Technology and Liquidity Provision: The Blurring of Traditional Definitions

Joel Hasbrouck and Gideon Saar

This version: November 3, 2004

Joel Hasbrouck and Gideon Saar are from the Stern School of Business, New York University, 44 West Fourth Street, New York, NY 10012 (Hasbrouck: suite 9-86, 212-998-0310, jhasbrou@stern.nyu.edu; Saar: suite 9-93, 212-998-0318, 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.

Technology and Liquidity Provision:
The Blurring of Traditional Definitions

Abstract

The usual economic perspective on a limit order emphasizes its role in supplying liquidity. We investigate the trading of 300 Nasdaq-listed stocks on the Island ECN, an electronic communication network organized as a limit order book. We find that a substantial portion of the limit orders are cancelled within an extremely brief time. We term "fleeting orders" those limit order that are cancelled within two seconds of submission, and explore the role they play in trading strategies. Our principal finding is that fleeting limit orders are closer substitutes for market orders than for traditional limit orders. Our results suggest that the aim of a trader who submits a fleeting order is to demand immediacy. This contrasts with the traditional view of limit order traders as patient providers of liquidity. We hypothesize that a "new equilibrium" has arisen, driven by improved technology, the emergence of an active trading culture, and increased market fragmentation. The new environment transforms the market from one in which prices are posted (visible limit orders) into one where searches (for hidden liquidity) are needed in order to achieve better terms of trade.

1. Introduction

The usual economic perspective on a limit order emphasizes its role in supplying liquidity. In this capacity, it is often viewed as extending to the market a visible, ongoing and persistent option to trade. Unlike a market order, it is passive and patient.

This characterization of a limit order arises most naturally from the perspective that a customer limit order is functionally equivalent to a dealer quote. Models that investigate dealer behavior and extensions of these models to the limit order book context generally exhibit three features.[1] First, dealers/limit order traders are indifferent to whether or not their order is hit. Second, they have no cause to cancel or modify their orders except in response to a trade. A limit order, in other words, persists until there is a trade. Third, the models feature a clean dichotomy between liquidity suppliers (who enter limit orders) and liquidity demanders (who use market orders).

The strategic order choice models suggest an alternative characterization.[2] Although the outcome of a trader's decision in these models is either a market order or a limit order, this choice usually depends in part on the trader's impatience or private value as well as the current state of the book. Thus, traders in these models cannot be identified a priori as liquidity suppliers or liquidity demanders. Under certain circumstances (a wide spread, for example), even a trader with a strong desire to trade might submit a limit order. Unlike the dealer models, traders who submit limit orders are not indifferent to execution. They unambiguously prefer that their bids or offers be hit. Still, these models share with the dealer models the notion that limit orders are submitted in order to remain in the book and have a chance to interact with incoming market orders.[3]

---

[1] For dealer models see, for example, Copeland and Galai (1983), Glosten and Milgrom (1985), and Easley and O'Hara (1987). Models with a similar perspective in the context of a limit order book include Glosten (1994) and Seppi (1997).

[2] See Cohen, Maier, Schwartz, and Whitcomb (1981), Angel (1994), Chakravarty and Holden (1995), Harris (1998), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2001), Goettler, Parlour, and Rajan (2003), Kaniel and Liu (2004), and Rosu (2004).

[3] Models that discuss a dynamic trading strategy (Harris (1998), Large (2004), and Rosu (2004)) suggest a pecking order whereby limit orders supply liquidity until such a time as traders revise their expectations with respect to the probability of execution and move to use more aggressive orders. See also the

The evidence presented in this paper calls into question the traditional view of limit orders as suppliers of liquidity. We argue that it is no longer the case that traders who demand immediacy use market orders and the more patient traders use limit orders. Instead, limit orders are heavily used to demand immediacy in the marketplace. Recognizing these changing roles is important for both theoretical modeling of traders in financial markets and empirical evaluation of the quality of trading venues.

To shed light on these issues we investigate the trading of 300 Nasdaq-listed stocks on the Island ECN, an electronic communication network organized as a limit order book.[4] On average, non-marketable limit orders account for 83% of all incoming orders (89.2% in terms of shares). However, the fill rate of limit orders on Island is remarkably low: only 18.4% of the limit orders get partial or full execution, representing 12.6% of the shares in submitted limit orders. A closer inspection reveals that this situation is driven by a substantial portion of the limit orders that are cancelled within an extremely brief time: 27.7% of the limit orders (32.5% in terms of shares) are cancelled within two seconds of submission. We term these "fleeting orders" and explore the role they play in trading strategies. Their sheer numbers and apparent defiance of easy classification in the usual framework of patient limit orders and impatient market orders poses a puzzle and a challenge to academic theories in this area.

We find that fleeting orders seem to be priced more aggressively than other limit orders. For example, fleeting buy orders are typically priced above the bid price but below the ask price (and hence they are non-marketable). On average, about 84% of fleeting limit orders are priced better than the prevailing Island bid or offer. What might lead a trader to submit such a brief aggressively priced order? One possibility is to achieve an execution against hidden depth. Like many other markets organized as limit order books, Island allows its traders to submit hidden orders that sit in the book but are

---

experimental evidence in Bloomfield, O'Hara, and Saar (2004) on the dynamic trading strategies of informed and uninformed traders.

[4] The Island limit order book was recently merged with that of Instinet. The merged entity, INET, has approximately 25% market share in Nasdaq-listed equity trading.

not visible to traders. We find that about 14% of the executions on Island take place against hidden depth.

To investigate the economic role of fleeting orders we carry out both a cross-sectional and a dynamic multinomial logit analyses. Our principal finding is that fleeting orders are closer substitutes for market orders than for traditional limit orders. For example, a rapid move of prices in one direction leads to a decrease in the propensity to use regular limit orders but an increase in usage of fleeting and marketable orders, both of which testify to more aggressive intentions that aim at effecting an execution. In a similar vein, a larger prevailing NBBO spread is associated with a higher propensity to submit regular limit orders and fewer fleeting and marketable orders.

Our findings suggest that the aim of a trader who submits a fleeting order is to demand immediacy. This contrasts with the traditional view of limit order traders as patient providers of liquidity. A trader who wishes to get a quick execution could send a marketable order, but could also send a regular limit order priced inside the quote to search for hidden liquidity. If such hidden depth inside the quote exists, we observe an execution. If there is no such depth, the trader cancels his order almost immediately and either submits a marketable order or looks for a counterparty on other ECNs or among Nasdaq dealers.

Why would traders use visible orders to explore hidden liquidity when they can use hidden orders for that purpose without disclosing their trading intentions? We hypothesize the emergence of a "new equilibrium" in trading strategies driven by (i) changes in technology that enable rapid submission and cancellation of orders, (ii) the evolution of an active trading culture, and (iii) fragmentation (the co-existence of multiple trading venues). This new environment transforms the market from being one in which prices are posted (visible limit orders) into one where search is required to achieve better terms of trade.

At the core of the new equilibrium are still two types of traders, patient and impatient. The patient traders consider posting limit orders but face costs imposed by the risks associated with order exposure (e.g., leakage of their private information or front running of their orders by other traders). The patient traders balance these costs against the benefit of attracting counterparts by publicizing their trading intentions. Using a hidden order can protect them somewhat from the risks associated with order exposure, but may cause impatient traders to look for liquidity on a different trading venue if they are not aware of the hidden depth.

The impatient traders need to balance their need for immediacy with their willingness to incur a high price impact. It is their optimal choice that changes most noticeably with advanced technology and a commitment to active trading. If submission and cancellation of limit orders are made easy (or even completely automated), impatient traders would be willing to engage in a search for hidden liquidity. Presumably, the new capabilities lower the cost of the search (in terms of their desire for immediacy that can be viewed as a discount rate) and at the same time provide for the opportunity to achieve better prices.

A new equilibrium therefore emerges where patient investors use hidden orders to supply liquidity to the book and impatient traders use limit orders priced inside the quote to search for the hidden liquidity. The impatient traders bear the cost of the search, but the larger supply of shares at better prices compensate them for it, and they are better off. The visibility of the fleeting order comes into play at this point—it serves as a signaling device that sustains the new equilibrium. The visible fleeting orders signal to patient traders that it is worthwhile for them to supply hidden liquidity by demonstrating that enough impatient traders are willing to search for it on this particular venue.

Our finding of a substantial number of visible fleeting orders inside the quote is consistent with the hypothesized new equilibrium. Furthermore, our multinomial logit analysis suggests that these orders arise from a desire to achieve immediacy on the part of impatient traders. Since patient traders can always choose to make their limit orders visible, the finding of a substantial portion of executions against hidden depth in the book suggests that some patient traders are better off submitting hidden orders. Together, these results are consistent with the behavior of traders we describe as the new equilibrium.

Our findings on the use of limit orders to demand immediacy and the hypothesized new equilibrium on the Island ECN have some important implications. First, they call into question results from theoretical models that characterize limit orders as persistent and their traders as patient. The new trading environment we observe requires a different framework for thinking about optimal order choices in markets. Second, our results challenge the manner in which execution quality of trading venues is evaluated. The Security and Exchange Commission's rule 11Ac1-5 requires market centers to report several measures in order to help investors figure out where to send their orders. One of the measures required by the SEC is the fill rate of limit orders. Presumably a higher fill rate of limit order testifies to a better market. The fill rate we document on Island is low, but Island is still the market of choice for many active and sophisticated traders. In the equilibrium we describe, the fill rate is a misleading and inappropriate metric of quality. We believe that recognizing the new ways in which trading and order choices have changed due to technology, active trading, and fragmentation is important to academics, regulators, and investors.

The rest of this paper is organized as follows. The next section provides a literature review, and Section 3 discusses our sample and the Island ECN. Section 4 provides an initial characterization of fleeting orders and hidden executions. The next two sections consider multinomial logit models of order categories. We consider separately cross-sectional variation (Section 5) and dynamic variation (Section 6) in the order mix. Section 7 discusses the developments that have led to broader use of fleeting orders and provides a historical perspective. A brief summary concludes the paper in Section 8.

