Limit Orders and Volatility in a Hybrid Market: The Island ECN

Joel Hasbrouck and Gideon Saar

November 12, 2001

Preliminary draft
Not for attribution
Comments welcome

Professor Joel Hasbrouck
Department of Finance
Stern School, NYU
44 West 4th St.
New York, NY 10012

tel. (212) 998-0310
e-mail: jhasbrou@stern.nyu.edu

Professor Gideon Saar
Department of Finance
Stern School, NYU
44 West 4th St.
New York, NY 10012

tel. (212) 998-0318
e-mail: gsaar@stern.nyu.edu

We thank the Island ECN for providing us with the data for this study, and are especially grateful for the help of Cameron Smith, Josh Levine, and Rob Newhouse. We also thank Tim McCormick and the NASD's Economic Research for providing us with odd-lot data, and Lei Yu for her research assistance.

All errors are our own responsibility.

Limit Orders and Volatility in a Hybrid Market: The Island ECN

Abstract

This paper presents an empirical analysis of trading activity on the Island ECN, an Alternative Trading System for US equities that is organized as an electronic limit order book. We focus on a cross-sectional investigation of the relations between different forms of volatility and various Island trading measures. We find that higher volatility is associated with

- a lower proportion of limit orders in the incoming order flow,
- a higher probability of limit order execution,
- shorter expected time to execution, and
- lower depth in the book.

In addition, we find that Island's market share for a given firm is positively related to the overall level of Nasdaq trading in the firm, and document substantial use of hidden limit orders (for which the submitter has opted to forgo display of the order). Finally, over one quarter of the limit orders submitted to Island are canceled (unexecuted) within two seconds or less. The extensive use of these "fleeting" orders is at odds with the view that limit order traders (like dealers) are patient providers of liquidity.

1. Introduction

The electronic limit order book has emerged as the most common form of security market organization worldwide. By choosing a market or limit order and selecting a limit price, an agent in such a market enjoys access to a range of strategies that trade off execution certainty against expected execution price. When the market has many participants, the collection of unexecuted limit orders (the book) may constitute a continuous source of liquidity, diminishing the role of professional intermediaries and maximizing direct interaction of the market's users. The factors that influence an individual's order choice and the aggregate properties of the limit order book are therefore of great interest.

For a market organized as an electronic limit order book, the volatility of the traded security stands out as one of the most potentially important factors. To illuminate the connections between volatility and book behavior, the present paper undertakes a cross-sectional empirical analysis of the trading process on the Island ECN, an electronic limit order book for U. S. equities.

The effects of volatility are complex and varied. For an individual trader contemplating the submission of a limit order, higher volatility will generally increase the probability of execution (a benefit), but may also increase the expected cost of getting picked off by a counterparty with better information and the expected cost of chasing a price that has moved away (see Angel (1994) and Harris (1998)). The source or type of volatility may matter. While execution probability may simply depend on overall volatility, the risk of getting picked off may be more properly measured by the trade-driven component of volatility. In moving from the individual choice problem to characterization of the equilibrium, additional interactions arise. For example, increased use of limit orders implies fewer market orders, and a consequent drop in the expected execution rate of the submitted limit orders, decreasing their desirability (Foucault (1999)).

It is also reasonable to conjecture causal effects running from characteristics of the trading mechanism to volatility. For example, holding constant the incoming flow of market orders, a deep book, i.e., one with large quantities at successive price levels, will exhibit smaller price changes than a thin book (Domowitz and Wang (1994)). This motivates a distinction between “true value” volatility (which derives from an asset’s fundamentals) and market price volatility (which impounds effects of the trading mechanism).

We use a cross-sectional approach to analyze the effects of true value and market price volatility in a typical limit order market. The Island ECN is one component of the Nasdaq stock market.[1] The market share of Island for the average stock in our sample is 3.5% in terms of share volume and 6.2% in terms of trades. [2] Two ECNs (Island and Instinet) dominate Nasdaq non-dealer trading. Whereas Instinet is predominantly used by institutions, however, Island is heavily used by retail investors.

That Island’s market share falls far short of dominating overall Nasdaq activity affects our analysis in several ways. Some of our conjectures derive from models in which the electronic book constitutes the entire market. In applying these models to a hybrid market, we are extending the implications of these models beyond their original formal scope. On the other hand, Island’s small market share helps justify the assumption that overall price determination and overall Nasdaq trading activity are exogenous with respect to measures of Island activity. It therefore lessens concerns about reverse causality in our econometric specifications.

For each stock in our sample, we examine the trading activity of investors on Island using three types of measures. First, we look at the manifestations of the trading

[1] In US securities law, an ECN (Electronic Communications Network) is a medium for disseminating (“publishing”) quotes (U.S. Securities and Exchange Commission (1996)). Because it offers executions, Island is also classified as an Alternative Trading System (ATS, US Securities and Exchange Commission (1998)).

[2] More recently, Island claims an overall Nasdaq share of 16.2% by trades and 7.6% by volume for June 2001 (Island (2001)).

strategies of investors by investigating measures of the flow of orders through the market: the proportion of limit orders in the order flow (the reminder are market orders), and the proportion of limit orders that are filled (the remainder are canceled or expire). Second, we examine average depths in the book that can be viewed as Island's supply and demand curves. These quantities are, of course, simply the accumulations of all limit order submissions, executions and cancellations at each price, but they conveniently summarize the steady state of the system. Third, we conduct a duration analysis of times to fill or cancel for limit orders. This characterizes the speed in real (wall-clock) time of the market's dynamics.

We then investigate the relation between Island's measures and the various measures of volatility that have arisen in the discussion above: total, permanent, systematic, unsystematic and trade-driven. We generally find that across all measures, higher volatility is associated with:

- a lower proportion of limit orders in the incoming order flow
- a higher probability of limit order execution
- shorter expected time to execution
- lower depth in the book.

These results confirm many of the direct effects present in the economic models of individual order strategy. They suggest that the offsetting effects hypothesized to arise in equilibrium are not large.

In addition, we examine how characteristics of the investor clientele, the percentage of shares held by institutions and the activity of odd-lot traders, relate to the Island measures. We find that that institutional ownership is positively related to the limit order submission proportion, and the proportion of limit orders priced at the quote or better, but negatively related to the proportion of limit orders priced behind (away from) the quote. This suggests that while institutions may favor limit orders, they do not contribute to depth away from the market.

Like many other electronic markets, Island permits undisclosed (also called hidden, invisible or "iceberg") orders. Our data support a partial characterization of such orders. We also examine an issue that is of particular interest to understanding the electronic limit order book, the phenomenon of "fleeting" limit orders, i.e., orders that are canceled almost immediately after submission. These constitute a substantial portion of the order flow in many stocks and their characterization is important to understanding this market structure. Finally, we look at Island's market share and how it relates to the volatility and investor measures, as well as to the presence of Island at the inside quote on Nasdaq.

The paper is organized as follows. In the next section, we review prior studies. Section 3 describes the Island system. Sample construction and data sources are discussed in Section 4. Section 5 discusses the econometric specifications. Results are presented in Section 6. Section 7 documents the importance on the Island system of hidden and "fleeting" limit orders, i.e., orders canceled almost immediately after submission. Section 8 provides an analysis of Island's market share. A summary concludes the paper in Section 9.

2. Literature survey

The large and growing importance of electronic limit order book systems in many securities markets has engendered much interest. In the following survey, we concentrate initially on the theoretical literature, with a view toward establishing relevant empirical predictions. We then highlight prior empirical work related to our analysis. Our discussion focuses on "volatility," but we note at the outset that meaning and usage of the term varies considerably across this literature.

a. Theoretical models

From one perspective, an agent who places a limit order is acting as a dealer. Models of dealer behavior are therefore useful, especially those featuring asymmetric information (beginning with Copeland and Galai (1983), Glosten and Milgrom (1985) and Easley and O'Hara (1987), and extending to many others). Building on this view, Glosten (1994) models a market organized as a consolidated limit order book. Seppi

(1997) and Rock (1990) consider the interplay between the book and a strategic dealer (specialist) who possesses a last mover advantage. Parlour and Seppi (2001) examine competition between one trading venue organized as a pure electronic book and one constituted as a book/dealer hybrid.

These models generally maintain that the information content of the incoming order flow is an important determinant of how aggressively a dealer will quote or how much liquidity will be supplied in the book. The analyses imply that an order flow characterized by a relatively high proportion of privately-informed traders will lead to a wider spread and smaller quantities available for sale or purchase at each price in the book. The present paper investigates this relation.

These models generally feature a sharp distinction between liquidity suppliers (limit order traders and dealers) and liquidity demanders (market order traders). The book fills to the point where no agent wishes to submit an additional limit order or to cancel an existing one. The dynamics of limit order arrivals are not explicitly modeled. The filling is assumed to occur instantaneously, or at least prior to the time at which a market order can arrive. For this reason, these models are sometimes described as static.

In reality, of course, an agent's role is often a matter of choice. An arriving trader makes the decision of whether and when to be a liquidity supplier (with a limit order) or a demander (with a market order). The practical importance of this decision motivates analysis of an agent's optimal order strategy.

In Cohen et al. (1981) agents are characterized by endowments and subjective assessments of the security's value and make choices between submitting market and limit orders. The model demonstrates the tradeoff between execution uncertainty and the cost of trading. Angel (1994) and Harris (1998) provide partial equilibrium models that analyze the optimal choice between market and limit orders and relate it to various factors in the economic environment. Some empirical implications of these models are intuitively clear: limit orders are more likely to be submitted if the spread is large (i.e., that the contra-side quote is distant) and the arrival rate of market orders is high.

The role of volatility in these models is more complex. Higher volatility generally increases the probability that a limit order will execute. This may be illustrated with the simple model suggested by Lo, MacKinlay, and Zhang (1997). A limit order is placed away from the current market price. It fills when the limit price is first met. This is formally the first passage time to the (limit price) barrier. If the security price follows a diffusion process with zero drift and variance per unit time σ^2, then the probability of an execution in any finite time increases with σ and converges to one (certainty) as time approaches infinity. Although this simple view suggests the direction of the effect, Lo, MacKinlay, and Zhang (1997) note that it leads to poor predictions of actual execution times.

The option perspective on limit orders, on the other hand, suggests that volatility is costly. The limit order is more likely to be picked off when the "true value" of the security has changed in such a way as to make it mispriced (say the value of the security increased above the price of a limit sell order). This pick-off risk is similar to the private information risk faced by a dealer in the sequential trade models. In a limit order context, however, the exposure can arise from information that might otherwise be considered public (e.g., a news release). The risk arises to the extent that the limit order can't be revised or canceled before the arrival of a market order.

Finally, higher volatility increases the expected distance that the price can move away from the order, decreasing its chances of ever executing. An agent who needs to complete the trade must chase the price, and perhaps ultimately enter a market order.

In Angel (1994)'s model, a buyer attempts to minimize the expected purchase price subject to the requirement of trade completion within a given time. There is no penalty for being picked off. With this consideration eliminated, Angel's numerical simulations suggest that the beneficial effect of volatility on execution probability outweighs the expected cost of chasing a price that has moved away. In all, volatility encourages use of limit orders. On the other hand, the objective function in Harris (1998)'s model does penalize bad fills. Here, the option effect dominates, implying that higher volatility makes limit orders less attractive.

Harris shows that a strategic trader who derives private value from a completed trade will use more aggressive limit orders than a risk-neutral liquidity trader, i.e., an agent who is closer to a dealer in the sequential trade models. The strategic traders, therefore, may determine the marginal prices and quantities in the book. This behavior also arises in Chakravarty and Holden (1995), where informed traders use limit orders to undercut the dealer's quotes.

Dynamic equilibrium is the focus in Parlour (1998), Foucault (1999) and Foucault, Kadan, and Kandel (2001). In these models, agents arrive sequentially and choose between market and limit orders. Their optimal strategies are conditioned on conjectures of other traders' optimal strategies. To make the analyses tractable, traders' problems and strategies are constrained. While traders have diverse private values, for example, none possess private information. More importantly, however, the agents are limited to one action in the market. They can't revise or cancel their orders and the time-in-force for a limit order is either one period or infinite. This rules out many of the dynamic strategies available to agents in Angel (1994) or Harris (1998).

The empirical predictions of the dynamic models typically involve transition probabilities or relative occurrences of events. Foucault (1999) provides the most definite empirical predictions concerning volatility. In his model the relevant volatility measure is that of fundamental ("true") value, rather than market price volatility. An agent's order choice is most immediately affected by the pick-off risk. An increase in volatility causes a trader who uses a limit order to price the order less aggressively. This causes an increase in the spread. The cost of using a market order rises, favoring a shift to limit orders. Thus, volatility should increase the proportion of limit orders, and decrease the fill (execution) rate of these orders.

