $1.24 for the 5% test), is explained largely by RTGAIN and DC.

5. Conclusion

The viability of an order driven market depends on limit order trading being profitable for a sufficient number of public participants. When transaction prices change solely in response to information, trading via limit order is sub-optimal for all traders because the advent of adverse news can trigger an undesired trade, while the advent of favorable news can result in the limit order not executing. We suggest that limit orders are placed by public traders only because liquidity driven price changes (caused by temporary order imbalances in the market) can offset the cost of being bagged by informed traders. Furthermore, in an order driven market, a paucity of limit orders can cause temporary order imbalances that lead to short-run changes in transaction price which eventually reverts to true price. Hence, equilibrium levels of limit order trading and of short-run price volatility can exist in the sense that an increase (decrease) in short-run volatility encourages (discourages) the placement of limit orders and, in turn, an increase (decrease) in limit order trading decreases (increases) short-run volatility.

We have tested the limit order trading strategy for the thirty Dow Jones Industrial stocks for the year 1988. The results are generally consistent, both with our theoretical framework and across the ten sub-period tests. Overall, the evidence indicates that limit order purchases (including closing purchases at the end of each trading window) and market order purchases result in comparable prices and returns. We find no evidence of systematic bagging; on the contrary, on average, prices tend to bounce back toward original levels after limit order executions. On the other hand, non-execution costs appear to be positive, but not always statistically significant. We have also examined limit order trading as a *pure* trading strategy and, in so doing, have documented the existence of a *limit order spread* that averages \$0.23 for the 1% strategy and monotonically increases to \$1.24 for the 5% strategy.

Our results show that some proprietary traders (those who, having minimal non-execution costs, are relatively patient) have an incentive to submit limit orders, that others (those who, having high non-execution costs, are relatively eager) prefer to submit market orders, and that this trading ecology can be self-sustaining. These results are consistent with the viability of an order driven market, where public participants supply liquidity to themselves without the intervention of dealers. We believe that future research on issues such as trade size and the profitability of executed limit orders, and on the cross-sectional differences in stock price volatility and the profitability of limit order trading, can provide further useful insights on the functioning of an order driven market.

References

Admati, Anat R. and Paul Pfleiderer, 1989, Divide and conquer: A theory of intraday and day-of-the-week mean effects, *The Review of Financial Studies* Vol 2, No 2, 189-223.

Bagehot, Walter (pseudonym), 1971, The only game in town, *The Financial Analysts Journal* 8, 31-53.

Copeland, Thomas and Dan Galai, 1983, Information effects on the bid-ask spreads, *The Journal of Finance*, 38, 1457-69.

Easley, David and Maureen O'Hara, 1987, Price, trade size and information in securities markets, *Journal of Financial Economics* 19, 69-90.

Glosten, Lawrence, 1989, Insider trading, liquidity, and the role of the monopolist specialist, *Journal of Business*, 62, 211-235.

Glosten, Lawrence, 1994, Is the electronic open limit order book inevitible? *The Journal of Finance*, 49, 1127-1161.

Grossman, Sanford and Merton Miller, 1988, Liquidity and market structure, *The Journal of Finance*, 43, 617-633.

Hamon, Jacques, Puneet Handa, Bertrand Jacquillat and Robert Schwartz, 1993, The profitability of limit order trading on the Paris stock exchange, working paper, Stern School, New York University.

Harris, Lawrence and Joel Hasbrouck, 1993, Market vs. limit orders: The SuperDOT evidence on order submission strategy, working paper, Stern School, New York University.

Lachenbruch, Paul A., 1967, An almost unbiased method of obtaining confidence intervals for the probability of mis-classification in discriminant analysis, *Biometrics*, 639-645.

Leach, Christopher and Ananth Madhavan, 1993, Price experimentation and security market structure, *The Review of Financial Studies*., Vol 6, No 2, 375-404.

Rock, Kevin, 1995, The specialist's order book and price anomalies, *The Review of Financial Studies* forthcoming.