2. Literature Review

The notion that limit orders supply liquidity to the market make them similar in nature to a dealers' quotes, and suggests that the economic forces affecting limit order strategies should be similar to those investigated in models of dealer markets. Dealers in the sequential trade models of asymmetric information are risk-neutral. They are subject to adverse selection, and the pricing of their bids and offers is ultimately determined by zero-expected profit conditions induced by competition (e.g., Copeland and Galai (1983), Glosten and Milgrom (1985), Easley and O'Hara (1987)). Some extensions of these models differentiate between dealers and limit order traders on the point that the latter cannot condition on the size of the incoming trade. The shape of the book under these circumstances is considered theoretically in Glosten (1994) and empirically in Sandas (2001). Seppi (1997) examines the interaction between a dealer and a limit order book.

Risk neutrality of dealers or limit order traders in the aforementioned models makes them indifferent to whether or not their quotes or orders are hit (although they may have preferences concerning the total size of the order that triggers the execution).[5] The sequential trade models also feature the rational expectations notion of “regret-free” prices, and therefore limit orders or quotes are changed only in response to new trades (or, as in Easley and O'Hara (1992), a period without trading). Since the interaction in a dealer market clearly distinguishes between liquidity suppliers and demanders, models of limit order books in this tradition also specify one class of traders who supply liquidity using limit orders and another one who demands liquidity.

Models that investigate the choice of traders between market and limit orders offer a different perspective. Cohen, Maier, Schwartz, and Whitcomb (1981), Angel (1994), and Harris (1998) focus on the trader’s choice of order type. This choice usually depends on expected limit order execution probabilities, which in turn depend on the order choices of other traders. The focus on execution probabilities and their importance in traders’ decisions motivates the equilibrium models of Chakravarty and Holden (1995), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2001), Goettler, Parlour, and Rajan (2003), Kaniel and Liu (2004), and Rosu (2004). The trader’s market/limit order decision in these models is affected by his degree of impatience, his private valuation of the asset, and the state of the book. These models come closer to the

[5] In many of these models the indifference to execution is mostly a consequence of assumed risk-neutrality. Risk-neutrality, however, is not essential to this result. The dealer in Stoll (1978) sets his bid to reflect the loss in expected utility in the suboptimal portfolio that will result if his bid is hit. At the optimum, however, expected utility conditional on the bid being hit is equal to that conditional on no trade. This implies an indifference to the execution.

true interaction in double auctions organized as limit order markets where each trader chooses between demanding and supplying liquidity.

The equilibrium order choice models do not for the most part attach importance to a limit order's duration. Typically, a randomly-drawn trader arrives at each instant and makes a choice between market and limit orders without the possibility of a subsequent trading opportunity. Order strategies are defined by type (market or limit), and (for a limit order) the price. Individual strategies balance the lower trading costs of a limit order execution against the costs of delay and/or non-execution. The duration of an order's exposure is not a key facet of these models because cancellation is non-strategic.[6]

A noteworthy feature of these models, however, is that a trader's order choice influences the choices of subsequently arriving traders. For example, the probability that the next trader will use a market order increases if the current trader enters a limit order. Limit orders in these models are therefore somewhat more active than the passive limit orders discussed above in connection with the dealer perspective.

Three papers in particular model dynamic strategies and enable the trader to cancel an order and resubmit a different one to actively seek an execution. Harris (1998) considers a trader trying to minimize the purchase price of a predetermined quantity, subject to a deadline. The optimal strategy is to initially place a limit order, then to reprice the order more aggressively as the deadline nears, and finally if necessary use a market order. That is, limit orders are entered and revised pre-deadline even by agents who are ultimately constrained to trade. Bloomfield, O'Hara, and Saar (2004) provide evidence confirming the utilization of these strategies by constrained liquidity traders in experimental settings. They also show that traders with private information about the (common) value of the security would tend to start trading using market orders but shift to using limit orders as prices adjust to reflect their private information.

[6] Orders expire in one period in the Foucault and Parlour analyses; they never expire in Foucault, Kadan and Kandel; in Goettler, Parlour, and Rajan they face random cancellation (with probabilities depending on the price path).

Large (2004) suggests that limit orders cancellations arise from the refinement (over time) of a limit order trader's beliefs about the arrival rate of market orders, which is directly related to the expected time until the order's execution. Rosu (2004) proposes a model in which traders can update (cancel and resubmit) existing limit orders instantaneously. By assumption, however, impatient traders only use market orders. While Rosu discusses "fleeting" limit orders, in his model a fleeting limit order is a limit order priced in such a way as to attract an incoming market order that executes it. Although one could view such a limit order as a means for demanding immediacy, our definition of fleeting orders differs from the one discussed in his paper. We define fleeting orders as those limit orders that get canceled quickly, while Rosu's fleeting orders are those that get executed quickly.

The strategies discussed to this point are set in the context of a single execution venue. Fragmentation may increase the cost of exposing a limit order. Competing traders can use other venues to price-match the order, reducing its probability of execution (since there is no time priority across venues). Shortening the order's exposure time may be a way of controlling these costs. Also, sequential strategies involving fleeting orders may be used across venues in a fragmented market. There has been to our knowledge no theoretical work on this problem that is specific to securities markets (although more broadly this could be considered a search problem).

A number of empirical studies have sought to characterize limit order markets (e.g., Hamao and Hasbrouck (1995), Biais, Hillion, and Spatt (1995), Ahn, Bae, and Chan (2001), Biais, Bisiere, and Spatt (2003), and Hollifield (2004)). Only a few studies model limit order durations. Cho and Nelling (2000) and Lo, MacKinlay, and Zhang (2002) estimate duration models, but their focus is on execution, with cancellation being taken as an exogenous censoring process. We use multinomial logit specifications to characterize order strategies. This approach is similar to that of Smith (2000), Ellul, Holden, Jain, and Jennings (2002), and Renaldo (2004). Our event classification, however, will involve outcomes as well as submission decisions.

Other studies have focused specifically on Island. Hasbrouck and Saar (2002) characterize the cross-sectional relation between volatility and Island characteristics. Hansch (2003) documents the extent of and variation in Island book depth. Nguyen, Van Ness, and Van Ness (2003) describe changes related to Island's decision to shift trade reporting from Nasdaq to the Cincinnati (now National) Stock Exchange. Hendershott and Jones (2004) consider changes in Island's market quality in exchange-traded funds when, for regulatory reasons, it did not disseminate quotes or book data.

3. Sample and data

a. Sample construction and descriptive statistics

The sample is drawn from among all Nasdaq National Market common stocks with data in the CRSP database from October 1 to December 31, 1999. The sample is the 300 largest firms based on equity market capitalization as of September 30, 1999.[7]

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 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 medians for three groups constructed by ranking on market capitalization, average number of daily trades (as a measure of trading activity), and standard deviation of daily returns (as a measure of volatility).

[7] 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.

b. The Island ECN

The Island ECN began operating in 1997. In 2002 Island merged with its largest competitor, Instinet, forming a combined entity named INET. Our sample period is the fourth quarter of 1999, and we will therefore use the name "Island". Island operates a pure agency market. All orders must be priced. A trader who seeks immediate execution must price the limit order to be marketable, e.g., a buy order priced at or above the current ask price. For all intent and purposes, a marketable limit order in a pure limit order book is equivalent to a market order in floor or dealer markets. Such an order is never displayed in the book; rather, it is immediately executed upon arrival to the system. We use the term market orders and marketable (limit) orders interchangeably in this paper.

Orders may be visible or hidden.[8] Execution priority follows price, visibility and time. All visible quantities at a price are executed before any hidden quantities are executed. During our sample period, Island's best (visible) bid and offer were incorporated into Nasdaq's montage, which determined the National Best Bid/Offer (NBBO) display, and all trades were reported to Nasdaq. Island's fee structure during our sample period differentiated between orders in the book and marketable orders that executed against the book. Upon execution, the former received a 0.1 cent per share rebate and the latter paid a 0.25 cent per share access fee.

Over our sample period, Island executions comprised about 11% of Nasdaq trades and 6% of Nasdaq's traded volume, with the disparity reflecting the small size of most Island trades. Furthermore, during this period, Island's presence was concentrated in the subset of very active Nasdaq stocks. Island's market share for the average stock in our sample is 6.23% in terms of trades and 3.52% in terms of share volume.

[8] The option of complete invisibility differentiates Island's limit orders from the reserve ("iceberg") orders found in Euronext, where at least a portion of the limit order must be visible at all times.

c. Island data and statistics

The Island data we use are identical to those supplied in real time to Island subscribers. These data are comprised of time-sequenced messages that completely describe the history of trade and book activity. The process may be summarized as follows. When an arriving order is marketable, i.e., it can be matched (in whole or part) against existing orders in the book, the system sends an Order Execution message. If the order can't be matched, i.e., it is a "regular" limit order in our terminology, 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.[9] 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 hidden orders, may be constructed by cumulating these messages from the start of the day onwards. Although the arrival time and quantity of a hidden order are never made available, the execution of such an order is signaled by a special trade message.

In presenting statistics based on the Island data, we take the firm as the unit of observation. That is, we first compute estimates for each firm, and then report summary statistics across firms. Table 2 presents summary statistics on the number and size of orders that arrive to Island.[10] The average number of daily limit orders increases with market capitalization (in the ranked group means), trading activity, and volatility. 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 trading activity groups, which may suggest that retail activity is more concentrated in the largest, most active Nasdaq stocks. Market orders (i.e., marketable orders that are executed immediately upon reaching Island without entering into the book) tend to be

[9] Orders with subscriber-only display condition are still visible to anyone with access to the Island book (all Island subscribers). The only difference from orders without this display condition is that even if these orders constitute the best prices on Island they are not incorporated into Nasdaq's NBBO.

[10] We only consider data from the regular trading session of the Nasdaq Stock Market (from 9:30 a.m. to 4:00 p.m.).

smaller than limit orders, with a mean of only 335 shares. As with limit orders, the average size of market orders decreases with market capitalization and trading activity.

Table 3 describes the incoming order mix and execution proportions of limit orders (fill rate). On average limit orders account for 83.0% of the incoming orders, or 89.2% of the shares in incoming orders (the rest are marketable orders). Of these (non-marketable) limit orders, 18.4% are at least partially filled, but only 12.6% of the shares in these orders are executed. As we indicate in Section 7, this fill rate is very low relative to estimates based on samples from ten or fifteen years ago. We believe that the current low fill rate reflects the new ways in which traders utilize limit orders on Island. To investigate this issue we now turn to the phenomenon of fleeting orders.