Domowitz and Wang (1994) examine the behavior of a limit order market in which orders arrive at various price levels with Poisson intensities that are partially endogenous. A general decrease in the arrival rates causes a drop in book depth, an increase in volatility of quoted and transaction prices, and an increase in the rate of market failure (intervals where the book is empty). Due to this mechanism, the model implies an association between volatility and properties of the book. The book in this model, however, generally achieves a stationary limiting distribution, which is incompatible with a diffusion process for the fundamental asset value. Accordingly, volatility in this framework derives solely from disturbances that are transitory (such as bid-ask bounce).

b. Empirical studies

Economic logic suggests that since limit orders forgo immediate execution, they should realize a cost advantage (on average) relative to market orders. Harris and Hasbrouck (1996) find this to be the case in a sample of NYSE orders. Investigating simulated strategies imposed on actual data, Handa and Schwartz (1996) find that when the costs of nonexecution are ignored (an assumption applicable to patient traders), the returns to limit orders are positive.

Harris (1998) notes that limit orders are priced more aggressively by liquidity traders nearing their trade completion deadline and by informed traders with stronger signals. By implication, both agents would be more aggressive than a dealer who is uninformed and indifferent to trade completion. [3] Chung, Van Ness, and Van Ness (1999) find that bid and ask quotes on the NYSE frequently represent the book instead of the specialist. Harris and Panchapagesan (1999) conclude that the state of the book is informative, in the sense of predicting future short-term (though not long-term) price movements.

A number of studies examine various features of markets organized primarily as consolidated limit order books. Among these studies, Sandas (1999) constitutes the sole empirical characterization of a limit order book in the equilibrium characterized by Glosten (1994). The results suggest that the book on the Swedish Stock Exchange

[3] A dealer nearing the close of a trading session with an inventory imbalance, however, may more closely resemble a liquidity trader subject to a completion deadline. Consistent with this supposition, Chan, Christie, and Schultz (1995) find that Nasdaq quotes narrow toward the end of the day.

provides less liquidity than would be predicted on the basis of the information in the order flow. For incoming buy orders, for example, the supply curve is too steep relative to the price revisions that these orders ultimately cause.

Other studies characterize the incoming order mix. A positive relation between the prevailing spread and the probability that an incoming order is a limit order is found on the Paris Bourse (Biais, Hillion, and Spatt (1995)), the Toronto Stock Exchange (Griffiths et al. (2000)) and for an anonymous Nasdaq wholesaler (Smith (2000)). These findings are consistent with the predictions of the order strategy models.

The evidence on volatility is mixed. Ahn, Bae, and Chan (2001) examine transitory volatility on the Hong Kong Stock Exchange. They measure volatility over 15-minute intervals and find that depth on the book rises subsequent to a volatility shock. On the other hand, Coppejans, Domowitz, and Madhavan (2001) find that a volatility shock reduces depth. Goldstein and Kavajecz (2000) note that during the October 1997 break, book depth declined dramatically.

Although some of these studies employ cross-sectional variables, they generally investigate variation over time. [4] The present paper employs a cross-sectional (across firm) perspective, which is particularly appropriate for investigating the attributes of firms that are related to the behavior of the book. Some of these attributes, and volatility is certainly one of them, also exhibit time-series variation; but even in these cases, it cannot be presumed that the cross-sectional and dynamic relations are the same. In fact, it is much more reasonable to examine the implications on the book of "true value" volatility using a cross-sectional rather than a time-series specification.

Time, in the formal analyses of limit orders, is primarily a notional construct. It typically indexes the sequence of agents' moves, rather than the passage of real ("wall

[4] Griffiths et al. (2000) examine firm size as a determinant of order aggressiveness. In similar specifications, Smith (2000) includes price and volatility.

clock") time. [5] It is nevertheless clear that in actual trading situations, real time may play a more distinctive role, due to institutional features (such as regular trading hours), decision cycles or monitoring costs that are measured in clock time. In limit order analyses, real time effects have been studied using duration models to characterize the time-to-fill or time-to-cancel of an order (Lo, MacKinlay, and Zhang (1997) and Cho and Nelling (2000)). If time were important solely as a volatility scale factor, a duration model would have a simple form, an accelerated failure time representation with volatility (per unit time) as the only important determinant. In fact, numerous other firm and market variables contribute explanatory power, suggesting a more complex relationship. For these reasons, the relation between duration and volatility is investigated in the present study.

Other ECN-related papers examine quotes or trades, but do not have the order-level data that we use in this paper: Huang (2000) investigates the contribution of ECN's to price discovery for the ten most actively traded NASDAQ stocks; Simaan, Weaver, and Whitcomb (1998) examine the behavior of market makers and ECN's following the tick size change to sixteenths; Conrad, Johnson, and Wahal (2001) examine institutional trading costs on ECN's and crossing systems; and Barclay, Hendershott, and McCormick (2001) compare execution costs between ECN's and market makers.

3. The Island ECN: Background and trading protocol

Since this is, to the best of our knowledge, the first academic study of the Island ECN, it is worthwhile to document the system and data. The Island ECN was founded in 1996 and began operating on January 1997, becoming one of the two largest ECN's in the market today in terms of both share volume and number of traders (the other major ECN is Instinet). In terms of market share, about 11% of the trades in Nasdaq stocks were executed on Island during our sample period (the last quarter of 1999), representing close

[5] The distinction between real time and event or "informational" time is a recurring theme in studies of financial markets (see Clark (1973) and Russell and Engle (1998), for example).

to 6% of Nasdaq's volume. The disparity between the market share in terms of trades and share volume testifies to the small size of most Island trades. In addition, Island's market share is not the same for all stocks, and seems to be higher for a small number of very active stocks. The market share of the average stock in our sample (that is comprised of the top 300 Nasdaq firms by market capitalization) is 6.23% in terms of trades and 3.52% in terms of share volume.

Island operates a pure agency market. The system is active (i.e., orders can be submitted and trades can take place) from 7:00 in the morning to 8:00 in the evening. [6] Island accepts only priced limit orders. Market orders as such are not accepted. A trader who seeks immediate execution must submit an order at a limit price that meets or crosses the best opposing price (a marketable limit order). Each time a limit order is received and the book contains a matching order, the limit order is immediately executed. If there is no matching order, the limit order is placed in the book until a matching order is received or the limit order is canceled.

All orders are matched based on strict price-time priority without regard to the number of shares in the order. The Island display is anonymous—the identities of the investor or the broker are not visible—with only the price and the number of shares made available to the market. Island's top orders are also represented in the Nasdaq quote montage, and are therefore incorporated into the National Best Bid/Offer (NBBO) display.

An unmatched limit order is normally displayed. That is, the price and size of the order are visible to Island subscribers. At the trader's discretion, however, display may be limited to Island subscribers or suppressed entirely. In neither of these cases is the order incorporated into the Nasdaq montage or NBBO.

Since Nasdaq forbids locking or crossing their market, subscriber-only orders are a convenient way of attempting to buy or sell a stock outside the Nasdaq quote without

[6] During our sample period, Island offered a continuous session from 8:00 a.m. to 8:00 p.m.

violating Nasdaq rules. [7] The display requirements of SEC's Regulation ATS dictate that if an ECN executes more than five percent of the total volume in a given stock during four out of the last six months, then the ECN is large enough that it should be required to display all its visible orders to the public marketplace. Island does not accept subscriber-only orders in the list of stocks that are subject to the ATS display requirements. This regulation does not apply to invisible orders because they are not seen on the Island book.

A subscriber can also specify the minimum number of shares of an order that can be executed. This feature is primarily aimed at subscribers who do not want to get odd-lot executions. However, orders that specify a minimum number of shares that is higher than 100 are not reflected in Island's quote on Nasdaq. An order that either specifies a minimum number of shares or is invisible has a lower priority than an order that is not restricted in these two ways. The lower priority means that if an order with a restriction is entered before an unrestricted order at the same price, the unrestricted order will execute first (i.e., restricted orders lose time priority).

An Island subscriber can submit limit orders without charge. If a limit order sits in the book and subsequently is executed by an incoming order, it is considered to have added liquidity to the book, and the subscriber receives a 0.1 cent rebate per executed share. The incoming order that removed liquidity from the Island book is charged 0.25 cent per executed share. While Island subscribers pay a fee for getting a data feed that allows complete construction of the book in real time, anyone with an Internet browser can observe the top 15 orders on each side of the book (for any stock) in real time on Island's web site.

---

[7] Island operates solely as an agency market that automatically executes matching buy and sell interest, irrespective of quotes displayed by other market participants. Hence, routing an order to Island does not guarantee receiving the best price in the market. Island maintains that it is the subscriber's responsibility to ensure best execution for their transactions by selecting the appropriate market venue. Also, subscribers bare sole responsibility to complying with Nasdaq's short sale rule, as Island does not check orders or executions to ensure compliance with the rule. The Island system is programmed to comply with the SEC short sale rule for NYSE-listed securities.

4. Sample and data

a. Sample construction and descriptive statistics

The sample was drawn from all Nasdaq National Market common stocks with data in the CRSP database from October 1 to December 31, 1999. [8] The sample is the 300 largest firms based on equity market capitalization as of September 30, 1999. [9]

Table 1 presents summary statistics. The smallest firm has an average market capitalization over the sample period of 824 million dollars, while the median firm is just over 3 billion dollars and the largest firm is close to 495 billion dollars. The sample also spans a range of trading activity and price levels. The most active firm has a daily average of 28,654 trades, while the median firm has about 1,066 trades on an average day, and the least actively traded firm in the sample has (on average) only 16 trades per day. Average daily CRSP closing prices range from \$8.40 to \$326.58, with a median of \$45.66. To provide a sense of the cross-sectional characteristics of the variables, we report means for subsamples constructed by ranking on market capitalization, average number of daily trades and standard deviation of daily returns, σ_r.

To characterize the trading clientele for the firms, we examine the level of institutional holdings from the Value Line Investment Survey. This ranges from 6.02% to 97.93%, with a median (and mean) of 51%. As a proxy for overall level of trading activity, we use daily turnover (the ratio of the day's volume to the number of shares outstanding), and take its median over our sample. To measure retail investor activity we

[8] The Nasdaq Stock Market is comprised of two separate market categories—Nasdaq National Market (NNM) and Nasdaq SmallCap Market (SCM). The two market categories differ mainly with respect to the listing requirements (but also with respect to a few details of trading protocol). The NNM has stricter listing requirements and generally includes larger firms.

[9] We also required that firms do not have more than one series of common stocks traded. Two firms (Associated Group Inc. and Molex Inc.) were excluded from the sample on this basis. We also excluded Comair Holdings Inc., which was in the process of being acquired by Delta Air Lines during the sample period.

use the daily average number of odd-lot trades provided to us by the NASD's Economic Research. The range of this quantity is from 19.81 odd-lot trades per day up to 8,348.22, with a median of just over 125. [10]

b. Island data and statistics

The data supplied by Island are identical to those supplied in real time to Island subscribers. These data comprise time-sequenced messages that completely describe the history of trade and book activity. The process may be summarized as follows. When an arriving order can be matched (in whole or part) against an existing order in the book, the system sends an Order Execution message. If all or part of the order can't be matched, the system sends an Add Order [to the book] message. An Add Order message contains the direction (buy or sell), number of shares, limit price, a display condition (normal or subscriber-only), and a unique identification number. If and when the order is executed, this number is reported in the Order Execution message. When an existing order is canceled or modified (in size), the system generates a Cancel Order message. The book, excepting the invisible orders, may be constructed by cumulating these messages from the start of the day onwards. Although the arrival time and quantity of an invisible order are never made available, the execution of an invisible order is signaled by a special trade message. In the rare event that a previous trade report was in error, the system sends a Broken Trade message.

Table 2 presents summary statistics on the number and sizes of orders that arrive to Island. We only consider data from the regular trading session of the Nasdaq Stock Market (from 9:30 a. m. to 4:00 p. m.). This was done to ensure that we are indeed looking at the Island system only when it is part of a much larger market and captures a relatively small fraction of the order flow and not when it is the one of a handful of

[10] The reason that the minimum AvgTrd from CRSP can be 15.73 while the minimum OddTrd can be 19.81 is that odd-lot trades are not reported to the tape, and are therefore not counted in the CRSP number (nor do they appear for that matter in the TAQ or NASTRAQ databases).

venues for trading during pre-opening and after-market hours. The average number of daily limit orders increases with market capitalization (in the ranked group means), average daily trades, and σ_r. The average size of limit orders on Island is 572 shares, testifying to the retail nature of trading on the system. The average size decreases slightly across capitalization and average trade subsamples, which may suggest that retail activity is more concentrated in the largest, most active Nasdaq stocks.

The Island system does not accept unpriced market orders. Traders seeking immediate execution must enter a limit price at or better than the opposing quote (e.g., a limit order to buy at the ask). Although these orders are formally marketable limit orders, we will henceforth consider them to be market orders. Table 2 shows that market orders tend to be smaller than limit orders, with a mean of only 335 shares. As with limit orders, the average size decreases with market capitalization and trading activity.

Nasdaq and Island trading activity is illustrated in Figure 1. For both Nasdaq and Island, activity is concentrated in the higher market capitalization deciles. Figure 2 describes Island's orders across market capitalization deciles. Limit orders outnumber market orders. Most limit orders are priced away from (less aggressively than) Island's quote.

c. Constructed Island variables

The observable variables that are closest to their counterparts in the theoretical models are the number of limit orders submitted (as a proportion of the total, limit and market orders), and their execution proportions. It is also interesting and useful to characterize the aggressiveness of the limit orders. Accordingly, we examine the number of limit orders priced at Island's quote or better, e.g., buy orders priced at Island's bid or better, and those less aggressive orders priced behind Island's quote. We also compute similar statistics based on the number of shares in the orders.