Table 1

Average Standardized Purchase Price Of A Limit Order For The Four Limit Order Tests, For The Thirty Dow Jones Industrial Stocks During 1988.[a]

A. All Limit Orders

		Limit Order Test[b]			
		0.5%	**1%**	**2%**	**3%**
Subperiod	1	100.522*	100.420	100.339	100.588
	2	99.933	99.861	100.398*	99.858
	3	100.124	100.160	100.029	100.125
	4	100.056	100.039	99.409**	99.697
	5	100.083	100.138	99.934	100.878**
	6	99.728**	99.467**	99.891	98.962**
	7	99.913	99.963	99.702	100.226
	8	100.062	100.044	100.562**	100.420*
	9	100.032	99.970	99.875	100.579*
	10	99.890	99.813*	100.003	100.030
Overall 1988		100.034	99.987	100.019	100.136
Wilcoxon p-value[c]		0.231	0.250	0.265	0.094

B. Unexecuted Limit Orders

		Limit Order Test[b]			
		0.5%	**1%**	**2%**	**3%**
Subperiod	1	105.122**	103.748**	102.488**	103.253**
	2	101.273**	100.772**	101.086**	100.883**
	3	101.470**	101.293**	101.423**	100.935*
	4	101.067**	100.796**	100.515**	100.340
	5	101.545**	101.077**	100.557*	101.783**
	6	100.933**	100.295*	100.418**	99.773
	7	100.996**	100.724**	100.634**	101.165**
	8	100.937**	100.621**	100.874**	100.961**
	9	101.002**	100.632**	100.401*	101.108**
	10	100.460**	100.225*	100.526**	100.646**
Overall 1988		101.299**	100.885**	100.854**	101.037**
Wilcoxon p-value[c]					

[a] In general, the purchase price of a limit order is its limit price if the stock price crosses the limit price during the trading window, and is the opening price of the day following the trading window if the limit price is not crossed during that period. All prices are expressed as percentages of the opening price at the start of a trading window (which is the purchase price of the market order).

[b] The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order (x = 0.5, 1, 2, 3) is x% below the opening price, rounded to the nearest one-eighth.

[c] A non-parametric test of the null hypothesis that the median standardized purchase price (of the ten subperiod means) equals 100, the standardized market order price.

* (**) Significantly different from 100, the standardized market order purchase price, at the 5% (1%) level. The sub-period tests use individual subperiod variances. The overall test uses the variance of all observations pooled across subperiods and stocks.

Table 2

Average Return On Market Order And Limit Order Purchases For Dow Jones Industrial Stocks During 1988.[a]

		Limit Order Test[b]			
		0.5%	1%	2%	3%
The average $ size of the limit order test parameter		0.266	0.523	1.037	1.551
Percentage of limit orders that executed		67.4%	45.8%	28.9%	22.7%
Average Return					
All market orders	r_o^m	0.039	0.039	-0.081	0.375**
All limit orders	r_o^l	0.209**	0.259**	0.273**	0.295**
Executed limit orders	r_e^l	0.288**	0.429**	0.532**	1.820**
Unexecuted limit orders	r_n^l	0.046	0.115	0.167	-0.152

a The return reported is the average for the ten subperiods in 1988. Significance tests are based on the variance of returns around their grand mean across all subperiods and stocks.

b The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order (x = 0.5, 1, 2, 3) is x% below the opening price, rounded to the nearest one-eighth.

* Significantly different from 0 at the 5% level.

** Significantly different from 0 at the 1% level.

Table 3

Differential limit order returns when the limit price is reached in a trading window compared to the *unconditional average* market order returns, for the Dow-Jones Industrial stocks for ten sub-periods during 1988.

Differential return to executed limit orders $r_e^l - r_o^m$		Limit Order Test[a]			
		0.5%	1%	2%	3%
Subperiod	1	0.308	0.361	1.626**	1.890**
	2	0.338	0.558	1.584*	3.286**
	3	0.161	0.314	-1.590*	2.152
	4	0.254	1.049**	1.236*	1.732**
	5	0.985**	1.451**	2.527**	3.122**
	6	0.680**	0.974**	0.403	1.619**
	7	0.130	-0.146	0.472	-0.446
	8	0.420	0.786*	1.416	2.313**
	9	-0.168	0.006	1.792**	1.046
	10	0.027	0.084	1.054**	0.798
Overall 1988		0.332**	0.561**	0.897**	1.786**
Wilcoxon p-value[b]		0.004	0.003	0.009	0.001

a The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order ($x = 0.5, 1, 2, 3$) is x% below the opening price, rounded to the nearest one-eighth.

b A non-parametric test of the null hypothesis that the median return (of the ten subperiod means) equals zero.

* (**) Significantly different from zero at the 5% (1%) level. The sub-period tests use individual subperiod variance. The overall test uses the variance of all observations pooled across subperiods and stocks.

Table 4

Differential limit order returns when the limit price is not reached in a trading window compared to the *unconditional average* market order returns, for the Dow-Jones Industrial stocks for ten sub-periods during 1988.