4. Fleeting orders and hidden executions: an initial characterization

To motivate our definition of fleeting orders, it is useful to consider the timing of cancellations. Let τ be the time that elapses between an order's submission and its cancellation. The survival function at time t is $S(t) = \Pr(\tau > t)$. These functions are estimated for each firm using the life-table method and taking execution as the censoring event. Figure 1 plots the cross-firm means. The time scale is nonlinear to show detail for shorter times.

Most strikingly, a large number of limit orders are canceled very shortly after their submission. $S(3)$, the probability of surviving to at least three seconds is 27%. The data are time-stamped to the second, so this implies that 27% of the orders are canceled in two seconds or less. Roughly 42% are canceled in ten seconds or less. For completeness, the figure also depicts the mean survival probability of an execution, taking cancellation as the censoring event. Execution is clearly the less probable event, particularly in the few seconds immediately after submission.

The large proportion of cancellations at short durations motivates consideration of these orders as a separate category. We use, somewhat arbitrarily, two seconds as the break-point: an order that is cancelled in two seconds or less is defined as “fleeting”.[11] Why would traders cancel so quickly such a large percentage of the limit orders they submit? In the introduction we mentioned the “new equilibrium” where patient traders supply liquidity using hidden orders and impatient traders search for that liquidity using fleeting orders. When they fail to execute immediately against hidden depth, fleeting orders are left in the book for two seconds to signal to patient traders that there is interest in hidden liquidity (to sustain the equilibrium). This, in contrast to the usual characterization of non-marketable limit orders as patient providers of liquidity.

For the purpose of our empirical investigation, our event classifications are defined as follows. An execution against a displayed quantity is termed a “market order”; an execution against a non-displayed quantity is a “hidden execution”. Visible orders added to the book represent limit orders. If they are cancelled within two seconds or less, they are “fleeting”; others are “regular limit orders”. The state of a limit order two seconds subsequent to its arrival therefore determines its event category.

Classification in this scheme therefore depends in part on outcome as well as order choice. This differs from the practice in Smith (2000), Ellul, Holden, Jain, and Jennings (2002), and Renaldo (2004), wherein events are defined solely by reference to the order and market condition at the time the order arrived. Since we are attempting to gain insight into order strategy, there is a risk of spurious inference arising from an event classification that also depends on outcome. Our classification scheme, however, should understate the number of fleeting orders. For example, a trader might submit a non-marketable limit order with a two-second time in force (an intended fleeting order). But if the order should execute one second after being posted to the book, it would be classified as a regular limit order.

[11] Since fleeting orders are characterized by the speed with which they are cancelled, it is useful to describe how such cancellation might occur. On Island (and in most limit order markets) a limit order can be cancelled by prearranged conditions set when the order is submitted, most importantly the time in force (TIF) attribute. Alternatively, a trader can continuously monitor, and enter a cancellation request in response to market conditions. Ideally we would like to know the intended time in force (TIF) of the order, i.e., the value actually submitted with the order or the value that has been programmed into the trader’s order management system. Our data do not contain this, however, and our inferences must therefore be based on the time the order was actually in the book.

In the classification schemes that do not depend on outcome, cancellation of an order is treated as a separate event, disconnected from (indeed, statistically independent of) the submission of the order. When this connection is suppressed, no conclusions can be drawn about the order's intended duration. We view duration as an important strategic variable, and therefore employ in this paper a classification scheme that utilizes outcome to investigate fleeting orders. We attempt to minimize misclassification and spurious inference by restricting the classification to an extremely brief outcome window (two seconds), and when appropriate and necessary, qualifying our conclusions.

Although our event classification is a complete partition of observable orders and outcomes (within two seconds), there is one event that is impossible in some circumstances. A market order can only occur when the book has visible depth on the opposite side. For many of the smaller firms in our sample this is frequently not the case. We therefore partition our sample depending on existence of visible opposing depth. Note that hidden executions are a possibility even in the absence of any visible depth.[12]

More specifically, each order is put into one of two subsamples for a particular firm: (i) orders that arrive to Island when there is opposing visible depth, and (ii) orders that arrive to Island when there is absence of opposing visible depth on the Island book. We carry out all analysis separately for the orders in these two subsamples. However, if there are fewer than 500 orders for a certain firm in one subsample, say, those orders arriving when there is no opposing depth, this firm is dropped from the multiple-firm analysis of orders in that subsample.

Table 4 summarizes the distribution of order frequencies across firms. Panel A of the table is based on the 279 firms that had at least 500 orders in the visible-opposing-depth subsample. The most frequent event is a regular limit order (58 percent). Fleeting limit orders are next (21 percent) followed by market orders (18 percent). Panel B of Table 4 summarizes outcome frequencies for the 101 firms that had at least 500 orders

[12] A similar consideration arises when cancellation is modeled as a separate event. Care must be taken to ensure that the statistical model does not imply a positive probability for order cancellation even when the book is empty.

during times when a market order would not have been feasible.[13] In this circumstance, regular limit orders are used about two-thirds of the time (on average), and fleeting limit orders are used about one third of the time. Hidden executions are quite rare in Panel B, suggesting that when there is no visible depth there isn't likely to be much hidden depth either. Relative to the frequencies in Panel A, there appears to be a shift or substitution from market orders to fleeting limit orders. This is consistent with the intuition that when no opposing depth exists, strategies that would otherwise have employed market orders utilize fleeting orders instead.

Table 4 also reports median order frequencies across the capitalization-, activity-, and volatility-ranked groups. There are no clear patterns in Panel B (no opposing depth), but this may reflect a lack of precision due to the small numbers of firms. The subsample with opposing depth contains substantially more firms. In this subsample, most strikingly, the frequencies of regular limit orders are relatively flat across the market capitalization groups. Noticeable changes do occur in the frequencies of fleeting and market orders, however, and they tend to be in offsetting directions. For capitalization-, activity-, and volatility-ranked groups, moving from low to high values is associated with a monotonic decrease in the frequency of fleeting limit orders and a monotonic increase in the frequency of market orders. The relative constancy of the regular limit order frequency and the offsetting changes in fleeting and market orders suggest that fleeting limit orders are not equivalent to regular limit orders, and that they may be closer substitutes for market orders.

Table 5 describes the pricing of limit orders relative to Island's quote. Both regular and fleeting limit orders tend to be priced ahead of (more aggressively than) the bid or offer. This tendency is more pronounced for fleeting limit orders (an average of 84% priced ahead, median 87%) than for regular limit orders (72% priced ahead, median

[13] There is some overlap. A number of firms had at least 500 orders when there was no opposing depth and at least 500 orders when there was opposing depth. Thus, some of the differences between the depth and no-depth subsamples are due to dynamic variation in market conditions for a single firm. Some of the differences are also cross-sectional, however. The smallest firms in the study rarely had 500 orders when visible opposing depth existed; the largest firms rarely had 500 orders when there was no depth.

73%). If pricing aggressiveness reflects urgency, this suggests that fleeting orders are more urgent. Such a large percentage of fleeting orders priced inside the quote is also consistent with the idea that they are used to “fish” for better-priced hidden liquidity in the book.

Table 6 provides another perspective on hidden executions and fleeting orders. First, although hidden executions constitute a small proportion of all events, they are substantial relative to all executions. The first two columns of the table show that on average, they account for 13.7% of all executions, and 11.8% of all executed shares. The last two columns of the table examine fleeting orders relative to all visible orders. By frequency, fleeting orders constitute on average 27.7% of all visible orders. By shares, this proportion is on average 32.5%, indicating that fleeting orders tend to be slightly larger.

Table 6 also shows that there is a pronounced monotone pattern in fleeting orders and hidden executions across the capitalization-, activity-, and volatility-ranked groups. Larger, more active, and more volatile stocks tend to have more hidden executions and fewer fleeting orders. The patterns highlight the relationship between these two events. Say a trader demands immediacy by sending an order inside the quote aimed at executing against hidden depth. The order will be categorized as a hidden execution if there was hidden depth in the book. If there was no hidden depth, however, the order will be posted to the book but quickly cancelled (so that the trader can demand immediacy by submitting a marketable order or seek out another pool of liquidity on another trading venue) and will be categorized as fleeting. Therefore, more hidden executions are associated with fewer fleeting orders in the cross-section of stocks. The table suggests that there is more hidden depth in the book of a larger, more active, or more volatile stock.

5. Cross-sectional multinomial logit analysis

The subsample tabulations discussed above suggest that the cross-sectional variation in fleeting order frequency differs markedly from that of regular limit orders. This section examines these effects in a multivariate analysis.

The econometric approach is multinomial logit on the categories defined in the previous section: {*Regular limit order, Fleeting limit order, Market order, Hidden execution*}. Let $i = 0,1,2,3$ denote an index corresponding to these events, and let *j* index firms. The probability of event *i* for firm *j*:

$$\log\left(\frac{\Pr_{i,j}}{\Pr_{0,j}}\right) = X_j \beta_i \text{ for } i = 1,2,3$$

where X_j is a vector of explanatory variables and β_i a vector of coefficients. Event $i = 0$, (the "regular limit order" event) is taken as the reference event: the occurrence probabilities for the other events are modeled relative to the probability of this event.

The explanatory variables include an intercept, daily return volatility, log average price, log market capitalization, median daily turnover, and the Herfindahl-Hirschman index (HHI) of trading volume concentration (computed over all Nasdaq market makers and ECNs). The importance of volatility as a cross-sectional determinant of order mix is suggested by Foucault (1999). Hasbrouck and Saar (2002) describe other effects of volatility on preferences for limit orders. The next three variables are included to control for ancillary and incidental effects. Among other things, capitalization may be related to investor characteristics and frequency of information events. 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. HHI is included as a proxy for market fragmentation.