Table 3 presents summary statistics on the submission proportions. First note that most of the order submissions are limit orders: a median of 82% (by number of orders), and 89% (by number of shares). In the ranked subsample means these proportions decrease with capitalization, average number of trades, and σ_r. This behavior also characterizes the more aggressive limit orders (priced at or better than Island's quote). The reverse is true, however, for the less aggressive limit orders. That is, as capitalization, average number of trades, and σ_r increase, traders tend to submit less aggressive limit orders. Proportions defined in terms of shares behave in a similar fashion.

Table 3 also presents summary statistics for execution proportions. [11] The mean execution rate is 18% (by orders) and 13% (by shares). In the subgroup means, the execution rates increase with market capitalization, average number of trades, and σ_r. Surprisingly, execution rates for more aggressive orders (those priced at the quote or better) are generally *lower* than the execution rates for less aggressive orders (behind the quote). There are a number of considerations that could potentially account for this, notably strategic order management. In particular, many of the more aggressive orders are canceled after one or two seconds, thus depriving them of the chance for execution. We discuss this behavior more extensively in Section 7.

The second type of analysis we provide is that of depth in the book. We construct a snapshot of the ask side of the book at five minute intervals by recording the cumulative dollar depth at $1/16 intervals, ranging from $0.50 better than the National Best Offer (NBO) up to $5.00 worse than the NBO (88 groups, with the last group including "$5 or higher"). We construct similar snapshots on the bid side. As an example, Figure 3 presents a box-plot summary of the ask depth snapshots for Microsoft. The length of the boxes and "whiskers" for each interval is large, which indicates substantial variation in depth over time. The mean ask-side depth for Microsoft is $4,416,000; the standard deviation is $2,130,000, implying a coefficient of variation slightly above two.

[11] Note that execution proportions cannot simply be defined as the number of market orders over the number of limit orders due to the difference in size between market and limit orders (limit orders are larger on average than market orders). We are able to follow each limit order that enters the system and therefore can produce an exact characterization of the execution proportion of orders.

Our cross-sectional analysis requires a summary measure of book depth for each firm. The most obvious measure is simply the mean (over time) of the total dollar value of the orders on the ask and bid sides. We computed this, but also estimated quadratic polynomial functions. The final results were substantially similar whether the summary measure was total dollar value, the intercept in the quadratic fit, or the slope coefficient. (The quadratic coefficient was small and variable.) Similar results were also obtained when the depth snapshots were measured in shares, and also when the interval widths were specified as a fraction of the average stock price. Accordingly, we report results only for total dollar value of orders.

Table 4 presents summary statistics (across firms) for mean total dollar depth. The mean coefficient of variation is roughly one, suggesting substantial time variation in depth for most firms. The standard deviation (across firms) of mean dollar depth is also relatively high, which indicates variation across firms. The patterns across market capitalization-ranked groups and trade-ranked groups conform to expectations. Larger and more actively traded firms have deeper books. Within the standard deviation-ranked groups, however, depth is not monotonic.

The third Island characteristic investigated in this paper is the timing of order events, and in particular the (elapsed) time between an order's submission and its first fill. We will refer to this duration simply as time to execution. In contrast to the proportions and depths discussed above, there is no convenient firm-specific summary statistic for this variable. Since many limit orders are canceled or expire unexecuted, the average duration of executed orders is an optimistic (unrealistically low) indication of what the submitter of a limit order should expect. The analysis therefore relies directly on individual limit orders, without an intermediate aggregation at the firm level.

This reliance leads to two difficulties. First, the number of observations becomes unworkably large. To deal with this, we employed for each firm a maximum of 2,000 orders (uniformly drawn from the all of the firm's orders). The second difficulty is that a modest number of firms did not have 2,000 orders in the sample. For the duration models, each observation is weighted by the inverse of the number of observations for the firm. This procedure effectively weights firms equally.

Figure 4 depicts “failure” functions for executions and cancellations. Intuitively, the failure function is the cumulative probability of event occurrence. Both executions and cancellations are of interest, so our treatment of the two types of events is symmetric. Specifically, in applying the (standard) Kaplan-Meier correction, cancellation is treated as censoring in the execution estimation, and execution is treated as censoring in the cancellation estimation. These functions are not adjusted to correctly weight firms with fewer that 2,000 observations; such firms are effectively underweighted. The time scale is nonlinear (to show detail for shorter times).

The cumulative execution probability rises fairly slowly, reaching approximately 70% at two hours. The function is almost certainly biased upwards. The standard framework assumes that the censoring process is independent of the event process. In the present case, this is tantamount to assuming that a limit order that is canceled has the same probability of execution (going forward) as an order that isn’t canceled. It is violated, for example, if traders are more likely to cancel limit orders when the price has moved away after submission.

The cumulative cancellation probability exhibits two notable features. Most strikingly, a large number of limit orders are canceled very shortly after their submission. Roughly 25% have been canceled after two seconds, and about 40% after ten seconds. [12] This is inconsistent with the traditional view of a limit order as providing ongoing liquidity. We describe limit orders canceled shortly after execution as “fleeting”, and discuss them in Section 7. The second interesting feature is the existence of two relatively sharp jumps in the cancellation function, at exactly three and five minutes. The Island protocol allows traders to specify a time-in-force for the order. Apparently three and five minutes are frequent choices.

[12] Like the execution probability function, the cancellation function is biased upwards, but since price movements over ten seconds tend to be small, the bias at this end of the time scale is likely to be small.

d. Volatility measures

The literature surveyed above suggests a central role for volatility. Different models, however, use the term to characterize different concepts. We consequently employ multiple measures in the empirical analysis. Table 5 presents summary statistics for these measures.

The first volatility measure is simply the return standard deviation, introduced above as σ_r. A sensible refinement of this variable involves differentiation between systematic and unsystematic volatility. This distinction may be important for the usual reason (in many asset pricing models, only the systematic risk is priced). In the present situation, however, systematic volatility may also proxy for trading risk that is relatively easy to hedge. An indexed portfolio manager who needs to invest in stocks, for example, might initially enter into a long futures position, and then purchase the individual stocks over time (reducing the futures position commensurately).

Our measures of systematic and unsystematic risk are based on the market model:

$$r_{it} = \alpha_i + \beta_i r_{Mt} + e_{it}, \qquad (1)$$

where r_{Mt} is the CRSP value-weighted portfolio return. The specification is estimated using three prior years of daily data (from October 1, 1996 to September 30, 1999). Data limitations restricted these estimations to 211 firms. Our proxy for the systematic risk for firm *i* is $\beta_i \sigma_M$; unsystematic risk is $\sigma(e_{it})$.

The volatility measures discussed to this point are derived from transaction prices. They therefore impound trading-induced price movements, such as bid-ask bounce. Noting this, Foucault (1999) suggests that long-run volatility (estimated using the Hasbrouck (1991) procedure) is the preferred measure. From intraday TAQ data aggregated at a one-minute frequency, we estimate a vector autoregression (VAR). [13]

---

[13] The details of the procedures are as follows. All variables are one-minute time-aggregates: r_t is the change in the logarithm of the NBBO midpoint at the close of the minute; x_t is the sum of the trade volume, wherein each trade volume is signed by reference to the midpoint of the quote immediately preceding the trade; $Sign(x_t)$ is +1 if

The VAR estimates may be transformed to yield the variance of the random-walk component of the security price, σ_w^2. We use the standard deviation per minute, σ_w, scaled up by a factor of $\sqrt{6.5 \times 60}$ to reflect volatility over a 6.5 hour trading day.

Table 5 shows that the estimated mean of σ_w is lower than that of $\sigma_{r'}$ in the total sample and all subsample groupings. There are two likely explanations for this. First, $\sigma_{r'}$ includes the overnight period, while σ_w does not. In the present case, this is a point in favor of σ_w, since the limit order book is primarily active during regular trading hours. Second, σ_w has been purged of transient volatility.

The VAR also supports a decomposition of the random-walk variance: $\sigma_w^2 = \sigma_{w,r}^2 + \sigma_{w,x}^2$, where the two terms on the left derive respectively from return innovations and signed trade innovations. We employ $\sigma_{w,x}$, which reflects the contribution to permanent changes in the security price that can be attributed to new trade information (while $\sigma_{w,r}$ is the remaining portion due to return innovations). Table 5 reports summary statistics for these measures. On average, about one-quarter of the random-walk volatility is due to signed trades.

It was noted in the literature summary that volatility enters some limit order models as a determinant of the likelihood that an agent using a limit order will have to "chase" a moving market. For these agents, typified by Angel's buyer and Harris's motivated trader, volatility increases the aggressiveness of their order strategies. In this view, the source of volatility is not important, and a total volatility measure such as $\sigma_{r'}$ or σ_w is appropriate. A trade-driven volatility measure, however, may be a more appropriate measure of pick-off risk.

---

x_t>0; –1 if x_t<0; and, 0 if x_t=0. $x_t^{1/2}$ is the sum of the signed square-roots of the trade volumes. The VAR comprises the variables $\{r_t, Sign(x_t), x_t, x_t^{1/2}\}$, with first and second lags included. The model allows contemporaneous effects running from the trade variables to returns.

5. Specifications

Corresponding to the three sorts of Island variables (execution and submission proportions, depths and times to execution), this study estimates three types of cross-sectional specifications. Each specification features a linear regression in which the regressors are the firm-specific variables. This commonality facilitates the presentation and discussion of results. The actual statistical models and their underlying assumptions are varied.

The analyses of the submission proportions, execution proportions and depths are conducted in two steps. For a proportion, we regress the overall estimate for each firm against the firm-specific variables. To deal with the restricted range of the proportions (between zero and one, inclusive), the proportions were first transformed by the logit function $f(x) = \log\left[x/(1-x)\right]$ for $0<x<1$. (Observations for which x=0 or 1 were deleted.) For the depth analyses, we compute the mean total dollar depth for each firm, which is then used as a dependent variable in the regression specification. It is important to note that these procedures effectively weight all firms equally. All regressions are estimated using OLS with White's heteroskedasticity-consistent standard errors. [14]

The analyses of time to execution are conducted in one step, i.e., without first constructing summary measures at the firm level. The data here consist of individual orders (up to 2,000 for each firm). This motivates a duration analysis of the sort suggested by Lo, MacKinlay, and Zhang (1997). Specifically, we employ an accelerated failure time model, wherein the logarithm of the time to execution is modeled as a linear function of the explanatory variables. Limit orders that are canceled or expire correspond to censored observations. [15] To maintain equal-firm weighting, we use the following

---

[14] For regression specifications (except the duration models), we also used two-stage Least Trimmed Squares (see Rousseeuw and Leroy (1987)) to examine whether our results are affected by outliers. The results were almost identical to the OLS results, and are therefore omitted from the presentation in the paper.

[15] There are two important generic limitations of standard duration analysis in limit order applications. First, we view the observations (orders) as independent, whereas the dynamic trading strategy models suggest that they may be linked. More importantly,

procedure. Denote by n_i the number of observations (limit orders) for firm *i*. Since n_i<2,000 for some firms, we weight each observation by $1/n_i$.

In each estimation we include as regressors, in addition to the volatility measures, a standard set of three control variables: the log of the average market capitalization; the average price per share; and, the log of the median daily turnover. Among other things, capitalization controls for agents' general beliefs about holding and trading characteristics of the firm. The average price is included to pick up discreteness effects in the price grid. Median turnover is intended to control for the market-wide "normal" level of trading in the stock. The median is used instead of the mean in order to have a measure of the typical trading intensity in a stock that is less sensitive to information shocks.

Many of the variables we seek to model (e.g., limit order execution rates), as well as many of our explanatory variables (such as turnover) are derived from trading data over the same sample period. This raises the possibility of simultaneity (causal effects running from the modeled variable to the explanatory variables) or correlated measurement errors. Our modeled variables, however, are derived solely from Island data, while the explanatory variables are computed using all Nasdaq trading activity and Nasdaq-wide prices. Since Island accounts for a relatively small portion of overall Nasdaq activity, problems stemming from reverse causality or correlated measurement errors are likely to be small. We provide additional evidence on this point in Section 8.

6. Results

Table 6 reports estimations in which the Island variables are modeled as functions of the control variables and a total volatility measure. As discussed in the preceding section, σ_r is a conventional daily volatility measure, while σ_w measures the implicit random-walk volatility. Despite this difference, the estimated effects of both variables are similar in all specifications. Most importantly, higher volatility is associated with a lower

---

though, the framework assumes that the cancellation and execution processes are independent (conditional on the explanatory variables). This independence is violated by a strategy in which limit orders are canceled if the price has moved out of range.

overall limit order submission proportion, but with a higher rate of execution. These findings are most consistent with the (non-equilibrium) order strategy models. In these models higher volatility increases the option value (a cost) of the limit order (or alternatively the pick-off risk). This reduces the desirability of limit orders. The shift to market orders increases the execution proportion of the remaining limit orders (as predicted by Lo, MacKinlay, and Zhang (1997)). [16]

The equilibrium view holds that limit order traders protect themselves against the higher pick-off risk by pricing their orders less aggressively. The estimates for the limit order submission rates differentiated by price support are consistent with this hypothesis. Higher volatility is associated with a higher proportion of limit orders priced behind Island's quote. In the equilibrium view, however, this increases the cost of market orders, leading to fewer market orders and lower execution rates for limit orders. The results are not consistent with these predictions.