Differential return to unexecuted limit orders $r_n^l - r_o^m$		Limit Order Test[a]			
		0.5%	1%	2%	3%
Subperiod	1	0.323	0.543	-1.537**	-4.134**
	2	0.257	0.221	0.787**	1.537**
	3	-0.129	-0.233	0.122	-1.716**
	4	-1.086**	-1.119**	0.825*	-0.781
	5	-0.284	-0.150	1.593**	-0.037
	6	-0.141	0.502	-0.436	1.129**
	7	-0.090	0.138	-0.005	-0.795**
	8	-0.545	-0.345	-0.462**	-0.827**
	9	0.158	0.031	0.363	-0.975**
	10	0.217	0.241	0.003	-0.263
Overall 1988		-0.164	-0.067	0.132	-0.626**
Wilcoxon p-value[b]		0.344	0.212	0.156	0.452

a The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order ($x = 0.5, 1, 2, 3$) is x% below the opening price, rounded to the nearest one-eighth.

b A non-parametric test of the null hypothesis that the median return (of the ten subperiod means) equals zero.

* (**) Significantly different from zero at the 5% (1%) level. The sub-period tests use individual subperiod variance. The overall test uses the variance of all observations pooled across subperiods and stocks.

Table 5

Unconditional limit order returns compared to the unconditional market order returns, for the Dow-Jones Industrial stocks for ten sub-periods during 1988.

Differential return to all limit orders $r_o^l - r_o^m$		Limit Order Test[a]			
		0.5%	1%	2%	3%
Subperiod	1	0.311	0.418	-0.040	-1.574*
	2	0.315	0.390	0.962**	2.003**
	3	0.059	0.028	-0.574	-0.910
	4	-0.259	-0.240	1.006**	-0.278
	5	0.604**	0.553**	1.826**	0.569
	6	0.529**	0.793**	-0.247	1.272**
	7	0.064	0.018	0.163	-0.714*
	8	0.013	0.032	-0.261	-0.382
	9	-0.043	0.021	0.668**	-0.706
	10	0.113	0.191	0.220	-0.077
Overall 1988		0.171**	0.220**	0.355**	-0.080
Wilcoxon p-value[c]		0.012	0.007	0.069	0.351

[a] The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order ($x = 0.5, 1, 2, 3$) is x% below the opening price, rounded to the nearest one-eighth.

[b] A non-parametric test of the null hypothesis that the median return (of the ten subperiod means) equals zero.

* (**) Significantly different from zero at the 5% (1%) level. The sub-period tests use individual subperiod variance. The overall test uses the variance of all observations pooled across subperiods and stocks.

Table 6

Overall and subperiod results from the panel regression:

$$\eta_{it} = \gamma_1 + \gamma_2 L_{it} + \varepsilon_{it}$$

where i represents the stock, t represents the window, η_{it} is the market-adjusted differential return, and L_{it} is an indicator variable that equals 1 when the limit price is reached, and 0 otherwise.[a]

		Limit Order Test							
		0.5%		1%		2%		3%	
		γ_1	γ_2	γ_1	γ_2	γ_1	γ_2	γ_1	γ_2
Subperiod	1	-0.808	1.017*	-0.448	0.663	-0.429	0.851	-1.028**	2.201**
	2	-0.116	0.174	-0.147	0.312	-0.061	0.453	-0.431	2.103**
	3	-0.602	0.916*	-0.399	0.817*	-0.064	0.020	-0.392	1.666**
	4	-0.093	0.106	-0.362	0.822	-0.470	1.177**	-0.359	1.747**
	5	-0.763**	1.131**	-0.528**	1.252**	-0.065	0.629	-0.415	2.383**
	6	-0.546	0.696	0.027	0.011	-0.219	0.814	-0.043	0.437
	7	-0.535	0.755*	-0.155	0.337	0.004	-0.046	-0.205	0.708
	8	-0.527*	0.896**	-0.260	0.745*	-0.221	1.678**	-0.337	2.218**
	9	0.216	-0.371	0.082	-0.242	-0.305	1.509**	-0.225	1.362
	10	-0.088	0.152	-0.028	0.084	-0.302	1.406**	-0.129	0.728
Overall 1988		-0.341**	0.505**	-0.213**	0.465**	-0.209**	0.718**	-0.339**	1.494**

[a] The vector of market-adjusted differential returns (η_{it}) for the entire year was obtained from a first-pass market model regression given by equation (6) in the text.

* Significantly different from 0 at the 5% level.