One order constitutes one observation in our analysis, and orders are assumed to be independent.[14] To achieve equal weighting for all firms, each observation is weighted by the inverse of the total number of events for the firm. Given the large number of observations, virtually all of the logit coefficients are statistically significant at the usual levels. To facilitate interpretation, we compute the event probabilities for a representative firm, i.e., one for which all explanatory variables are equal to their sample means, $X_j = \bar{X}$. This is considered the "base" case. We then examine the probabilities implied by the model when each of the variables, taken one at a time, increases by one standard deviation. These calculations are reported in Table 7. The table reports probabilities for the base case and the case where the variables increase by one standard deviation, as well as differences in probabilities between the two cases.[15] The table reports only the results corresponding to the subsample with available depth on the opposite side of the order. The reason is that we are interested in specifically examining the relation between limit, fleeting, and market orders. The analysis without opposing depth does not have market orders, making all changes in limit orders mirrored by opposing changes in fleeting orders.[16]

Table 7 shows that the frequency of regular limit orders does not generally change when we consider increases of one standard deviation in the independent variables. The largest differences tend to be offsetting changes between market order and fleeting limit order frequencies. In particular, an increase in volatility, capitalization or turnover is associated with decreased probability of a fleeting limit order. This decreased probability is mostly offset by an increase in the probability of a market order. These effects are consistent with the ranked groups results presented in Table 4.

The results pertaining to volatility are particularly noteworthy. Existing theoretical studies offer various predictions concerning the effects of volatility. In the model of Foucault (1999) volatility is positively related to the pick-off risk faced by the

[14] Independence across firms would be violated by common factors in order flows.
[15] The base probabilities approximate, but do not exactly equal, the mean frequencies reported in Table 4. The differences arise because the logit probabilities are nonlinear functions of the explanatory variables.
[16] This analysis is available from the authors upon request.

limit order submitter. This effect militates against limit orders. Traders compensate, however, not by entering fewer limit orders, but by pricing their limit orders less aggressively. In equilibrium, the wider spread makes market orders more costly, leading to a higher proportion of limit orders (and a lower fill rate).

Handa and Schwartz (1996) model limit order execution as the first passage of a random walk (the stock price) hitting a barrier (the limit order). With zero-drift, the expected first passage time (execution duration) is negatively related to volatility. Lo, MacKinlay, and Zhang (2002) find this negative relation in a sample of NYSE limit orders, although they note that the random-walk model does not give accurate point estimates. With respect to the brevity of execution durations, therefore, an increase in volatility should make limit orders more attractive.

In Table 3 we presented limit order submission proportions across firm groups ranked by volatility. As volatility increases, this proportion declines, an apparent refutation of Foucault's prediction. The estimates in Table 7, however, suggest that while increased volatility induces a lower frequency of limit orders, the drop occurs entirely in fleeting limit orders. The frequency of regular limit orders actually increases slightly. This pattern is also evident in the volatility-ranked groups frequencies reported in Table 4. With respect to regular limit orders, therefore, the volatility results do not contradict (nor do they support) the prediction of Foucault (1999). Higher volatility is presumably associated with greater presence of active traders and therefore increased use of (the more sophisticated) hidden orders.[17] As the likelihood of finding hidden depth inside the quote increases, there is an increase in the frequency of hidden executions and a decrease in the frequency of fleeting orders.

Table 7 also shows that an increase in concentration of trading activity (HHI), i.e., lower fragmentation, implies a small decrease in the usage of fleeting orders. Since HHI takes into account all execution venues (including all alternative trading systems), this

[17] For example, high volatility is claimed to be a prerequisite for profitable day trading (see Bernstein (1998)).

result is consistent with the hypothesis that fleeting orders are more widely used when it is necessary to search among multiple pools of liquidity.

6. Dynamic analysis

Although many attributes of firms vary in the cross-section, changes in market conditions may also lead to different order strategies over time. To assess these effects, we estimate separate logit specifications for each firm over time:

$$\log\left(\frac{\Pr_{i,j,t}}{\Pr_{0,j,t}}\right) = X_{j,t}\beta_{i,j} \text{ for } i = 1,2,3$$

where i indexes event types, j indexes firms, and t indexes events. As in the cross-sectional analysis, the set of events is {*Regular limit order, Fleeting limit order, Market order, Hidden execution*}, with a regular limit order corresponding to event i=0.

The explanatory variables comprise measures intended to capture dynamic variation in market conditions. The prevailing Nasdaq NBBO spread reflects the cost of a market order. Volume over the prior five minutes is intended to capture variation in the general pace of market activity. Specification of other variables takes into account the direction (buy or sell) of the order. (Log) depth on the same side of the market (e.g., depth on the bid side for a buy order) and (log) depth on the opposite side (e.g., depth on the offer side for a buy order) are suggested by Parlour (1998). We also include a momentum variable defined as the return over the preceding five minutes for a buy order, and the negative of the return over the preceding five minutes for a sell order. This essentially measures the price change in the direction of the order. In utilizing these variables, we are viewing order direction as a predetermined, exogenous characteristic of the order, and modeling the event outcome (regular limit, fleeting limit, etc.) conditional on this direction. Dummy variables for the first and last hour of trading are included to capture deterministic intraday patterns.

As in the cross-sectional analysis, orders are assumed independent. This assumption is more suspect than in the cross-sectional analysis for at least two reasons. First, some of explanatory variables (e.g., depth) are determined at a given time by the history of earlier events. Second, aspects of "market conditions" that aren't modeled are nevertheless likely to persist over periods encompassing multiple events.

The specification is estimated separately for each firm. As in Section 5, we report in a table only the results corresponding to the subsample of 279 firms with available depth on the opposite side of the order. For ease of computation, we impose an upper limit of 10,000 observations for each firm. Maximum likelihood estimation converged without warnings in all but 12 of the firms, which were dropped from subsequent analysis.

For each firm we compute base probabilities, the probabilities associated with an increase of one standard deviation in the independent variables, and the changes in probabilities (similar to the cross-sectional analysis). Here, however, we use the mean and standard deviation specific to each firm. For the purpose of summarizing these estimations, we compute the cross-firm means of the probabilities. These are reported in Table 8.

In the cross-sectional analysis, the changes in the explanatory variables generally had little impact on the probabilities of regular limit orders. This is not the case in the dynamic analysis. As measured by the overall magnitude of the changes in probabilities, the most important effect is return momentum. For buy orders, an increase of one standard deviation in the return over the prior five minutes is associated with a 7.2 percent decrease in the probability of a regular limit order, a 2.9 percent increase in the probability of a fleeting limit order, a 3.7 percent increase in the probability of a market order, and a 0.7% increase in the probability of a hidden execution. The net effect is a substitution away from regular limit orders and towards fleeting and market orders. This is intuitively reasonable: when prices are moving in the same direction as the trading intention, an order that is not aggressive runs a high risk of missing the market. Here we clearly observe that fleeting orders are used in a similar fashion to the more aggressive market orders and unlike the "patient" limit orders.

We observe that an elevated Nasdaq spread means more regular limit orders and an offsetting decline in fleeting orders and to a lesser extent a decline in market orders.

Higher spreads seem to make traders more cost conscious. Therefore, they shift to less aggressive and more patient orders that are less expensive. Since fleeting limit orders demand immediacy and liquidity, their use declines and regular limit orders use increases. Here as well, limit and fleeting orders look very different.

A one-standard-deviation increase in volume during the preceding five minutes is associated with a relatively large increase in the probability of a hidden execution (2.3 percent). Regular market orders increase by 1.7 percent, regular limit orders decrease by 2.2 percent, and fleeting limit orders decrease by 1.7 percent. These results suggest that periods of high trading activity are associated with greater hidden depth. As hidden depth increases, limit orders submitted inside the quote to search for hidden liquidity are executed and we observe the increase in the hidden executions category. With the resulting higher fill rates for incoming limit orders inside the quote, fewer would be added to the book and subsequently cancelled, and the incidence of fleeting limit orders in our classification scheme would decline.

It is also well known that volume is autocorrelated: High volume now is more likely if volume five minutes ago was high. The usual interpretation of this observation is that when an information event or a demand shock occurs in the market, execution of orders of different investors (or the same investors who break down their orders) takes time. We find that if volume was high in the previous five minutes, there is increased likelihood of market orders and hidden executions, the two events in our classification that result in reported volume (and therefore a high volume now).

Looking at the Island variables in the model, we observe that when opposite-side depth increases by one standard deviation, the probability of a market order increases by 4.5 percent. This is offset by a 2.7 percent decline in regular limit orders and a 2.3 percent decline in fleeting limit orders. As depth on the opposite site of the book increases, it is easier and presumably cheaper to demand liquidity using a market order, and so there is no reason to incur the cost (in time and potential price movement) of searching for hidden orders. Therefore, both fleeting and regular limit decrease and market orders increase. Parlour (1998) suggests that for a buy order, for example, elevated depth on the sell side implies strong competition from sellers, and increased likelihood of a subsequent market sell order, favoring the immediate submission of a limit buy order. Our results are not consistent with her prediction.

As depth on the same side of the order increases, those who demand liquidity shift to more aggressive orders, and we observe a decline in fleeting orders and an increase in market orders. The difference between regular limit and fleeting orders is visible here as well, as the probability of fleeting limit orders declines while the probability of regular limit orders increases. The increase in the likelihood of observing a regular limit order presumably stems from traders splitting orders or pursuing correlated strategies. In other words, the increase in depth on the same side is due to recently arriving limit orders, and so if traders split orders or different traders follow correlated strategies there is increased likelihood that the next order arriving would be a limit order as well.

7. Discussion

The preceding sections show that the patterns and determinants of cross-sectional and dynamic variations in fleeting order utilization resemble those of market orders more closely than those of regular limit orders, i.e., that fleeting limit orders are a closer substitute for market orders. In this section we give a broader perspective.

We first note that our results cannot be explained by a simple model of fast and slow markets in which all clock-time event processes are subordinated to a common driving process (e.g., "information intensity"; see Clark (1973)). If this were the case, a "fast" market would result in more limit orders being classified as fleeting, but would not change the relative market and overall limit order proportions. That is, all substitution would occur between fleeting and regular limit orders. Our results clearly suggest otherwise.[18]

[18] Large (2004) suggests that fleeting orders convey information regarding market order arrival rates. From Table 2 the average daily number of market orders is 340. With a 6.5 hour trading day, this implies a mean inter-arrival time of approximately 68 seconds. This appears long relative to the two-second criterion used to classify a fleeting order. Thus, the Large mechanism would seem to require extreme dynamic variation in expected arrival rates. We do not model time-varying arrival rates in this paper, however, and so do not test his mechanism directly.

Are fleeting orders a universal feature of limit order markets? Existing limit order data samples are relatively recent and brief, so it is difficult to provide a definitive answer. We may nevertheless draw some tentative conclusions from U.S. samples that are available.