The next set of estimates in Table 6 deals with book depth. (For the sake of brevity, only ask-side estimates are presented. Bid-side results are similar.) Both measures of volatility are negatively associated with book depth, but the relation is not statistically significant. The associations between time to execution and the volatility measures are strongly negative: increased volatility is associated with more rapid execution. This is consistent with the analysis and results of Lo, MacKinlay, and Zhang (1997). It should also be acknowledged, however, that positive relation might arise as a methodological artifact. A volatile price path is more likely to move away from the limit price, increasing the likelihood that the order will be canceled. As noted in the previous

---

[16] Angel (1994) predicts that the execution probability of limit orders will be increasing in the rate of order flow arrival. We tested his prediction using the total number of limit and market orders as a proxy for the rate of arrival of orders. Since this measure is highly correlated with capitalization and turnover, , we used as controls only average price and σ_r. The results were supportive of Angel's prediction: all measures of execution proportions were increasing in the arrival rate of orders and were highly statistically significant.

section, this censoring is not, as the model specification would require, independent of the conditional execution probability.

In all specifications, the two volatility measures have similar properties. This feature is significant. In Handa and Schwartz (1996), limit order traders benefit from transient liquidity pressures. The σ_r measure impounds the transient price components, while the σ_w measure in principle does not. That these two measures play similar roles suggests that volatility effects do not arise solely from transient price movements.

To further explore volatility effects, we next consider specifications in which systematic and unsystematic risk are differentiated. The estimates, reported in Table 7, do not suggest that the distinction is an important one. In most specifications, the signs of the coefficients of these variables are identical to those reported for the total volatility measures. The effects are, however, statistically weaker.

The final set of volatility proxy measures volatility that is trade-related, i.e., volatility in the sense of pick-off risk. Table 8 presents these estimates. The absolute level of trade-related volatility, $\sigma_{w,x}$, appears to behave in these specifications in much the same manner as σ_r or σ_w.

In summary, the pattern of volatility associations is generally uniform across a wide range of volatility measures and proxies. Higher volatility is associated with a diminished use of limit orders in general and of limit orders priced at the quote or better, but with higher use of limit orders behind the quote. It is furthermore associated with increased likelihood of execution, shorter time to execution, and diminished book depth.

Table 9 reports estimates of specifications that include as a regressor either the percentage of institutional holdings (as a proxy for institutional trading) or the average number of odd lot trades (a proxy for retail trading). The coefficient of institutional holdings is positive for the submission proportions of all limit orders and also that of limit orders priced at the quote or better, but negative for the proportion priced behind the quote. The odd-lot coefficients are generally of the opposite sign, though of lower significance. These estimates suggest that while institutions are relatively heavy users of limit orders, they are less likely to provide depth away from the market.

7. Hidden and fleeting orders

Limit orders are sometimes viewed as supplying liquidity in a manner similar to (and competing with) dealer quotes. This analogy presumes that limit orders are relatively visible and persistent, like the bids and offers of a dealer who is maintaining a market presence. In fact, however, many limit orders are either hidden or short-lived. This section discusses such orders.

a. Hidden orders

The Island trading protocol allows traders to designate that an order not be displayed. The no-display option is a common feature of electronic book systems. In Island (and most of these systems), the hidden quantities lose priority to visible quantities at the same price. From a market design viewpoint, they are thought to encourage traders to supply liquidity when they might be reluctant to disclose the full size of the amount sought.

Our data report executions of hidden orders, but not submissions or cancellations. Our estimates can only suggest, therefore, a lower limit to the usage of these orders. These are reported in Table 10. Executed hidden orders constitute only about 3% of submitted limit orders (defined as submissions of visible limit orders and executed hidden orders), and about 2% by share amounts. They account, however, for almost 12% of all order executions and executed shares. This suggests a more significant presence.

b. Fleeting orders

We have noted that a large number of orders submitted to Island are canceled almost immediately. We term limit orders canceled within two seconds of their submission “fleeting”. Table 10 reports summary statistics. On average in the full sample, fleeting orders constitute 27.7% of all visible orders and 32.5% relative to shares in all visible orders. In the subgroup means, relative usage declines modestly with capitalization, average trades and σ_r. Table 11 presents summary statistics on the pricing of these orders. Fleeting orders are primarily submitted at prices that better the pre-existing bid or ask.

There are several possible explanations for the use of fleeting limit orders. One possibility is that Island receives these orders from automated order routing systems, which act as intelligent agents for customer orders. The strategies used by these systems frequently involve successive attempts to achieve execution at different market centers. For example, if Archipelago receives a marketable order at a time when Island's limit order book posts the best prices, Archipelago routes the order (or part of it) to Island for execution. If the order sent to Island fails to execute, say because the Island prices are no longer available, Archipelago essentially cancels the Island order and submits one to another market center.[17]

Searching for the best prices in the market may take time, and therefore the ability to cancel orders very quickly on Island (say by specifying a very short time-in-force for the order) is very important. Sophisticated systems can also create synthetic order types that take advantage of the ability to submit and cancel orders quickly. For example, Archipelago has a Now Order type that is matched against its book or routed for execution to a select group of market participants that have direct connections to Archipelago and can accept immediate-or-cancel orders. REDIBook, for example, has a special Limit Sweep Order that, when submitted to REDIBook, generates multiple orders seeking immediate execution that are routed to ECNs and market makers at multiple prices between the NBBO and the limit price. From these examples, it appears that many

[17] The function performed by an order routing system is essentially one of brokerage (as opposed to market making). Many of the systems, however, are implemented by the ECNs themselves or by brokers with close ties to ECNs. Archipelago and REDIBook, for example, incorporate order routing functions into their interfaces. These systems are sometimes generically referred to as smart order routing technology (SORT) systems. Both Smart Order Routing Technology and SORT, however, are service marks of MarketXT.

of the limit orders generated by these systems are directed at removing liquidity from the market, rather than supplying it.[18]

Another possible reason for a fleeting limit order is that the submitter wants to fish for hidden orders that better the opposing quote. A buyer, for example, might submit an order priced just short of the ask quote, hoping to trade against any hidden sell orders. Here as well a fleeting limit order represents a liquidity demander, rather than a supplier. The question then arises as to why the buyer's order in the above example needs to be visible, even briefly. A hidden order would accomplish the same thing without revealing the buyer's interest. Our data cannot characterize the extent of such practices. The fact that many of the fleeting orders are visible, though, suggests that finding hidden sellers is not the only motive, and that the brief display serves some purpose. The display might signal tentative buying interest to prospective sellers, without going so far as to provide them with a firm option.

A final possibility is a manipulative tactic known as "spoofing". To manipulate, a trader places a visible order in the opposite direction of the trade that is genuinely desired. For example, a seller might post a small buy order priced above the current bid, in hopes of convincing other buyers to match or outbid. If this occurs, the trader can sell into this (higher) price. It is necessary here that the order be visible. The practice resembles "shilling" by an auction seller, but there are some significant differences. In the stock market, the manipulator runs the risk that the spurious bid will be hit by some other seller, increasing the manipulator's long position. On the other hand, the Nasdaq market includes one group of buyers who are compelled to match the manipulator's spurious bid: dealers whose order preferencing arrangements require them to fill at the best prevailing price. This might make the manipulative strategy an appealing one. Both the NASD and SEC are conducting investigations and maintaining surveillance, however,

---

[18] In light of the ambiguity in classifying fleeting limit orders into liquidity demanding or supplying, we repeated the analysis of submission and execution proportions without fleeting limit orders. The results were qualitatively similar to those presented and discussed in Section 6.

(see Connor (2000)). The possibility of detection and prosecution is significant, and for this reason we doubt that such tactics lie behind the bulk of the fleeting orders.

8. Island's market presence

Island is only one venue in a broader market that comprises other ECN's and traditional dealers. In this section we examine the relative extent of Island activity, and firm and market characteristics to which it is related.

Table 12 presents summary statistics on Island market share. For the average firm in our sample, Island's market share is roughly 6.2% by trades and 3.5% by volume. In the ranked subgroups, this share increases with capitalization, the number of average trades, and σ_r. Figure 5 presents a log/log plot of the average number of Island trades vs. the total number of Nasdaq trades. The slope of the log/log best fit line is 1.7, which suggests that within the sample Island trades rise as the 1.7^{th} power of Nasdaq trades. In other words, Island's share increases for more active stocks.

The estimates in Table 12 suggest that Island's market share is larger for more volatile stocks. In cross-sectional regression analyses, this was confirmed in the presence of the control variables (capitalization, price and turnover) for all of our volatility proxies.

The positive relation between market share and volatility may reflect several mechanisms. The growth in ECN trading volume has been attributed in the popular press to increased day trading. While we have no direct evidence on this, our market share estimates are positively correlated with odd-lot trading volume (a measure of retail activity) and negatively correlated to institutional ownership. There is also evidence that day traders prefer volatile stocks. (Several popular how-to guides cite high volatility as a requirement for a stock to be an attractive candidate for day-trading.) This is consistent with our evidence.

An alternative measure of Island's impact is the extent to which Island sets or matches the market price (the NBBO). Table 13 shows that on average Island matched the best bid roughly 20% of the time and the best offer roughly 19% of the time. Much less frequently, however, was Island alone at the bid or the ask (4% of the time). Only 0.2% of the time was Island alone at both the bid and ask. The market share and quoting figures suggest that Island does not dominate trading in these stocks. This supports our empirical presumption that market variables used as explanatory variables are exogenous to our analysis.

9. Conclusions

The Island ECN is an electronic limit order book that trades Nasdaq National Market stocks, and offers traders an alternative to Nasdaq's traditional dealer-mediated trading. The analysis in this paper focuses on the cross-sectional relationships between volatility and measures of Island's activity.

Economic models of individual behavior identify several effects of volatility on order choice. For a given limit order, higher volatility is usually assumed to lead to higher probability of execution and shorter expected time to fill. Lo, MacKinlay, and Zhang (1997) find this to be the case for NYSE limit orders. We confirm this finding for Island orders. On the other hand, higher volatility may increase the cost of being picked-off or chasing a price that has moved away from the limit order. To quantify pick-off risk as distinct from total volatility, we employ as one of our measures the return variance that can be attributed to trades. Handa and Schwartz (1996) point out that transitory volatility is a source of profit for limit orders. We therefore examine both total transaction price volatility (which includes mechanism-related volatility, e.g., that arising from bid-ask bounce) and the volatility of the implicit random-walk component of the stock price (which excludes mechanism-related volatility). We furthermore separately examine the effects of systematic and unsystematic risk.

For all measures, we find that higher volatility is associated with a lower proportion of limit orders in Island's incoming flow. This suggests that for the representative trader, expected costs of pick-off risk or chasing the stock outweigh the benefit of increased execution likelihood or expected profits from transitory volatility. We also find for all measures that higher volatility is associated with a larger proportion of limit orders being entered at prices away from the market, i.e., priced less aggressively.

Equilibrium economic analysis introduces additional considerations. In Foucault (1999), an increase in the pick-off risk component of volatility causes limit order traders to price their orders less aggressively. This is consistent with our findings on limit order pricing noted above. Foucault's analysis furthermore suggests an equilibrium effect. When limit orders are priced less aggressively, the spread widens and book depths drop, making market orders more costly. This engenders a shift in the order mix toward limit orders and a lower probability of execution. Our findings are not consistent with a major role for this mechanism.

Island offers limit order submitters a "no display" option, wherein an order is entered on the book, but is not visible to any other market participants. Submissions of such hidden orders are not reflected in our data set. Executions of hidden orders are noted, however, and these suggest substantial usage of these orders. Executions of hidden orders constitute roughly ten percent of all executions. Furthermore, many limit orders are canceled almost immediately after they are submitted. We term orders canceled in two seconds or less "fleeting". Fleeting orders constitute 27.7% of all limit order submissions. Possible motives for these orders include probing for hidden orders, communicating tentative trading interest, and manipulative "spoofing".

Island's market share varies considerably across firms, and is positively related to overall Nasdaq activity in the stock. Thus, while Nasdaq activity is concentrated in firms that are larger (by market value), the concentration of Island's trading is even more pronounced.

These results suggest several directions for subsequent research. First, the concentration of Island's activity in larger firms raises concerns about the viability of the electronic limit order book as the primary mechanism for low-capitalization or low-activity firms. The importance of this issue for public policy warrants further examination.

Glosten (1994) suggests that there are powerful forces favoring the consolidation of trading activity in one electronic limit order book. In this connection, it is worth emphasizing that Island is but one of two highly successful ECNs. The market share of Instinet (the other) is larger than Island's on the basis of trading volume, but smaller on the basis of number of trades. The success of both ECN's indicates the likely importance of trading clientele attributes as determinants of ECN viability.