** Significantly different from 0 at the 1% level.

Table 7

Robustness tests of limit order trading with the price record run backward (Inverted Buy) and with a limit sell order strategy (Regular Short-Sell), for the Dow-Jones Industrial stocks during 1988.[a]

Robustness Test[b]	Limit Order Test Parameter[c]	Differential return to executed limit orders $r_e^l - r_o^m$	Differential return to unexecuted limit $r_n^l - r_o^m$	Differential return to all limit orders $r_o^l - r_o^m$
Regular Buy	1%	0.561*	-0.067	0.220**
	2%	0.897**	0.132	0.355**
	3%	1.786**	-0.626**	-0.080
Inverted Buy	1%	0.363**	-0.040	0.154*
	2%	1.132**	-0.006	0.447**
	3%	0.967**	-0.680**	-0.282*
Regular Short-Sell	1%	-0.259	-0.328**	-0.301**
	2%	-0.244	-0.566**	-0.470**
	3%	0.623**	0.054	0.199

a The robustness tests are designed to examine the sensitivity of the results to trending in prices. The test period 1988 was marked by rising prices and it is not clear what role this plays in our results on the profitability of limit order trading. The robustness tests reverse the price trend (by either running the price series backward or by selling instead of buying) without altering the correlation structure of the prices.

b The robustness test compares a *regular buy strategy* with two alternative strategies, (1) an *inverted buy strategy* where the entire 1988 transaction price series for each stock is run backward and buy limit orders compared with buy market orders, and (2) a regular short-sell strategy where the price series is not disturbed but instead limit sell orders are compared to market sell orders, with the short position being covered at the end of the investment window.

c The limit order test parameter sets the limit order price with respect to the opening price (which is also the price at which the market order executes). Specifically, the limit price for the x% limit order ($x = 1, 2, 3$) is x% below the opening price, rounded to the nearest one-eighth.

* Significantly different from 0 at the 5% level.

** Significantly different from 0 at the 1% level.

Table 8

Limit Order Spread from a multiple limit order trading strategy when a network of limit orders is placed around the opening price, from one to five percentage points apart, and all transactions are closed at the end of an eighty trading-day period, for three trading periods in 1988 for all Dow Jones Industrial stocks.[a]

Percentage Parameter[b]	Round-trip gain[c] (*RTGAIN*) (in $)	Number of round trips[c] (*RTRIPS*)	Differential per share at closure[c] (*DC*) (in $)	Share imbalance at closure[c] (*IMB*)	Bagging Cost[c] (*BC*) (in $)	Non-execution Cost[c] (*NEC*) (in $)	Total Return[c] (Π) (in $)	Limit Order Spread[c] (*LOS*) (in $)
1% *($0.50)*	0.54 *(11.28)*	83.18 *(13.52)*	-0.31 *(-1.13)*	4.50 *(4.31)*	-36.86 *(-11.73)*	20.53 *(4.93)*	16.33 *(3.37)*	0.23 *(3.74)*
2% *($1.00)*	1.02 *(10.79)*	25.27 *(15.40)*	-0.31 *(-1.12)*	2.49 *(4.12)*	-19.40 *(-11.10)*	12.41 *(3.08)*	6.99 *(1.42)*	0.47 *(4.03)*
3% *($1.50)*	1.49 *(10.61)*	12.28 *(10.48)*	-0.24 *(-0.88)*	1.72 *(4.00)*	-12.89 *(-11.55)*	7.84 *(2.91)*	5.05 *(1.66)*	0.67 *(4.51)*
4% *($2.00)*	1.83 *(9.89)*	6.83 *(8.78)*	-0.31 *(-1.21)*	1.25 *(3.39)*	-10.03 *(-11.09)*	5.37 *(2.87)*	4.67 *(2.12)*	0.95 *(5.05)*
5% *($2.375)*	2.26 *(11.03)*	4.05 *(11.91)*	-0.12 *(-0.46)*	0.78 *(3.18)*	-7.49 *(-11.00)*	3.68 *(2.26)*	3.80 *(2.14)*	1.24 *(5.43)*

[a] We test the profitability of a *multiple limit order trading strategy* by monitoring purchases and sales during the period and also the profitability of the closing trade. Our hypothetical limit order trader is expected to gain from round trips very much like a market maker gains from the spread, and to lose at the closing trade when the inventory imbalance is corrected. For a definition of each quantity that was monitored, please see the text.

[b] Figures in brackets refer to the average rounded(to an eighth) dollar value of the percentage parameter.

[c] The t-statistic corresponding to the null hypothesis that the value is zero, is reported in parenthesis.