The earliest widely-available limit order dataset is TORQ, which covers a sample of NYSE stocks, October, 1989 through January, 1990 (see Hasbrouck (1992)). We formed a sample of all TORQ limit orders that were not tick-sensitive or otherwise qualified (roughly 300,000 orders). The cross-firm median execution rate (proportion of orders for which the first event is a partial fill) is 56% (vs. 18% for the Island sample). Thus, the Island sample exhibits a much higher proportion of cancellation.

Figure 2 depicts the median survival profiles of executions and cancellations in the Island and TORQ samples. Not surprisingly, the overall pace of activity is faster in the Island sample: both executions and cancellations tend to occur sooner than in the TORQ sample. It is striking, however, that the relative intensities of executions and cancellations differ. In the TORQ sample, the survival rate for execution events is higher than that of cancellations. In the Island sample this is reversed and the cancellation intensity is much higher than that of executions.

Lo, MacKinlay, and Zhang (2002) analyze a sample of NYSE limit orders from 1994–1995. Their sample is drawn from among the limit orders handled by one institutional broker, ITG, and hence are more representative of a specific trading clientele than an entire market. Lo et al. report that 53% of the limit orders were at least partially filled. Again, this implies a much lower rate of cancellations than in the Island sample.

To what might the differences across these samples be attributed? One possibility is the transparency of the trading process. During the TORQ sample period, information about the limit order book was unavailable away from the floor. Boehmer, Saar, and Yu (2004) find that when the NYSE began distributing information on the book in real time (January, 2002), the cancellation rate of limit orders increased by 17% and the average time to cancellation decreased by 25%. This suggests that NYSE trading practices have moved in the direction of Island's. Another possibility is tick size. The minimum tick during the TORQ and Lo et al. sample periods was 1/8. During our sample, the minimum tick was 1/16th on Nasdaq systems, but 1/256th on Island. A smaller tick size decreases the protection time priority offers to limit orders in the book, which may result in shorter durations.

Since fleeting orders appear to be a recent phenomenon, it is logical to consider the relevant ways in which markets have changed. We propose three factors as the major influences: technology, active trading, and fragmentation. The first of these, technology, particularly that related to automated order management systems and algorithmic trading, appears to be a necessary precondition. It is difficult to imagine human traders managing fleeting orders efficiently, especially given their low probability of execution.

It is impossible to date the "invention" of automated order management, and in any event the growth of these systems has occurred over time, but one indicative milestone might be the formation in 1995 of the initial FIX consortium. FIX, the Financial Information Protocol, is a standardized messaging protocol used to communicate orders and related information that is widely thought to have facilitated the construction of automated order management systems. Also, the advent of sophisticated order routing algorithms at the end of the 1990's (like Tradescape.com's Smart Order Routing Technology) has dramatically enhanced traders' ability to carry out automated trading strategies. In particular, these systems, aimed at achieving superior order execution, enabled traders to automatically sweep limit order books of different ECNs and engage in an effective search process that is at the heart of the new equilibrium.

The second contributing factor, active trading, refers to a broad array of practices (day trading, breaking down of large orders by buy-side trading desks, "statistical arbitrage", etc.) that require frequent, repeated and ongoing interaction with the market. The last factor, market fragmentation plays a role in order strategy by increasing the use of sequential strategies, decreasing the probability of order execution posted on any one book, and increasing the cost of limit order exposure. We believe that this environment—where technology enables rapid submission and cancellation of orders coming from active traders who can tap multiple pools of liquidity—gives rise to the "new equilibrium” in which traders engage in a search process to achieve the best possible execution.

The visibility of fleeting orders seems essential to this search process. ECNs like Island allow traders to hide their orders because they want to encourage liquidity provision. Hiding orders help traders manage the risks associated with exposure of limit orders, but is useful only if other traders search for the hidden liquidity. Market fragmentation gives rise to a coordination problem because patient traders need to decide on where to post their hidden orders and impatient traders need to decide on where to search for the hidden depth. The visible fleeting orders on Island serve to solve the coordination problem by signaling to the patient traders that impatient traders search this trading venue. This in turn encourages patient traders to submit hidden orders to the book, and both patient and impatient traders are able to fulfill their trading needs.

While we propose the new equilibrium as an explanation of the fleeting limit orders phenomenon, there may be other possible explanations. It could be that even a one- or two-second exposure time may suffice to encourage offsetting orders. A human trader could certainly react with this speed, but only if she were already monitoring the market. This seems so costly that we view it as more likely that the possible offsetting orders would be submitted by automated trading systems.

Another possibility is that the visible fleeting orders serve to induce competition among other quote setters, with the ultimate goal being execution of an order *in the opposite direction*. In its most blatant form, this has been considered market manipulation. Specifically, wholesale market makers often guarantee execution of retail orders at the NBBO prices. A seller of, say, 2000 shares, might place an aggressively priced *buy* limit order for 100 shares, with the intent of establishing a favorable sale price, only for as long as necessary to affect the trade with the wholesaler. SEC litigation reports document several of these “spoofing” cases in recent years. The numbers are small (ten or fewer cases), but this may reflect the difficulty of detection. In light of the enormous number of fleeting orders on Island, however, we doubt that spoofing is at the core of this phenomenon.

Among the factors we view as driving fleeting orders, none seems likely to abate in the near future. Accordingly, we view the fleeting order phenomenon as one of ongoing importance. This motivates consideration of fleeting orders' implications for overall welfare. If one considers the technology, active trading, and fragmentation environment as exogenous, then both patient and impatient traders are better off in the new equilibrium. Clearly patient traders can choose to submit visible limit orders, and their choice of supplying hidden liquidity must therefore be optimal. Similarly, the costly search of the impatient traders using fleeting orders makes sense only if they are better off carrying out the search than submitting marketable orders. Still, fragmentation and active trading could be viewed as responses to changing technological and regulatory environments, and hence endogenous to the equilibrium. It is therefore unclear whether all traders in the new equilibrium fare better than in an alternative economy with a centralized trading venue that operates more as a posted price market than a search market.

Note that in a fragmented market with heavy dealer participation, fleeting orders may have a negative byproduct. For example, when the SEC sought comment regarding the limit order display rule in 1996, some members of the dealer community viewed the rule as likely to result in "flickering" quotes, and a consequent decline in "market price integrity" (see final rule release 37619A). One interpretation of this point is that dealers establish reputations for maintaining quotes that are firm, ongoing and deep. These reputations may help them attract orders. When they are forced to display rapidly-changing bids and offers with small sizes, say due to the need to match fleeting orders on competing ECNs, this reputation is diminished. While this may suggest that dealer welfare may decline, the overall welfare implications for investors of this effect are unclear.

8. Conclusion

A common economic perspective, and arguably the historical reality, is that a customer limit order closely resembles a dealer quote. In the present paper, however, we focus on a key attribute of customer limit orders that in the current environment strongly differentiates them from dealer quotes, specifically the duration of the order's visibility. Dealers seeking to establish and maintain reputations as dependable liquidity suppliers will tend to publish quotes that are highly persistent. We demonstrate that many customer limit orders, on the other hand, are cancelled after an extent that at first appears inexplicably brief.

We investigate the new economic role of limit orders using a comprehensive dataset of order submissions and executions on the Island ECN. We find that over a quarter of the (non-marketable) limit orders submitted to Island are cancelled within two seconds. We term these "fleeting orders" and explore the puzzle they pose for the traditional thinking about limit order markets. We observe that fleeting orders are priced more aggressively than other limit orders, suggesting that they aim at exploring the existence of hidden depth in the book. We carry out both a cross-sectional and a dynamic multinomial logit analyses and our investigation demonstrates that fleeting orders are closer substitutes for market orders than for traditional limit orders. In other words, our results suggest that the aim of a trader who submits a fleeting order is to demand immediacy, in contrast to the traditional view of limit order traders as patient providers of liquidity.

The limit order data used in Lo, MacKinlay, and Zhang (2002) (which dates from 1995) and the TORQ database (1989-1990) do not exhibit a comparable frequency of extremely brief cancellations. It therefore appears that fleeting orders are a relatively new phenomenon. We conjecture that they have arisen due to the interplay of improved technology, the emergence of an active trading culture, and increased market fragmentation. Which, if any, of these factors is the primary driver remains a question for future research.

This new market environment, however, induced a change in traders' equilibrium strategies. Fleeting limit orders seem to be an important part of the search for liquidity and help sustain the new equilibrium. In this equilibrium, patient investors use hidden orders to supply liquidity to the book and impatient traders use limit orders priced inside the quote to search for the hidden liquidity. The impatient traders bear the cost of the search, but the larger supply of shares at better prices compensate them for it, and they are better off. The visibility of fleeting orders, those failed attempt at immediate executions, signals to patient traders that impatient traders are willing to search for liquidity in this specific trading venue. The signaling function has an important coordination role in the current environment (especially for Nasdaq stocks) because the Island ECN is only one among multiple destinations (dealers and alternative trading systems) for executing trades.

We view our documentation of fleeting limit orders and their role in the trading process as significant for understanding the new realities of trading and for modifying the conventional dichotomy of market vs. limit orders. The “smart order” in modern markets is essentially a trading strategy that can use multiple orders sent at different prices to different venues in order to achieve a single execution. The new “smart order” reality demands a new framework for thinking about optimal order choices and new theoretical models to guide us on trader behavior in markets organized as limit order books.

Our results also caution against applying the traditional thinking in evaluating the execution quality of a trading venue. It is notoriously difficult to define “best execution,” and this difficulty created a host of measures that are used by investors and regulators to judge execution and make order routing decisions. It is important to realize that some of the statistics used for this purpose may not be reflecting best execution in the new environment.

For example, the SEC mandates that each market center reports to the public on a monthly basis the fill rate of limit orders. This requirement is part of rule 11Ac1-5 that was adopted in November 2000 to help investors assess execution quality and determine the best destination for their orders. Historically a low fill rate in a particular venue might have indicated the absence of activity sufficient to generate high quality executions. We believe that this interpretation is no longer valid. Island is the destination of choice to many traders because it is extremely fast and thus enables traders to develop trading strategies that utilize modes of interaction, such as an efficient search for hidden liquidity, that are not feasible on other systems. As such, it provides a valuable service to traders and has therefore emerged as one of the dominant execution venues for Nasdaq stocks despite low reported fill rates. Certainly one goal of future research is the development of market-quality measures more suitable for the new environment.