In many economic models limit orders are characterized as being widely visible and persistent, much like dealer quotes. Furthermore, regulatory initiatives such as the SEC's Order Handling Rules focus on protecting the rights of limit order traders against dealers. From this perspective, limit orders compete with, and are therefore in some sense equivalent to, dealer quotes as sources of liquidity. Many of the Island limit orders, however, are hidden, and a large fraction are canceled almost immediately after submission. These orders are quite different, therefore, from dealer quotes. Economic analysis of such orders and the strategies that rely on them constitute another worthwhile research direction.

The analysis in this paper is cross-sectional, and attempts to relate firm-specific characteristics to average attributes of Island activity. There is also, however, substantial dynamic variation in activity. The depth (available liquidity) on Island's book, for example, is highly variable over time. We are in the process of exploring the nature of this variation and the manner in which the limit order book adjusts to market shocks.

References

Ahn, H.-J., Bae, K.-H., Chan, K., 2001. Limit orders, depth and volatility: Evidence from then 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.

Barclay, M. J., Hendershott, T., McCormick, D. T., 2001. Electronic communications networks and market quality. Unpublished working paper. University of Rochester.

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.

Chakravarty, S., Holden, C. W., 1995. An integrated model of market and limit orders. Journal of Financial Intermediation 4, 213-241.

Chan, K. C., Christie, W. G., Schultz, P. H., 1995. Market structure and the intraday pattern of bid-ask spreads for NASDAQ securities. Journal of Business 68, 35-60.

Cho, J.-W., Nelling, E., 2000. The probability of limit order execution. Financial Analysts Journal 56, 28-33.

Chung, K. H., Van Ness, B. F., Van Ness, R. A., 1999. Limit orders and the bid-ask spread. Journal of Financial Economics 53, 255-287.

Clark, P. K., 1973. A subordinated stochastic process model with finite variance for speculative prices. Econometrica 41.

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.

Connor, John. 28 February 2000. No joke: NASD plans crackdown on "spoofing," placing and canceling a quote to spark a move. *Wall Street Journal*.

Conrad, J., Johnson, K. M., Wahal, S., 2001. Alternative trading systems. Unpublished working paper. University of North Carolina, Kenan-Flagler Business School.

Copeland, T., Galai, D., 1983. Information effects and the bid-ask spread. Journal of Finance 38, 1457-1469.

Coppejans, M., Domowitz, I., Madhavan, A., 2001. Liquidity in an automated auction. Unpublished working paper. Department of Economics, Duke University.

Domowitz, I., Wang, J., 1994. Auctions as algorithms. Journal of Economic Dynamics and Control 18, 29-60.

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

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.

Glosten, L. R., Milgrom, P. R., 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71-100.

Goldstein, M. A., Kavajecz, K. A., 2000. Liquidity provision during circuit breakers and extreme market movements. Unpublished working paper. Department of Finance, Wharton School, University of Pennsylvania.

Griffiths, M. D., Smith, B. F., Turnbull, D. A. S., White, R. W., 2000. The costs and determinants of order aggressiveness. Journal of Financial Economics 56, 65-88.

Handa, P., Schwartz, R. A., 1996. Limit order trading. Journal of Finance 51, 1835-1861.

Harris, L., 1998. Optimal dynamic order submission strategies in some stylized trading problems. Financial Markets, Institutions and Instruments 7.

Harris, L. E., Hasbrouck, J., 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. Journal of Financial and Quantitative Analysis 31, 213-31.

Harris, L. E., Panchapagesan, V., 1999. The information content of the limit order book:

Evidence from NYSE specialist actions. Unpublished working paper. Marshall School of Business, University of Southern California.

Hasbrouck, J., 1991. The summary informativeness of stock trades: An econometric analysis. Review of Financial Studies 4, 571-95.

Huang, R., 2000. Price discovery by ECNs and Nasdaq market makers. Unpublished working paper. Owen School of Management, Vanderbilt University.

Island. 2001. "Island's Market Leadership." Web page, [accessed 20 August 2001]. Available at http://www.island.com/pressroom/leadership.htm.

Lo, A. W., MacKinlay, A. C., Zhang, J., 1997. Econometric models of limit order execution. Unpublished working paper. NBER.

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.

Rousseeuw, P. J., Leroy, A. M., 1987. Robust Regression and Outlier Detection. John Wiley & Sons, New York.

Russell, J. R., Engle, R. F., 1998. Econometric analysis of discrete-valued irregularly-spaced financial transactions data using a new Autoregressive Conditional Multinomial model. Unpublished working paper. University of Chicago, Graduate School of Business.

Sandas, P., 1999. Adverse selection and competitive market making: evidence from a pure limit order book. Unpublished working paper. Wharton School, University of Pennsylvania.

Seppi, D. J., 1997. Liquidity provision with limit orders and a strategic specialist. Review of Financial Studies 10, 103-150.

Simaan, Y., Weaver, D. G., Whitcomb, D. K., 1998. The quotation behavior of ECN's and Nasdaq market makers. Unpublished working paper. Zicklin School of Business, Baruch College.

Smith, J. W., 2000. Market vs. limit order submission behavior at a Nasdaq market maker. Unpublished working paper. NASD Economic Research.

U.S. Securities and Exchange Commission. 1996. *Order execution obligations*. Release No. 34-37619A .

US Securities and Exchange Commission. 1998. *Regulation of exchanges and alterative trading systems*.
Release No. 34-40760.

Table 1. Summary statistics

Summary statistics (across firms) over the 64 trading days in the fourth quarter of 1999. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999). Capitalization, market price, and trading volume are taken from CRSP; σ_r is the standard deviation of daily CRSP returns; spreads are derived from Nastraq; institutional holdings are from the Value Line Investment survey; odd-lot trades were supplied by NASD Economic Research.

		Avg equity mkt. cap. ($MM)	Avg daily trades	σ_r (daily return)	Avg daily volume (1,000 shares)	Median daily turnover	Average price ($/share)	Institutional holdings	Spread ($/share)	Spread (% of quote midpt)	Avg daily odd lot trades
Total Sample	Mean	10,205	2,677	0.0436	1,873	1.288	63.03	50.8%	0.2563	0.46%	346
	Median	3,081	1,066	0.0433	877	1.107	49.82	51.3%	0.1871	0.44%	125
	SD	38,104	4,413	0.0169	3,504	0.946	45.66	25.6%	0.2180	0.25%	786
	Min	824	16	0.0018	7	0.028	8.40	6.0%	0.0520	0.07%	5
	Max	494,932	28,654	0.1083	30,073	5.208	326.58	97.9%	1.9103	2.79%	8,348
	Nobs	300	300	300	300	300	300	290	300	300	300
Means for mkt. cap. groups	Low	1,500	654	0.0386	549	0.952	39.16	50.8%	0.2520	0.63%	75
	Medium	3,169	1,474	0.0470	1,051	1.322	58.26	50.7%	0.2757	0.49%	147
	High	25,947	5,904	0.0452	4,017	1.590	91.67	50.7%	0.2413	0.28%	817
Means for trade groups	Low	1,953	326	0.0326	314	0.541	41.84	48.4%	0.2730	0.63%	49
	Medium	3,772	1,202	0.0491	933	1.257	62.19	57.1%	0.2849	0.48%	134
	High	24,891	6,504	0.0491	4,371	2.066	85.06	47.1%	0.2109	0.28%	855
Means for σ_r groups	Low	18,100	2,290	0.0257	2,038	0.565	44.20	51.9%	0.1989	0.49%	317
	Medium	7,304	2,548	0.0429	1,872	1.424	63.11	57.3%	0.2162	0.41%	271
	High	5,212	3,193	0.0622	1,708	1.874	81.79	42.4%	0.3538	0.49%	450

Table 2. Island summary statistics

Summary statistics (across firms) over the 64 trading days in the fourth quarter of 1999. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999). All data are tabulated from Island ITCH data. On the Island system, all orders carry a limit price. Market orders are defined as orders that are matched upon arrival (and so never appear in the book).

		Avg daily no of limit orders	Avg size of limit order (shares)	Avg no of cancelations	Avg size of cancelation (shares)	Avg no of filled limit orders	Avg size of limit order fill (shares)	Avg daily no. of market orders	Avg market order size (shares)	Avg daily no. of odd-lot trades
Total Sample	Mean	965.9	572.4	672.0	617.7	275.7	389.7	339.9	335.0	57.7
	Median	285.3	585.2	221.0	627.3	51.1	380.1	60.7	329.3	7.6
	SD	1,764.8	158.1	1,144.6	157.5	602.5	133.9	760.7	110.2	153.2
	Min	3.9	214.1	2.7	241.9	0.0	148.8	0.0	123.3	0.0
	Max	11,992.4	985.3	6,963.5	1,032.0	4,726.7	931.8	6,123.6	742.7	1,498.7
	Nobs	300	300	300	300	300	299	300	299	300
Means for mkt. cap. groups	Low	157.7	612.5	119.3	644.6	34.0	414.9	40.5	363.5	6.4
	Medium	461.3	567.5	337.7	613.9	113.7	379.9	136.8	329.2	19.5
	High	2,278.6	537.2	1,558.9	594.6	679.3	374.5	842.4	312.5	147.3
Means for trade groups	Low	67.1	631.5	58.4	653.1	7.7	423.1	8.6	377.5	1.0
	Medium	332.5	553.3	257.2	599.5	69.5	372.2	80.7	322.1	12.2
	High	2,498.0	532.5	1,700.3	600.4	749.8	374.1	930.4	305.8	159.9
Means for σ_r groups	Low	680.0	668.1	488.9	694.0	177.9	463.2	229.8	404.5	30.5
	Medium	913.6	592.6	659.9	644.1	239.1	403.6	289.6	344.7	44.2
	High	1,304.0	456.6	867.1	515.0	410.0	303.0	500.4	256.5	98.4

Table 3. Submission and execution proportions for Island limit orders.

Summary statistics (across firms) over the 64 trading days in the fourth quarter of 1999 for visible (non-hidden) Island limit orders. The firm sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999). All data are tabulated from Island ITCH data and the Nastraq database. The order sample is all visible (non-hidden) limit orders entered into the Island system that are not matched upon arrival.

		Limit order submissions relative to all orders:			Limit order submissions (shares) relative to shares in all orders:			Limit order executions			Limit order executions (shares)		
			Price relative to quote			Price relative to quote			Price relative to quote			Price relative to quote	
		All prices	At or better	Away	All prices	At or better	Away	All prices	At or better	Away	All prices	At or better	Away
Total Sample	Mean	84.3%	53.8%	30.5%	90.3%	58.5%	31.8%	16.0%	15.4%	17.8%	11.1%	10.3%	13.1%
	Median	83.4%	49.2%	34.0%	90.1%	55.7%	33.8%	17.2%	15.2%	19.4%	11.0%	9.7%	13.7%
	SD	7.8%	25.0%	17.6%	5.1%	23.9%	19.3%	8.2%	8.6%	7.7%	6.3%	6.3%	6.4%
	Min	66.3%	13.4%	0.0%	76.3%	15.6%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
	Max	100.0%	100.0%	59.1%	100.0%	100.0%	66.7%	34.7%	38.9%	33.5%	31.0%	33.7%	31.8%
	Nobs	300	300	300	300	300	300	300	300	298	300	300	298
Means for mkt. cap. groups	Low	88.9%	69.4%	19.4%	93.1%	73.7%	19.4%	11.3%	10.5%	13.8%	7.7%	6.9%	10.3%
	Medium	84.0%	53.8%	30.2%	90.1%	58.9%	31.2%	16.4%	15.6%	18.6%	11.3%	10.2%	13.7%
	High	80.1%	38.3%	41.9%	87.7%	42.9%	44.7%	20.3%	20.3%	21.0%	14.3%	13.7%	15.3%
Means for trade groups	Low	92.3%	80.9%	11.4%	95.3%	84.3%	11.1%	7.6%	7.0%	10.6%	5.0%	4.5%	8.0%
	Medium	83.6%	51.6%	32.0%	89.8%	57.4%	32.4%	16.9%	15.6%	18.9%	11.5%	10.1%	13.8%
	High	77.1%	29.0%	48.2%	85.7%	33.9%	51.9%	23.5%	23.7%	23.8%	16.8%	16.1%	17.5%
Means for σ_r groups	Low	90.8%	75.6%	15.2%	94.2%	78.4%	15.8%	9.1%	9.1%	11.1%	6.5%	6.4%	8.6%
	Medium	83.0%	48.4%	34.5%	89.4%	53.6%	35.8%	17.4%	16.3%	19.4%	12.0%	10.8%	14.2%
	High	79.2%	37.4%	41.8%	87.2%	43.5%	43.8%	21.5%	20.9%	22.7%	14.9%	13.6%	16.5%

Table 4. Depth summary statistics

Summary statistics (across firms) over the 64 trading days in the fourth quarter of 1999 for visible (non-hidden) Island limit orders. The firm sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999). All data are tabulated from Island ITCH data and the Nastraq database. The order sample is all visible (non-hidden) limit orders entered into the Island system that are not matched upon arrival. For firm *i* and five-minute interval *t*, we compute $D_{i,t}^{ask}$, the total dollar value of all limit orders on the ask side of the book at the end of *t*. We then compute for each firm (across time) the mean, median, standard deviation and coefficient of variation of $D_{i,t}^{ask}$. Bid-side statistics are computed analogously.