Equilibrium in market microstructure models usually characterizes a balance in the supply and demand for immediacy. Fleeting orders, however, do not readily fit into this perspective. They are more naturally viewed as intermediate steps in dynamic strategies that are actively pursuing executions. Except for their brevity, however, fleeting orders are still functionally ordinary limit orders. It may be the case that higher efficiency would result from vastly different order types. Fischer Black suggested that orders in an electronic market would eventually be characterized by direction and "urgency", with the latter term reflecting a desired *rate* of execution and an acceptable cost (in terms of price deterioration), (Black (1995)). Limit orders would be indexed, i.e., "continually adjusted to reflect market conditions." Brown and Holden (2002) explicitly analyze the welfare implications of adjustable limit orders and find in simulations that they benefit both the limit order traders and those seeking immediacy via market orders.[19]

Our analysis also relates to one of the conclusions from the double auction literature that high efficiency of prices (or convergence to the competitive equilibrium) depends more on the institution itself—the double auction—than on the optimality of the traders' actions (e.g., Easley and Ledyard (1993), Friedman (1993), Gode and Sunder (1993), Rust, Miller, and Palmer (1993)). This is most vividly demonstrated by Gode and Sunder (1993), who show how markets populated by "zero-intelligence" traders demonstrate high efficiency and converge to the competitive equilibrium. Friedman (1993) notes that many different forms of double auctions exist in the world. He hypothesizes that efficiency could be very sensitive to environmental details, and that the most appropriate variant of double auction is different in each circumstance. It may be that some features of ECNs, like the ability to submit hidden orders, were instituted in

[19] It is interesting to note that INET recently introduced special pegged orders that enable clients to price orders relative to the current market price of a security.

response to the changing environment we discuss. However, our results suggest that the behavior of traders also adjusts to the new environment and creates a whole new way of interacting in the double auction institution. Therefore, it may be that in a complex environment, both changes in the rules of the double auction institution and the optimal responses of traders are important in achieving efficient outcomes.

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.

Biais, B., Bisiere, C., Spatt, C., 2003. Imperfect competition in financial markets: Island vs. Nasdaq. Unpublished working paper. University of Toulouse.

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.

Black, F., 1995. Equilibrium exchanges. Financial Analysts Journal 51, 23-29.

Bloomfield, R., O'Hara, M., Saar, G., 2004. The "make or take" decision in an electronic market: evidence on the evolution of liquidity. Journal of Financial Economics, forthcoming.

Boehmer, E., Saar, G., Yu, L., 2004. Lifting the veil: an analysis of pre-trade transparency at the NYSE. Journal of Finance, forthcoming.

Brown, D. P., Holden, C. W., 2002. Adjustable limit orders. Unpublished working paper. University of Indiana.

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

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

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.

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

Easley, D., Ledyard, J. O., 1993. Theories of price formation and exchange in double oral auctions. In: Friedman, D. and Rust, J. (Eds.), The Double Auction Market Institutions, Theories and Evidence, Proceedings Volume XIV, Santa Fe Institute. Addison-Wesley, Reading, MA, 63-97.

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

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. Journal of Finance 47, 576-605.

Ellul, A., Holden, C. W., Jain, P., Jennings, R., 2002. Determinants of order choice on the New York Stock Exchange. Unpublished working paper. Indiana University.

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.

Friedman, D., 1993. The double auction market institution: a survey. In: Friedman, D. and Rust, J. (Eds.), The Double Auction Market Institutions, Theories and Evidence, Proceedings Volume XIV, Santa Fe Institute. Addison-Wesley, Reading, MA, 3-26.

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.

Gode, D. K., Sunder, S., 1993. Allocative efficiency of markets with Zero-Intelligence traders: markets as a partial substitute for individual rationality. Journal of Political Economy 101, 119-137.

Goettler, R. L., Parlour, C. A., Rajan, U., 2003. Welfare in a dynamic limit order market. Unpublished working paper. GSIA, Carnegie Mellon University.

Hamao, Y., Hasbrouck, J., 1995. Securities trading in the absence of dealers: trades and quotes on the Tokyo Stock Exchange. Review of Financial Studies 8, 849-78.

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

Hansch, O., 2003. Island tides: exploring ECN liquidity. Unpublished working paper. Smeal College of Business, Penn State University.

Harris, L., 1998. Optimal dynamic order submission strategies in some stylized trading problems. Financial Markets, Institutions and Instruments 7, 1-76.

Hasbrouck, J., 1992. Using the TORQ Database. Unpublished working paper. New York Stock Exchange.

Hasbrouck, J., Saar, G., 2002. Limit orders and volatility in a hybrid market: the Island ECN. Unpublished working paper. Department of Finance, Stern School, NYU.

Hendershott, T., Jones, C. M., 2004. Island goes dark: transparency, fragmentation and regulation. Unpublished working paper. Graduate School of Business, Columbia University.

Hollifield, B. M. R. A. S. P., 2004. Empirical analysis of limit order markets. Review of Economic Studies 71, 1027-1063.

Kaniel, R., Liu, H., 2004. So what orders do informed traders use? Unpublished working paper. Fuqua School of Business, Duke University.

Large, J., 2004. Cancellation and uncertainty aversion on limit order books. Unpublished working paper. Nuffield College, Oxford University.

Lo, A. W., MacKinlay, A. C., Zhang, J., 2002. Econometric models of limit order execution. Journal of Financial Economics 65, 31-71.

Nguyen, V. T., Van Ness, B. F., Van Ness, R. A., 2003. An examination of the dissemination of Island trades through the Cincinnati Stock Exchange. Unpublished working paper. College of Business, University of Mississippi.

Parlour, C., 1998. Price dynamics in limit order markets. Review of Financial Studies 11, 789-816.

Renaldo, A., 2004. Order aggressiveness in limit order book markets. Journal of Financial Markets 7, 53-74.

Rosu, I., 2004. A dynamic model of the limit order book. Unpublished working paper.

MIT.

Rust, J., Miller, J. H., Palmer, R., 1993. Behavior of trading automata in a computerized double auction market. In: Friedman, D. and Rust, J. (Eds.), The Double Auction Market Institutions, Theories and Evidence, Proceedings Volume XIV, Santa Fe Institute. Addison-Wesley, Reading, MA, 155-198.

Sandas, P., 2001. Adverse selection and competitive market making: evidence from a pure limit order book. Review of Financial Studies 14, 705-734.

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

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

Stoll, H. R., 1978. The supply of dealer services in securities markets. Journal of Finance 33, 1133-1151.

Table 1.
Sample Summary Statistics

Our sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). The following variables are calculated for each stock over the sample period using data in CRSP: AvgCap is the average daily market capitalization (the number of shares outstanding multiplied by the daily closing price), AvgTrd is the average number of daily trades, AvgVol is the average daily share volume, MedTurn in the median daily turnover (the number of shares traded divided by the number of shares outstanding), AvgPrc is the average daily closing price, and StdRet is the standard deviation of daily returns. The following variables are calculated from the intraday Nastraq database for each stock: AvgSprd (in $) is the average dollar spread (using all NBBO quotes in the sample period), and AvgSprd (in %) is the average relative spread (dollar spread divided by the quote midpoint). The table presents cross-sectional summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

		AvgCap (in million $)	**AvgTrd**	**AvgVol** (1,000s shrs)	**MedTurn** (in %)	**AvgPrc** (in $)	**StdRet** (in %)	**AvgSprd** (in $)	**AvgSprd** (in %)
Entire Sample	Mean	10,205	2,677	1,873	1.288	63.03	4.36	0.256	0.46
	Median	3,081	1,066	877	1.107	49.82	4.33	0.187	0.44
	Std. Dev.	38,104	4,413	3,504	0.946	45.66	1.69	0.218	0.25
	Min	824	16	7	0.028	8.40	0.18	0.052	0.07
	Max	494,932	28,654	30,073	5.208	326.58	10.83	1.910	2.79
	N	300	300	300	300	300	300	300	300
AvgCap Groups	Low	1,490	406	375	0.715	34.56	3.54	0.191	0.60
	Medium	3,081	1,058	884	1.188	51.22	4.65	0.210	0.45
	High	8,197	3,808	2,048	1.436	79.64	4.35	0.164	0.25
AvgTrd Groups	Low	1,583	295	290	0.479	36.45	2.97	0.194	0.59
	Medium	2,965	1,066	885	1.232	51.21	4.65	0.206	0.45
	High	7,730	4,174	2,514	1.998	75.49	4.96	0.163	0.26
StdRet Groups	Low	2,187	353	361	0.468	36.82	2.65	0.152	0.49
	Medium	3,882	1,421	1,115	1.333	52.37	4.33	0.184	0.40
	High	3,216	2,023	1,077	1.730	69.48	5.91	0.292	0.42

Table 2.
Island Summary Statistics

This table presents summary statistics on the orders submitted to the Island ECN. Our sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). All variables are calculated for each stock over the sample period using order-level data from Island: NumLMT is the average daily number of (non-marketable) limit orders (including hidden limit orders that are inferred from executions), SizeLMT is the average size of a limit order in shares, NumCanc is the average daily number of order cancellations, SizeCanc is the average size of a cancelled order in shares, NumMKT is the average daily number of market orders (where a market order is defined as an order that is matched upon arrival and so never appears in the book), and SizeMKT is the average number of a market order is shares. The table presents cross-sectional summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

		NumLMT	**SizeLMT** (in shares)	**NumCanc**	**SizeCanc** (in shares)	**NumMKT**	**SizeMKT** (in shares)
Entire Sample	Mean	966	572	672	618	340	335
	Median	285	585	221	627	61	329
	Std. Dev.	1,765	158	1,145	158	761	110
	Min	4	214	3	242	0	123
	Max	11,992	985	6,964	1,032	6,124	743
	N	300	300	300	300	300	299
AvgCap Groups	Low	76	626	67	649	9	356
	Medium	284	565	229	620	56	326
	High	1,329	527	1,005	584	362	302
AvgTrd Groups	Low	53	636	47	651	5	371
	Medium	285	564	217	615	61	323
	High	1,588	506	1,178	575	478	291
StdRet Groups	Low	65	668	58	683	6	395
	Medium	419	585	304	637	92	337
	High	591	431	394	494	212	237

Table 3.
Limit Order Submission Proportions and Fill Rates

This table presents summary statistics on the order mix and fill rate of limit orders on the Island ECN. The submission proportion of limit orders (LMT-SP) for orders is defined as the number of limit orders divided by the sum of limit and market orders (where market orders are those orders that are matched upon arrival and never enter the book). LMT-SP for shares is similarly defined in terms of the number of shares in submitted limit and market orders. The Fill Rate for orders is the number of limit orders that were at least partially executed divided by the total number of limit orders submitted. The Fill Rate for shares is analogously the number of shares executed divided by the total number of shares submitted in limit orders. All variables are calculated for each stock over the sample period using order-level data from Island. The table presents summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

		LMT-SP (in terms of orders)	**LMT-SP** (in terms of shares)	**Fill Rate** (in terms of orders)	**Fill Rate** (in terms of shares)
Entire Sample	Mean	0.830	0.892	0.184	0.126
	Median	0.820	0.887	0.193	0.128
	Std. Dev.	0.086	0.058	0.101	0.074
	Min	0.637	0.747	0.000	0.000
	Max	1.000	1.000	0.435	0.339
	N	300	300	300	300
AvgCap Groups	Low	0.893	0.937	0.104	0.067
	Medium	0.821	0.890	0.196	0.123
	High	0.776	0.856	0.245	0.168
AvgTrd Groups	Low	0.924	0.953	0.075	0.049
	Medium	0.817	0.886	0.197	0.128
	High	0.748	0.837	0.285	0.194
StdRet Groups	Low	0.925	0.950	0.072	0.052
	Medium	0.816	0.881	0.197	0.135
	High	0.763	0.852	0.265	0.173

Table 4.
Proportion of Orders in Event Categories

This table presents the proportions of orders in each of our four categories (orders in a category divided by the total number of orders in all categories). Regular Limit orders are non-marketable limit orders that are submitted to the book and are not cancelled in the first two seconds after submission (though they can be executed during that time). Fleeting Limit orders are non-marketable limit orders that are cancelled within two seconds after submission. Market orders are defined as orders that are matched upon arrival to Island (and so never appear in the book) with visible limit orders. Hidden Executions are defined as orders that are matched with hidden depth in the book upon arrival to Island. All variables are calculated for each stock over the sample period using order-level data from Island. Each order is put into one of two subsamples for a particular stock: (i) orders that arrive to Island when there is opposing visible depth, and (ii) orders that arrive to Island when there is no opposing visible depth. If there are fewer than 500 orders for a certain stock in one subsample, this stock is dropped from the analysis of orders in that subsample. We present the proportions separately for sample with visible opposing depth (in Panel A) and without visible opposing depth (in Panel B) because the category of Market orders is well defined only for the subsample where there is visible depth on the opposite side of the book when the order arrives. The table presents summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

Panel A: Orders that Arrive when there is Visible Depth on the Opposite Side of the Book

		Regular Limit	**Fleeting Limit**	**Market**	**Hidden Execution**
Entire Sample	Mean	0.582	0.212	0.178	0.028
	Median	0.589	0.200	0.182	0.019
	Std. Dev.	0.058	0.090	0.058	0.025
	Min	0.406	0.053	0.018	0.000
	Max	0.736	0.491	0.419	0.137
	N	279	279	279	279
AvgCap Groups	Low	0.583	0.238	0.151	0.012
	Medium	0.598	0.206	0.179	0.022
	High	0.586	0.158	0.203	0.030
AvgTrd Groups	Low	0.555	0.276	0.123	0.009
	Medium	0.602	0.211	0.168	0.019
	High	0.586	0.141	0.220	0.036
StdRet Groups	Low	0.586	0.255	0.135	0.008
	Medium	0.589	0.210	0.181	0.017
	High	0.589	0.147	0.206	0.043

Panel B: Orders that Arrive when no Visible Depth on the Opposite Side of the Book Exists

		Regular Limit	**Fleeting Limit**	**Market**	**Hidden Execution**
Entire Sample	Mean	0.636	0.361		0.003
	Median	0.647	0.349		0.001
	Std. Dev.	0.118	0.119		0.004
	Min	0.331	0.053		0.000
	Max	0.947	0.669		0.019
	N	101	101		101
AvgCap Groups	Low	0.644	0.355		0.000
	Medium	0.610	0.385		0.001
	High	0.672	0.325		0.003
AvgTrd Groups	Low	0.622	0.377		0.000
	Medium	0.669	0.326		0.003
	High	0.707	0.284		0.008
StdRet Groups	Low	0.635	0.365		0.000
	Medium	0.674	0.323		0.002
	High	0.633	0.361		0.007

Table 5.
Pricing of Limit Orders

This table presents the proportions of regular and fleeting limit orders that are priced behind, at, and ahead of the prevailing Island best bid or offer (BBO) prices. Regular Limit orders are non-marketable limit orders that are submitted to the book and are not cancelled in the first two seconds after submission (though they can be executed during that time). Fleeting Limit orders are non-marketable limit orders that are cancelled within two seconds after submission. All variables are calculated for each stock over the sample period using order-level data from Island. Each order is put into one of two subsamples for a particular stock: (i) orders that arrive to Island when there is opposing visible depth, and (ii) orders that arrive to Island when there is no opposing visible depth. If there are fewer than 500 orders for a certain stock in one subsample, this stock is dropped from the analysis of orders in that subsample. We present the proportions separately for sample with visible opposing depth (in Panel A) and without visible opposing depth (in Panel B). The table presents summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

Panel A: Orders that Arrive when there is Visible Depth on the Opposite Side of the Book

		Regular Limit Orders			**Fleeting Limit Orders**		
		% of Orders Priced Relative to Island's BBO			% of Orders Priced Relative to Island's BBO		
		Behind (in %)	**At** (in %)	**Ahead** (in %)	**Behind** (in %)	**At** (in %)	**Ahead** (in %)
Entire Sample	Mean	15.0	12.7	72.3	9.3	7.1	83.6
	Median	13.3	12.0	73.3	7.1	6.2	87.3
	Std. Dev.	10.2	6.2	15.6	8.0	5.1	12.1
	Min	0.2	0.3	26.7	0.0	0.0	39.2
	Max	48.2	31.9	99.5	44.3	25.9	100.0
	N	279	279	279	279	279	279
AvgCap Groups	Low	8.0	8.5	84.0	3.3	3.9	92.8
	Medium	12.1	11.4	76.6	7.0	5.9	87.9
	High	19.6	16.0	65.3	14.0	9.1	77.8
AvgTrd Groups	Low	4.9	6.7	88.8	2.2	2.7	94.9
	Medium	12.6	11.2	76.2	5.2	5.9	88.1
	High	23.9	17.8	59.0	15.8	10.6	74.0
StdRet Groups	Low	4.8	8.2	87.2	2.5	4.3	92.9
	Medium	13.0	13.3	72.9	6.8	5.9	87.6
	High	19.1	12.8	66.4	10.7	8.0	79.8

Panel B: Orders that Arrive when no Visible Depth on the Opposite Side of the Book Exists

		Regular Limit Orders			Fleeting Limit Orders		
		% of Orders Priced Relative to Island's BBO			% of Orders Priced Relative to Island's BBO		
		Behind (in %)	**At** (in %)	**Ahead** (in %)	**Behind** (in %)	**At** (in %)	**Ahead** (in %)
Entire Sample	Mean	3.4	5.6	91.0	2.7	2.9	94.4
	Median	2.4	4.7	91.9	1.8	1.7	95.5
	Std. Dev.	3.0	4.0	6.4	2.9	3.1	5.0
	Min	0.0	0.0	59.0	0.0	0.0	71.0
	Max	20.1	20.8	99.8	17.3	14.1	100.0
	N	101	101	101	101	101	101
AvgCap Groups	Low	1.9	3.8	93.7	1.4	1.2	97.0
	Medium	2.4	4.7	91.9	1.9	1.5	95.0
	High	3.9	8.9	87.4	2.5	4.3	91.7
AvgTrd Groups	Low	1.8	3.7	94.1	1.4	1.2	97.2
	Medium	4.3	8.5	86.8	2.8	3.9	93.2
	High	14.2	20.5	65.3	12.7	12.9	74.4
StdRet Groups	Low	1.7	4.0	94.7	1.4	1.5	97.0
	Medium	5.5	7.1	87.1	2.4	4.3	92.2
	High	5.1	4.8	89.9	5.0	1.0	93.4

Table 6
Fleeting Orders and Hidden Executions

This table presents summary statistics on the usages of fleeting orders and hidden executions on the Island ECN. We report two measures of the frequency of fleeting orders relative to all limit orders: one in terms of the number of orders and the other in terms of the number of shares in the submitted orders. We also report two measures of the frequency of hidden executions relative to all executions: one in terms of the number of executions and the other in terms of the number of shares executed. All variables are calculated for each stock over the sample period using order-level data from Island. The table presents summary statistics for the entire sample and separately the medians for three groups (low, medium, and high) sorted by market capitalization, number of daily trades, and return standard deviation.