		Ask side:				Bid side:			
		Avg. value ($1,000)	Median value ($1,000)	SD of value ($1,000)	Coeff. of variation	Avg. value ($1,000)	Median value ($1,000)	SD of value ($1,000)	Coeff. of variation
Total Sample	Mean	237.3	210.1	145.5	1.4	142.3	118.2	99.4	1.7
	Median	66.1	42.3	66.1	1.0	38.6	22.1	48.2	1.2
	SD	565.8	553.4	255.8	1.4	321.2	299.2	161.3	1.6
	Min	0.1	0.0	1.0	0.2	0.1	0.0	1.4	0.3
	Max	5,392.8	5,465.8	2,130.1	17.1	2,698.7	2,535.4	1,486.9	13.3
	Nobs	300	300	300	300	300	300	300	300
Means for mkt. cap. groups	Low	48.8	37.9	40.2	1.9	29.9	21.2	29.7	2.4
	Medium	117.8	98.0	86.3	1.5	69.6	52.5	61.6	1.6
	High	545.2	494.3	309.9	1.0	327.3	280.8	206.8	1.1
Means for trade groups	Low	15.8	8.3	22.2	2.4	10.3	3.9	17.5	2.9
	Medium	69.4	48.3	71.8	1.2	46.4	27.7	56.6	1.4
	High	626.6	573.7	342.3	0.7	370.1	322.9	224.0	0.8
Means for σ_r groups	Low	250.4	233.5	127.7	2.2	142.4	124.6	81.0	2.7
	Medium	193.8	169.2	134.7	1.0	118.0	95.4	94.1	1.2
	High	267.5	227.7	174.0	1.1	166.5	134.6	123.0	1.2

Table 5. Volatility summary statistics

The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999). σ_r is the standard deviation of the daily CRSP return. σ_w is the standard deviation of the random-walk component of the stock price; $\sigma_{w,x}$ is the standard deviation of the contribution to the random-walk component attributable to signed trades. σ_w and $\sigma_{w,x}$ are estimated using the Hasbrouck (1991) procedure applied to a vector autoregression of quote-midpoint returns and signed trades, aggregated over one-minute intervals. They are scaled to reflect volatility over a 6. 5 hour trading day. σ_r, σ_w, and $\sigma_{w,x}$ are estimated over the 64 trading days in the fourth quarter of 1999. $\beta_i\sigma_M$ is the standard deviation of systematic risk; $\sigma(e_{it})$ is the standard deviation of unsystematic risk. Both are based on the market model $r_{it} = \alpha_i + \beta_i r_{Mt} + e_{it}$, where r_{Mt} is the return on the CRSP value-weighted portfolio, estimated using daily CRSP data from October 1, 1996 to September 30, 1999. The table presents summary statistics across firms.

		σ_r	σ_w	$\sigma_{w,x}$	$\beta_i \sigma_M$	$\sigma(e_{it})$
Total Sample	Mean	0.0436	0.0383	0.0193	0.0133	0.0330
	Median	0.0433	0.0383	0.0197	0.0132	0.0331
	SD	0.0169	0.0138	0.0076	0.0050	0.0108
	Min	0.0018	0.0029	0.0015	0.0025	0.0136
	Max	0.1083	0.0890	0.0438	0.0310	0.0620
	Nobs	300	300	300	211	211
Means for mkt. cap. groups	Low	0.0386	0.0346	0.0167	0.0106	0.0314
	Medium	0.0470	0.0412	0.0211	0.0130	0.0348
	High	0.0452	0.0392	0.0202	0.0168	0.0332
Means for trade groups	Low	0.0326	0.0298	0.0142	0.0095	0.0277
	Medium	0.0491	0.0428	0.0221	0.0136	0.0352
	High	0.0491	0.0424	0.0217	0.0185	0.0380
Means for σ_r groups	Low	0.0257	0.0244	0.0123	0.0106	0.0251
	Medium	0.0429	0.0387	0.0199	0.0153	0.0374
	High	0.0622	0.0519	0.0258	0.0157	0.0427

Table 6. Island limit orders and total volatility

The table presents regression coefficient estimates (using the indicated specification) for submission proportions, execution probabilities, dollar value of limit orders on the ask side of the book, and limit order durations. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999) over the 64 trading days in the fourth quarter of 1999. “Regression” and “logit regression” specifications are estimated across firms in the sample. The duration specifications are estimated for a sample consisting of 2,000 randomly selected limit orders for each firm, adjusted to weight all firms equally. Numbers in parentheses are coefficient estimates divided by the asymptotic standard error of estimate. The latter standard errors are heteroskedasticity-consistent in the regression and logit regression specifications (but not in the duration specifications). σ_r is the standard deviation of the daily CRSP return. σ_w is the standard deviation of the implicit random-walk component of the quote midpoint, estimated with one-minute data and rescaled to reflect volatility over a 6. 5 hour trading day.

Model specification	Dependent variable	Intercept	σ_r	σ_w	Log(Avg cap)	Avg price	Log(Median turnover)	No. Obs.	R^2	Scale	Shape
Logit regression	Limit subm. prop.	10.158	-20.970		-0.347	0.004	-0.486	298	0.753		
	(all prices)	(20.44)	(-7.56)		(-17.81)	(7.42)	(-7.58)				
	Limit subm. prop.	15.901	-34.561		-0.657	0.006	-0.774	299	0.821		
	(quote or better)	(21.55)	(-8.10)		(-22.44)	(6.96)	(-8.18)				
	Limit subm. prop.	-15.205	34.834		0.584	-0.005	0.765	298	0.744		
	(away)	(-16.24)	(7.72)		(15.78)	(-5.78)	(7.89)				
	Limit exec. prop.	-10.393	22.396		0.356	-0.004	0.506	298	0.754		
	(all prices)	(-19.65)	(7.74)		(17.09)	(-7.35)	(7.87)				
	Limit exec. prop.	-11.401	22.098		0.400	-0.004	0.502	298	0.762		
	(quote or better)	(-21.24)	(7.91)		(18.60)	(-7.24)	(8.17)				
	Limit exec. prop.	-6.568	14.652		0.204	-0.003	0.325	283	0.570		
	(away)	(-14.23)	(6.72)		(10.82)	(-6.11)	(5.93)				
Regression	Ask Depth	-8,401.492	-992.833		402.860	-3.075	84.681	300	0.479		
		(-5.18)	(-1.00)		(5.25)	(-3.63)	(3.10)				
Duration	Time to execution	14.832	-23.159		-0.414	0.003	-0.466			3.049	-1.893
		(110.06)	(-63.91)		(-71.68)	(19.92)	(-67.28)			(365.68)	(-140.72)
Logit regression	Limit subm. prop.	10.436		-25.472	-0.357	0.004	-0.495	298	0.753		
	(all prices)	(20.56)		(-7.43)	(-18.53)	(7.74)	(-7.86)				
	Limit subm. prop.	16.482		-43.326	-0.677	0.006	-0.777	299	0.828		
	(quote or better)	(22.07)		(-8.31)	(-23.41)	(7.18)	(-8.54)				
	Limit subm. prop.	-15.994		45.651	0.610	-0.005	0.753	298	0.762		
	(away)	(-16.56)		(8.13)	(16.27)	(-6.17)	(8.32)				
	Limit exec. prop.	-10.693		27.242	0.366	-0.004	0.515	298	0.755		
	(all prices)	(-19.84)		(7.61)	(17.84)	(-7.77)	(8.17)				
	Limit exec. prop.	-11.703		26.936	0.410	-0.004	0.510	298	0.763		
	(quote or better)	(-21.50)		(7.82)	(19.40)	(-7.55)	(8.48)				
	Limit exec. prop.	-6.592		16.182	0.206	-0.003	0.341	283	0.547		
	(away)	(-13.87)		(5.88)	(10.84)	(-6.01)	(6.10)				
Regression	Ask Depth	-8,372.235		-1,359.334	401.898	-3.062	85.195	300	0.479		
		(-5.16)		(-1.15)	(5.23)	(-3.60)	(3.15)				
Duration	Time to execution	15.017		-27.181	-0.422	0.003	-0.484			3.052	-1.902
		(109.54)		(-61.20)	(-72.30)	(20.43)	(-70.01)			(364.36)	(-139.85)

Table 7. Island limit orders and systematic/unsystematic volatility

The table presents regression coefficient estimates (using the indicated specification) for submission proportions, execution probabilities, dollar value of limit orders on the ask side of the book, and limit order durations. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999) over the 64 trading days in the fourth quarter of 1999. "Regression" and "logit regression" specifications are estimated across firms in the sample. The duration specifications are estimated for a sample consisting of 2,000 randomly selected limit orders for each firm, adjusted to weight all firms equally. Numbers in parentheses are coefficient estimates divided by the asymptotic standard error of estimate. The latter standard errors are heteroskedasticity-consistent in the regression and logit regression specifications (but not in the duration specifications). $\beta_i\sigma_M$ is the standard deviation of systematic risk; $\sigma(e_{it})$ is the standard deviation of unsystematic risk. Both are based on the market model $r_{it} = \alpha_i + \beta_i r_{Mt} + e_{it}$, where r_{Mt} is the return on the CRSP value-weighted portfolio, estimated using daily CRSP data from October 1, 1996 to September 30, 1999.

Model specification	Dependent variable	Intercept	$\beta_i\sigma_M$	$\sigma(e_{it})$	Log(Avg cap)	Avg price	Log(Median turnover)	No. Obs.	R^2	Scale	Shape
Logit regression	Limit subm. prop.	8.272	-23.336		-0.281	0.003	-0.638	209	0.702		
	(all prices)	(14.82)	(-1.82)		(-8.53)	(2.86)	(-5.25)				
	Limit subm. prop.	12.270	-50.171		-0.518	0.005	-0.944	210	0.758		
	(quote or better)	(14.71)	(-2.46)		(-10.71)	(2.89)	(-5.40)				
	Limit subm. prop.	-11.470	53.201		0.442	-0.005	0.974	209	0.667		
	(away)	(-11.95)	(2.24)		(8.33)	(-2.59)	(5.17)				
	Limit exec. prop.	-8.379	23.754		0.285	-0.003	0.666	209	0.700		
	(all prices)	(-14.29)	(1.80)		(8.32)	(-2.71)	(5.36)				
	Limit exec. prop.	-9.488	22.484		0.334	-0.003	0.649	209	0.718		
	(quote or better)	(-16.50)	(1.81)		(10.06)	(-2.83)	(5.49)				
	Limit exec. prop.	-5.118	13.655		0.152	-0.002	0.414	194	0.467		
	(away)	(-9.21)	(1.30)		(4.95)	(-1.59)	(3.62)				
Regression	Ask Depth	-8,649.081	4,357.588		407.329	-2.735	13.569	211	0.493		
		(-4.43)	(0.46)		(4.34)	(-2.10)	(0.43)				
Duration	Time to execution	13.288	-33.995		-0.336	0.000	-0.595			3.266	-1.532
		(88.37)	(-16.09)		(-45.98)	(0.54)	(-57.55)			(313.29)	(-153.07)
Logit regression	Limit subm. prop.	10.872		-25.268	-0.376	0.004	-0.513	209	0.738		
	(all prices)	(14.26)		(-3.00)	(-15.62)	(4.47)	(-3.91)				
	Limit subm. prop.	16.954		-43.757	-0.697	0.006	-0.757	210	0.794		
	(quote or better)	(13.89)		(-3.44)	(-17.47)	(4.38)	(-4.00)				
	Limit subm. prop.	-16.341		45.256	0.629	-0.006	0.787	209	0.702		
	(away)	(-10.43)		(3.24)	(11.67)	(-4.00)	(3.87)				
	Limit exec. prop.	-11.152		27.205	0.386	-0.004	0.526	209	0.739		
	(all prices)	(-14.09)		(3.17)	(15.25)	(-4.46)	(3.99)				
	Limit exec. prop.	-12.127		25.917	0.429	-0.004	0.516	209	0.754		
	(quote or better)	(-15.83)		(3.19)	(17.18)	(-4.33)	(4.13)				
	Limit exec. prop.	-6.978		18.400	0.217	-0.002	0.302	194	0.512		
	(away)	(-9.95)		(2.64)	(9.12)	(-2.79)	(2.56)				
Regression	Ask Depth	-9,159.383		5,031.732	425.725	-2.891	-12.648	211	0.497		
		(-4.77)		(1.96)	(4.70)	(-2.22)	(-0.46)				
Duration	Time to execution	16.945		-35.203	-0.475	0.002	-0.462			3.235	-1.582
		(101.86)		(-40.18)	(-67.71)	(8.45)	(-46.07)			(311.97)	(-146.15)

Table 8. Island limit orders and trade-related volatility

The table presents regression coefficient estimates (using the indicated specification) for submission proportions, execution probabilities, dollar value of limit orders on the ask side of the book, and limit order durations. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999) over the 64 trading days in the fourth quarter of 1999. "Regression" and "logit regression" specifications are estimated across firms in the sample. The duration specifications are estimated for a sample consisting of 2,000 randomly selected limit orders for each firm, adjusted to weight all firms equally. Numbers in parentheses are coefficient estimates divided by the asymptotic standard error of estimate. The latter standard errors are heteroskedasticity-consistent in the regression and logit regression specifications (but not in the duration specifications). $\sigma_{w,x}$ is the standard deviation of the contribution to the random-walk component attributable to signed trades. $\sigma_{w,x}$ is estimated using the Hasbrouck (1991) procedure applied to a vector autoregression of quote-midpoint returns and signed trades, aggregated over one-minute intervals and scaled to reflect volatility over a 6. 5 hour trading day.