		Fleeting Orders / All Limit Orders	Fleeting Orders / All Limit Orders (in terms of shares)	Hidden Executions / All Executions	Hidden Executions / All Executions (in terms of shares)
Entire Sample	Mean	0.137	0.118	0.277	0.325
	Median	0.116	0.103	0.254	0.298
	Std. Dev.	0.105	0.095	0.117	0.120
	Min	0.000	0.000	0.059	0.058
	Max	1.000	1.000	0.884	0.915
	N	299	299	300	300
AvgCap Groups	Low	0.089	0.074	0.297	0.338
	Medium	0.131	0.117	0.257	0.299
	High	0.151	0.127	0.212	0.265
AvgTrd Groups	Low	0.066	0.061	0.340	0.381
	Medium	0.123	0.106	0.257	0.302
	High	0.172	0.138	0.191	0.245
StdRet Groups	Low	0.063	0.050	0.308	0.335
	Medium	0.112	0.102	0.258	0.309
	High	0.201	0.175	0.198	0.266

Table 7
Cross-Sectional Logit Model

This table reports the results of a cross-sectional multinomial logit model for the four categories of orders (regular limit, fleeting limit, market, and hidden execution). We report both probability levels (top portion of the table) and differences from the base case (bottom portion of the table) implied by the logit model. Our sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). Let $i = 0,1,2,3$ denote an index corresponding to the categories *{Regular Limit Order, Fleeting Limit Order, Market Order, Hidden Execution}*, and let *j* index firms. The probability of event *i* for firm *j* is:

$$\log\left(\Pr_{i,j}/\Pr_{0,j}\right) = X_j \beta_i \text{ for } i = 1,2,3$$

where X_j is the vector of explanatory variables. The following are the explanatory variables we use in the estimation: Intercept, StdRet is the standard deviation of daily return, LgMedTurn is the log of median daily turnover, LgAvgPrc is the log of average closing prices, LgAvgCap is the log of the daily average market capitalization, and HHI is the Herfindahl-Hirschman Index of volume across all market makers and alternative trading systems. All variables are calculated for each stock over the sample period using data from CRSP except the Herfindahl-Hirschman Index which is computed using monthly volume data provided by Nasdaq. The model is estimated on the subsample of orders that arrive when there is visible depth on the opposite side of the market. Observations for each stock are weighted to give equal weighting to all stocks. The column labeled “Base Case” reports the probabilities implied by setting all explanatory variables to their sample means. In each of the remaining columns, the indicated explanatory variable is set to one standard deviation above its mean.

			Changed Variables				
	Category	**Base Case**	**StdRet**	**LgMedTurn**	**LgAvgPrc**	**LgAvgCap**	**HHI**
Predicted Probability Levels	Regular Limit	0.594	0.600	0.591	0.582	0.600	0.597
	Fleeting Limit	0.207	0.166	0.181	0.232	0.164	0.204
	Market	0.177	0.199	0.202	0.159	0.209	0.178
	Hidden Execution	0.022	0.035	0.026	0.026	0.027	0.021
Probability Differences Relative to Base Case	Regular Limit		0.006	-0.003	-0.012	0.006	0.003
	Fleeting Limit		-0.041	-0.026	0.025	-0.043	-0.003
	Market		0.022	0.024	-0.018	0.032	0.000
	Hidden Execution		0.013	0.005	0.005	0.005	-0.000

Table 8
Dynamic Logit Model

This table reports the results of a dynamic multinomial logit model for the four categories of orders (regular limit, fleeting limit, market, and hidden execution). We report both probability levels (top portion of the table) and differences from the base case (bottom portion of the table) implied by the logit model. Our sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). Let $i = 0,1,2,3$ denote an index corresponding to the categories *{Regular Limit Order, Fleeting Limit Order, Market Order, Hidden Execution}*, *j* index firms, and *t* index events. A logit model is estimated separately for each firm:

$$\log\left(\frac{\Pr_{i,j,t}}{\Pr_{0,j,t}}\right) = X_{j,t}\beta_{i,j} \text{ for } i = 1,2,3$$

where $X_{j,t}$ is the vector of explanatory variables. The following are the explanatory variables we use in the estimation: Return Momentum is the return over the preceding five minutes for a buy order, and (–1)*return for a sell order, NBBO Spread is the Nasdaq NBBO prevailing at the time of the order, LgVolume is log dollar Nasdaq volume in the preceding five minutes, LgDepthSameSide is the log of Island book depth on the same side of the market (e.g., depth on the bid side for a buy order), and LgDepthOtherSide is the log of Island book depth on the opposite side of the market. We also include dummy variables for the first and last hour of trading. Return Momentum, NBBO Spread, and LgVolume are calculated using data from the Nastraq database while the Island book depth variables are computed from Island order-level data. The direction of the event (buy or sell) is taken as exogenous. The model is estimated on the subsample of orders that arrive when there is visible depth on the opposite side of the market. The column labeled “Base Case” reports the probabilities implied by setting all explanatory variables to their stock-specific sample means. In each of the remaining columns, the indicated explanatory variable is set to one stock-specific standard deviation above the mean. The table reports cross-stock averages of the probabilities that are computed separately for each of the stocks in the sample..

			Changed Variables				
	Category	**Base Case**	**Return Momentum**	**NBBO Spread**	**LgVolume**	**LgDepth SameSide**	**LgDepth OtherSide**
Predicted Probability Levels	Regular Limit	0.603	0.531	0.634	0.581	0.608	0.575
	Fleeting Limit	0.209	0.232	0.181	0.192	0.196	0.186
	Market	0.169	0.211	0.164	0.185	0.177	0.214
	Hidden Execution	0.020	0.028	0.022	0.043	0.019	0.026
Probability Differences Relative to Base Case	Regular Limit		-0.072	0.031	-0.022	0.006	-0.027
	Fleeting Limit		0.023	-0.028	-0.017	-0.013	-0.023
	Market		0.042	-0.005	0.017	0.008	0.045
	Hidden Execution		0.007	0.002	0.023	-0.001	0.005

Figure 1
Survival probabilities for cancellations and executions

This figure plots median survival probabilities for cancellation (with execution taken as the censoring event) and for execution (with cancellation taken as the censoring event) of Island limit orders. Our sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). The figure is based on order-level data from the Island ECN. If τ is the time of event occurrence, the survival function at time t is $\Pr[\tau > t]$. For each type of event, a survival function is estimated for each firm using the life-table (actuarial) method. The plot depicts cross-stock means.

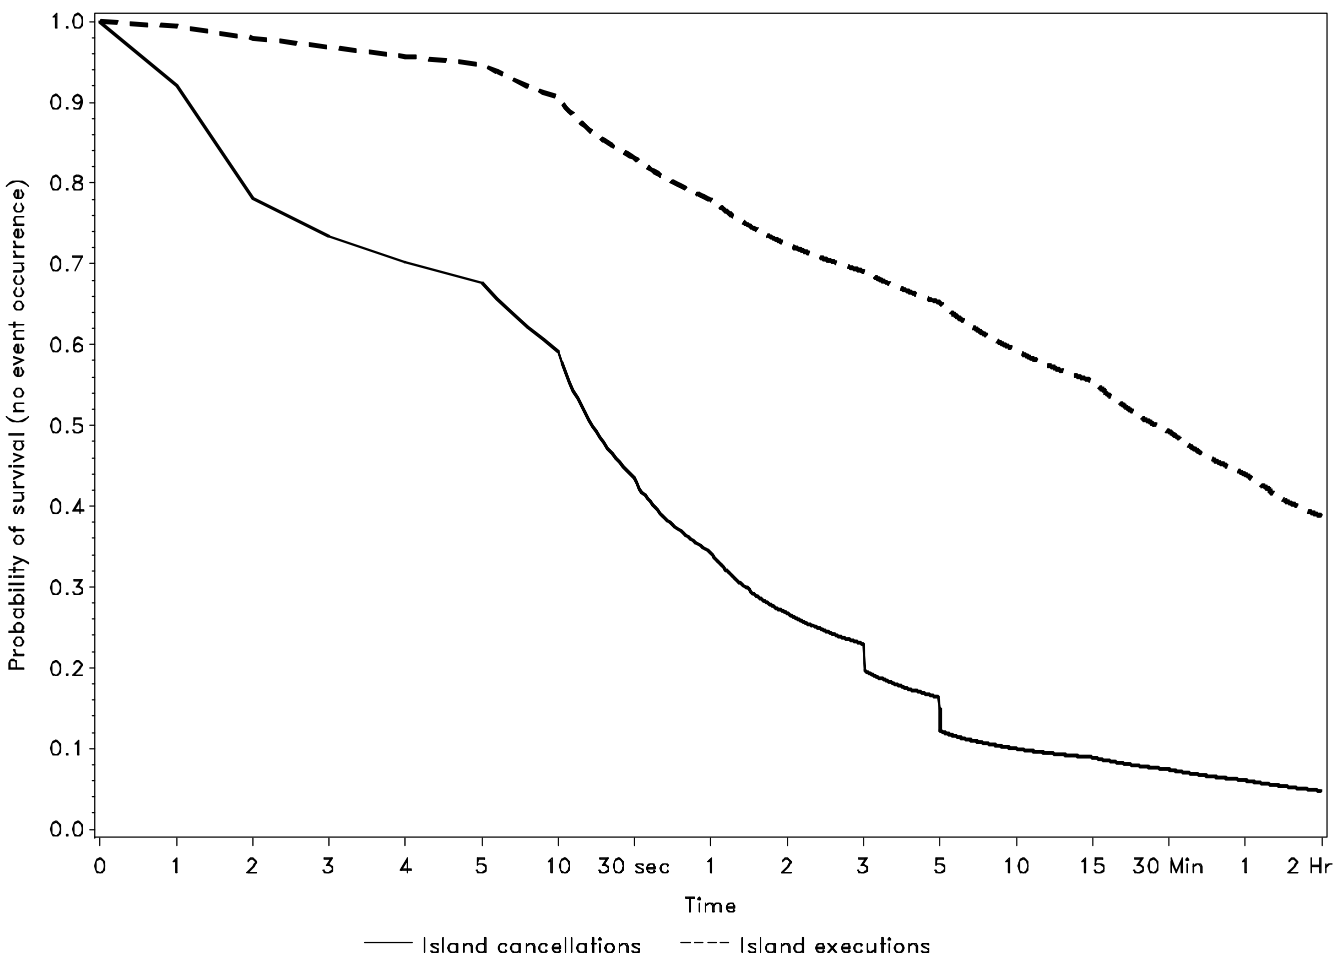

Figure 2
Survival probabilities for executions and cancellations

This figure compares cancellation and execution survival probabilities for limit orders on the Island ECN and the New York Stock Exchange. We plot median survival probabilities for cancellation (with execution taken as the censoring event) and for execution (with cancellation taken as the censoring event) of limit orders. We utilized two different samples for this plot. The Island sample consists of the 300 largest Nasdaq National Market stocks based on equity market capitalization on September 30, 1999. The sample period is the fourth quarter of 1999 (64 trading days). The NYSE sample is based on all limit orders in the TORQ database (150 firms, October 1989 through January 1990). If τ is the time of event occurrence, the survival function at time t is $\Pr[\tau > t]$. For each type of event, a survival function is estimated for each firm using the life-table (actuarial) method. The plotted survival probability at time t is the cross-stock median survival probability at that time.