Model specification	Dependent variable	Intercept	$\sigma_{w,x}$	Log(Avg cap)	Avg price	Log(Median turnover)	No. Obs.	R^2	Scale	Shape
Logit regression	Limit subm. prop.	9.223	-35.526	-0.311	0.003	-0.514	298	0.709		
	(all prices)	(19.93)	(-8.20)	(-16.41)	(3.87)	(-8.04)				
	Limit subm. prop.	14.371	-58.361	-0.599	0.004	-0.822	299	0.777		
	(quote or better)	(19.96)	(-9.29)	(-19.75)	(3.30)	(-8.83)				
	Limit subm. prop.	-13.886	64.464	0.530	-0.003	0.788	298	0.715		
	(away)	(-15.27)	(8.77)	(14.44)	(-2.90)	(8.23)				
	Limit exec. prop.	-9.416	38.485	0.318	-0.003	0.533	298	0.711		
	(all prices)	(-19.20)	(8.32)	(15.86)	(-3.83)	(8.21)				
	Limit exec. prop.	-10.430	37.801	0.362	-0.003	0.529	298	0.720		
	(quote or better)	(-20.79)	(7.96)	(17.45)	(-3.86)	(8.34)				
	Limit exec. prop.	-5.760	22.325	0.174	-0.002	0.345	283	0.506		
	(away)	(-12.69)	(6.45)	(9.08)	(-2.87)	(5.84)				
Regression	Ask Depth	-8,246.638	-6,597.081	399.715	-3.071	103.612	300	0.484		
		(-5.12)	(-3.01)	(5.22)	(-3.62)	(3.76)				
Duration	Time to execution	13.278	-31.166	-0.356	0.001	-0.461			3.088	-1.872
		(99.75)	(-39.29)	(-62.06)	(5.07)	(-63.16)			(365.59)	(-137.28)

Table 9. Island limit orders and investor characteristics

The table presents regression coefficient estimates (using the indicated specification) for submission proportions, execution probabilities, dollar value of limit orders on the ask side of the book, and limit order durations. The sample is the largest 300 Nasdaq National Market stocks (ranked by equity capitalization on September 30, 1999) over the 64 trading days in the fourth quarter of 1999. "Regression" and "logit regression" specifications are estimated across firms in the sample. The duration specifications are estimated for a sample consisting of 2,000 randomly selected limit orders for each firm, adjusted to weight all firms equally. Numbers in parentheses are coefficient estimates divided by the asymptotic standard error of estimate. The latter standard errors are heteroskedasticity-consistent in the regression and logit regression specifications (but not in the duration specifications) Institutional holdings are from the Value Line Investment survey; odd-lot trades were supplied by NASD Economic Research.

Model specification	Dependent variable	Intercept	Institutional holdings	No. of odd lot trades	Log(Avg cap)	Avg price	Log(Median turnover)	No. Obs.	R^2	Scale	Shape
Logit regression	Limit subm. prop.	7.718	0.006		-0.290	0.003	-0.743	288	0.702		
	(all prices)	(18.86)	(7.01)		(-15.57)	(5.03)	(-10.93)				
	Limit subm. prop.	11.706	0.012		-0.559	0.005	-1.202	289	0.776		
	(quote or better)	(19.16)	(9.35)		(-19.93)	(4.85)	(-12.60)				
	Limit subm. prop.	-11.103	-0.010		0.484	-0.004	1.169	288	0.674		
	(away)	(-14.03)	(-6.01)		(13.42)	(-3.52)	(11.23)				
	Limit exec. prop.	-7.784	-0.007		0.294	-0.003	0.779	288	0.699		
	(all prices)	(-17.87)	(-7.03)		(14.80)	(-4.90)	(11.20)				
	Limit exec. prop.	-8.792	-0.008		0.340	-0.004	0.781	288	0.728		
	(quote or better)	(-20.57)	(-8.39)		(17.51)	(-5.28)	(11.68)				
	Limit exec. prop.	-4.653	-0.005		0.157	-0.002	0.483	273	0.509		
	(away)	(-11.68)	(-6.68)		(8.60)	(-3.94)	(7.83)				
Regression	Ask Depth	-8,338.209	-3.515		407.293	-3.488	104.464	290	0.500		
		(-5.09)	(-4.21)		(5.24)	(-3.88)	(3.55)				
Duration	Time to execution	11.692	0.003		-0.322	0.001	-0.597			3.072	-1.922
		(91.02)	(16.38)		(-56.45)	(4.68)	(-83.35)			(351.45)	(-129.59)
Logit regression	Limit subm. prop.	6.162		-0.189	-0.103	0.005	-0.481	298	0.663		
	(all prices)	(7.30)		(-2.37)	(-1.31)	(5.06)	(-4.06)				
	Limit subm. prop.	8.448		-0.416	-0.160	0.008	-0.668	299	0.752		
	(quote or better)	(7.03)		(-3.66)	(-1.43)	(5.62)	(-4.03)				
	Limit subm. prop.	-7.470		0.445	0.059	-0.008	0.631	298	0.681		
	(away)	(-6.27)		(4.10)	(0.56)	(-5.48)	(4.03)				
	Limit exec. prop.	-6.223		0.191	0.105	-0.005	0.512	298	0.658		
	(all prices)	(-7.06)		(2.29)	(1.28)	(-4.83)	(4.17)				
	Limit exec. prop.	-7.308		0.186	0.154	-0.005	0.509	298	0.670		
	(quote or better)	(-8.56)		(2.30)	(1.95)	(-4.83)	(4.33)				
	Limit exec. prop.	-3.807		0.118	0.042	-0.003	0.324	283	0.464		
	(away)	(-5.39)		(1.92)	(0.68)	(-4.08)	(3.32)				
Regression	Ask Depth	-8,090.570		48.104	361.617	-3.777	28.249	300	0.482		
		(-4.86)		(2.41)	(4.51)	(-4.18)	(0.97)				
Duration	Time to execution	9.227		-0.270	-0.063	0.004	-0.363			3.050	-2.003
		(68.12)		(-41.24)	(-7.93)	(25.72)	(-43.33)			(348.53)	(-128.30)

Table 10. Hidden and fleeting orders

Summary statistics for the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. All data are tabulated from Island ITCH data. Hidden orders are those that were entered with a “no display” qualifier. Visible orders are limit orders not so qualified, that are not matched immediately on arrival. Fleeting orders are visible limit orders that are canceled (unexecuted) within two seconds of entry.

		Executions of hidden orders / All visible limit orders	Executed shares in hidden orders / Shares in all visible limit orders	Executions of hidden orders / All executions	Executed shares in hidden orders / All executed shares	Fleeting limit orders / All visible limit orders	Shares in fleeting limit orders / Shares in all visible limit orders
Total Sample	Mean	3.1%	1.8%	11.8%	11.8%	27.7%	32.5%
	Median	2.0%	1.1%	10.1%	10.3%	25.4%	29.8%
	SD	3.3%	1.9%	9.3%	9.5%	11.7%	12.0%
	Min	0.0%	0.0%	0.0%	0.0%	5.9%	5.8%
	Max	18.2%	11.7%	100.0%	100.0%	88.4%	91.5%
	Nobs	300	300	299	299	300	300
Means for mkt. cap. groups	Low	1.7%	1.0%	9.2%	9.2%	32.5%	36.7%
	Medium	3.0%	1.8%	11.5%	11.6%	27.6%	32.7%
	High	4.7%	2.7%	14.6%	14.5%	22.9%	28.1%
Means for trade groups	Low	0.9%	0.5%	7.7%	7.8%	36.9%	40.5%
	Medium	3.1%	1.8%	12.2%	12.3%	25.9%	31.6%
	High	5.4%	3.1%	15.3%	15.3%	20.2%	25.4%
Means for σ_r groups	Low	0.8%	0.5%	6.6%	6.4%	34.7%	37.5%
	Medium	2.8%	1.6%	10.7%	10.9%	26.9%	32.1%
	High	5.9%	3.3%	17.9%	18.0%	21.4%	28.0%

Table 11. Pricing of fleeting orders

Summary statistics for the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. All data are tabulated from Island ITCH data. Fleeting orders are visible limit orders that are canceled (unexecuted) within two seconds of entry.

		Buy orders relative to Island's bid			Sell orders relative to Island's ask		
		Better	At	Behind	Better	At	Behind
Total Sample	Mean	83.9%	6.7%	9.5%	85.6%	6.6%	7.9%
	Median	88.9%	5.1%	5.7%	87.9%	5.8%	5.8%
	SD	14.9%	6.1%	9.5%	10.5%	4.7%	7.0%
	Min	30.2%	0.0%	0.0%	47.7%	0.0%	0.0%
	Max	100.0%	33.3%	48.1%	100.0%	28.2%	40.6%
	Nobs	300	300	300	300	300	300
Means for mkt. cap. groups	Low	92.8%	3.2%	4.0%	91.7%	4.6%	3.8%
	Medium	85.8%	6.1%	8.1%	86.4%	6.4%	7.3%
	High	73.0%	10.8%	16.2%	78.7%	8.8%	12.6%
Means for trade groups	Low	96.4%	1.5%	2.2%	94.5%	3.1%	2.3%
	Medium	87.2%	5.7%	7.1%	86.7%	6.4%	6.9%
	High	68.0%	13.0%	19.1%	75.5%	10.1%	14.4%
Means for σ_r groups	Low	90.4%	4.0%	5.6%	91.0%	4.6%	4.4%
	Medium	83.3%	7.0%	9.8%	85.7%	6.7%	7.6%
	High	77.9%	9.2%	13.0%	80.0%	8.3%	11.7%

Table 12. Market share

Summary statistics for the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. Statistics are based on Island ITCH data and the Nastraq database.

		Island's market share in:	
		Trades	Volume
Total Sample	Mean	6.2%	3.5%
	Median	5.2%	2.7%
	SD	5.0%	3.5%
	Min	0.0%	0.0%
	Max	23.7%	20.3%
	Nobs	300	300
Means for mkt. cap. groups	Low	3.0%	1.6%
	Medium	6.0%	3.2%
	High	9.7%	5.8%
Means for trade groups	Low	1.8%	0.9%
	Medium	5.5%	2.8%
	High	11.3%	6.9%
Means for σ_r groups	Low	2.6%	1.3%
	Medium	6.6%	3.5%
	High	9.6%	5.7%

Table 13. Island's quotes and the NBBO

Summary statistics for the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. Statistics are based on Island ITCH data and the Nastraq database.

		Proportion of time Island quotes are at NBBO					
		At ask	At bid	At both	Alone at the ask	Alone at the bid	Alone at both
Total Sample	Mean	19.6%	18.6%	4.0%	3.9%	3.9%	0.2%
	Median	14.5%	14.4%	1.3%	2.9%	3.0%	0.1%
	SD	15.2%	13.9%	6.5%	3.3%	3.2%	0.5%
	Min	0.0%	0.7%	0.0%	0.0%	0.0%	0.0%
	Max	73.3%	69.0%	45.8%	21.6%	21.9%	4.3%
	Nobs	300	300	300	300	300	300
Means for mkt. cap. groups	Low	10.1%	9.6%	1.0%	2.2%	2.5%	0.1%
	Medium	17.0%	16.7%	2.7%	3.9%	4.0%	0.2%
	High	31.7%	29.3%	8.2%	5.7%	5.3%	0.4%
Means for trade groups	Low	7.6%	7.3%	0.3%	2.1%	2.5%	0.0%
	Medium	14.6%	14.4%	1.5%	3.6%	3.7%	0.1%
	High	36.5%	33.9%	10.0%	6.1%	5.6%	0.5%
Means for σ_r groups	Low	13.7%	12.6%	2.8%	2.0%	2.1%	0.0%
	Medium	21.2%	20.2%	4.0%	3.9%	3.9%	0.2%
	High	23.8%	22.8%	5.1%	5.8%	5.7%	0.5%

Figure 1. Number of Trades

The sample is the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. The figures is based on Island ITCH and Nasdaq Nastraq data.

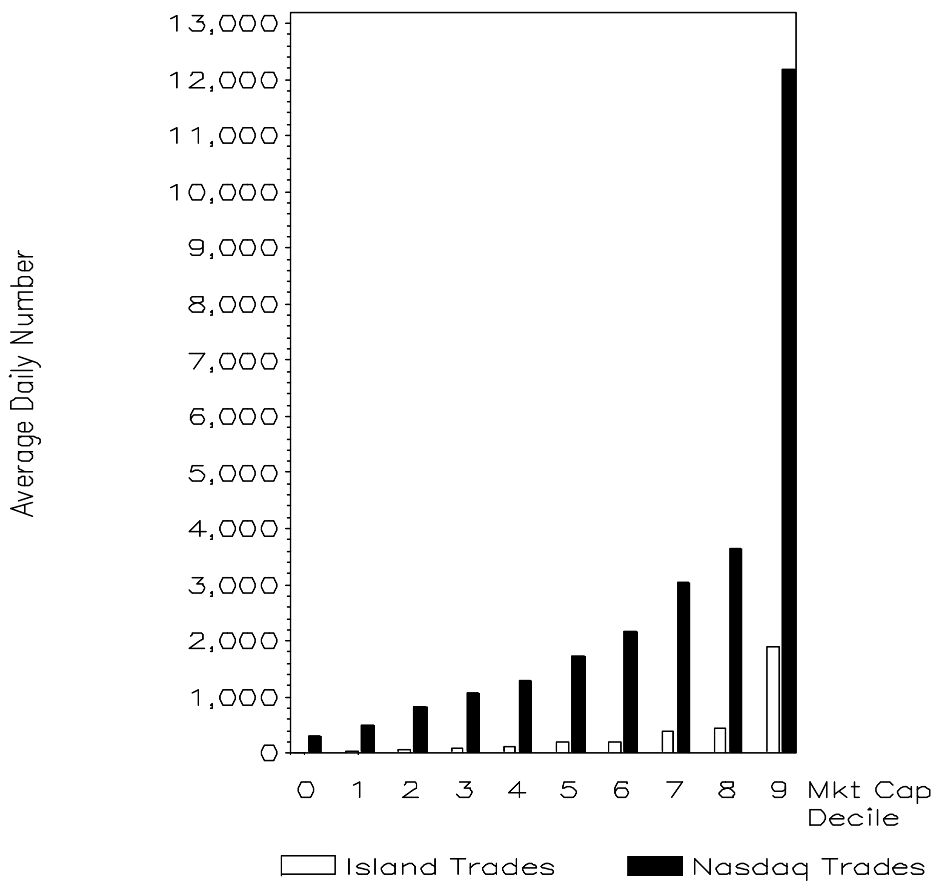

Figure 2. Island market and limit orders

The sample is the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. The figures is based on Island ITCH data. .

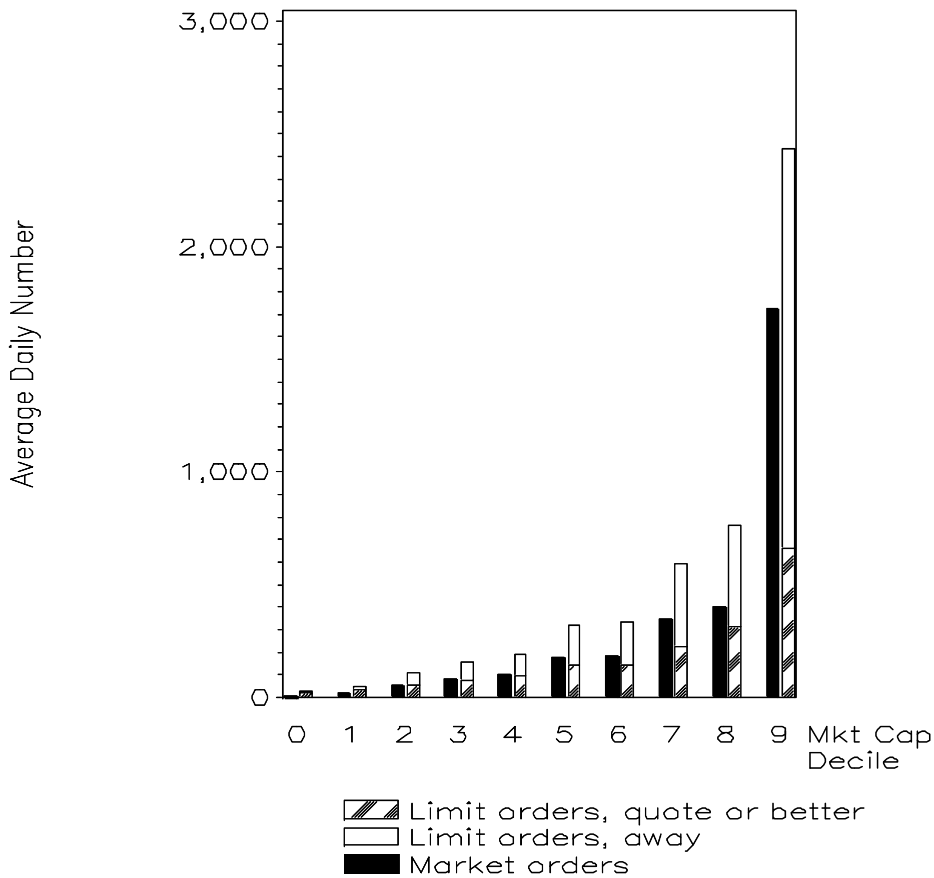

Figure 3. Ask Depth for Microsoft

Cumulative dollar depth on the ask side is recorded for Microsoft at five-minute intervals during regular trading hours from October 1, 1999 through December 31, 1999. Depth is measured in intervals of width $1/16. The highest group ("$5. 00") also includes offerings priced above $5. 00. At each interval, the line in the middle of the box indicates the median; the bottom and top of the box indicate 0. 25 and 0. 75 quantiles; and, the ends of the lines indicate 0. 01 and 0. 99 quantiles. The dark line depicts a smoothed fit through all of the points.

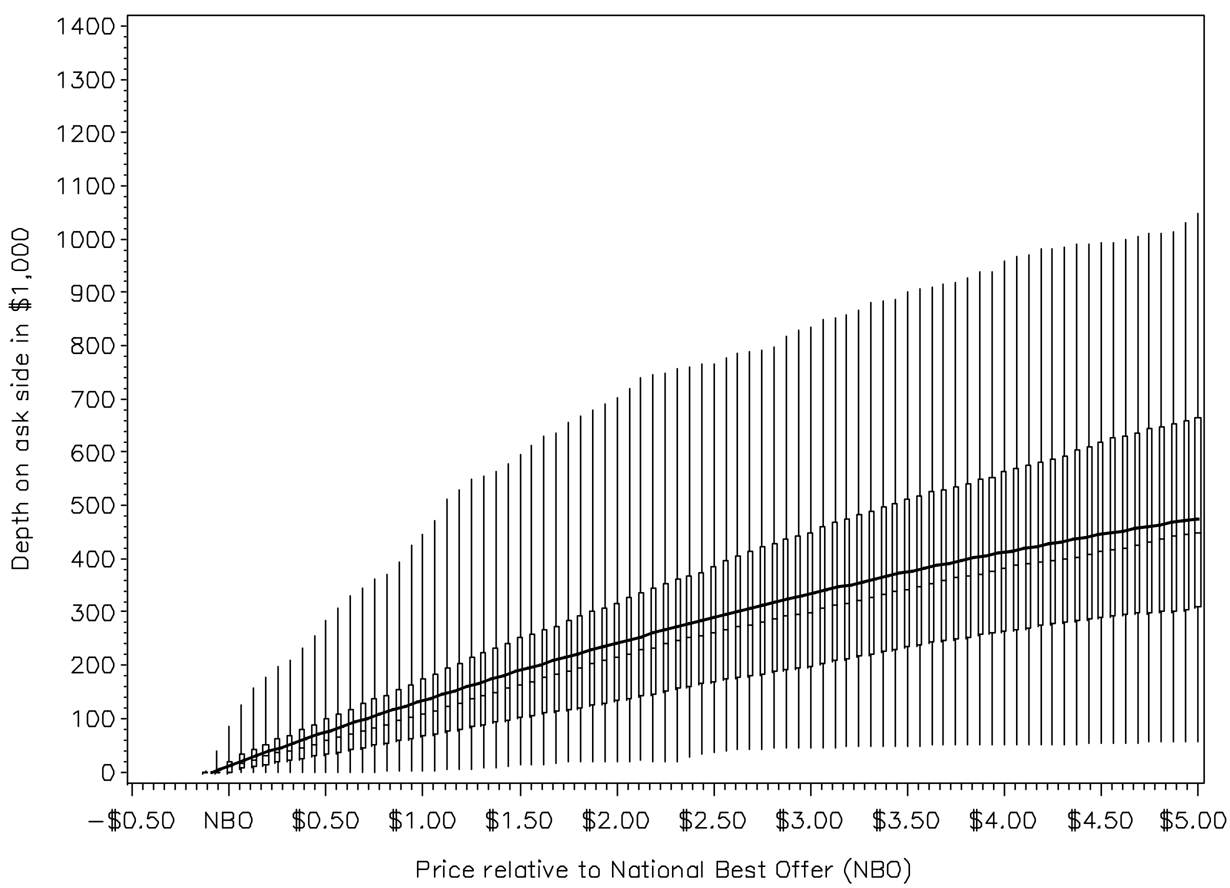

Figure 4. Executions and cancellations over time

The sample is the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. Failure functions (cumulative probabilities of occurrence) for executions and cancellations of limit orders over time, estimated with the Kaplan-Meier correction for censoring. In estimating the function for execution, cancellation was treated as equivalent to censoring. In estimating the function for cancellation, execution was treated as equivalent to censoring.

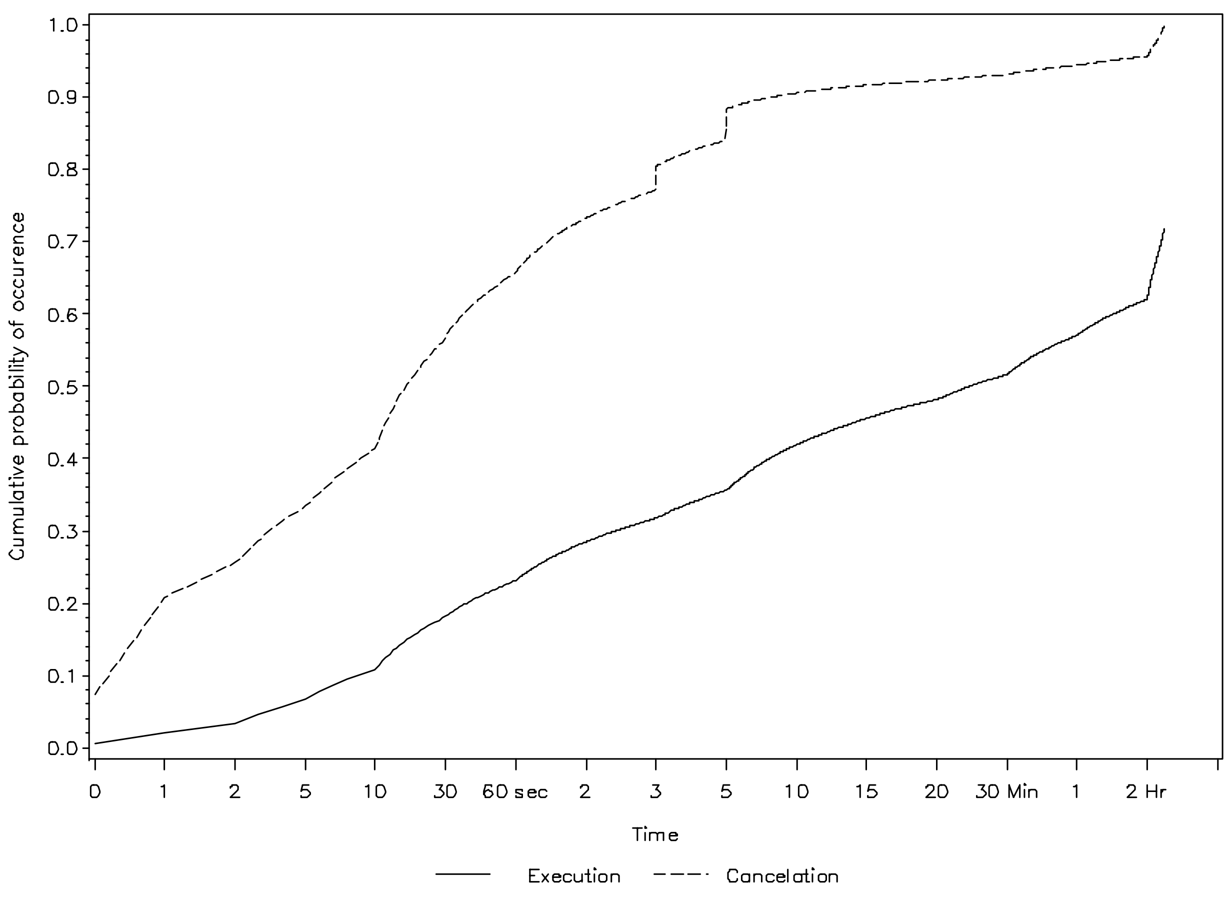

Figure 5. Nasdaq and Island Volume

The sample is the 300 largest firms in the Nasdaq National Market (ranked by equity capitalization as of September 30, 1999), over the 64 trading days in the fourth quarter of 1999. The figures is based on Island ITCH data and the Nastraq database.