The Determinants of Limit Order Cancellations

Petter Dahlström, Björn Hagströmer, Lars L. Nordén*

Heavy limit order traffic congests limit order books and triggers regulatory responses such as limit order fees, minimum quote lives, and speed bumps. We postulate that limit order cancellations occur because of reductions in the expected profits of the orders. Using a model of liquidity supply with adverse selection costs, we predict that limit order cancellations relate to fluctuations in the bid-ask spread, the bid-ask depth imbalance, and the order queue position. We find strong empirical support for the predictions, in particular for the order queue position, which is new to the literature. Our evidence indicates that the frequent cancellations by high-frequency traders are due to their close monitoring of fluctuations in fundamental value.

*All authors are at Stockholm Business School, Stockholm University; SE-106 91 Stockholm; Sweden. Contact by email at petter.dahlstrom@sbs.su.se; bjorn.hagstromer@sbs.su.se; lars.norden@sbs.su.se. Björn Hagströmer and Lars L. Nordén are affiliated with the Swedish House of Finance and are grateful to the Jan Wallander and Tom Hedelius foundation and the Tore Browaldh foundation for research support. We thank Frank Hatheway, Patrik Sandås, Bart Zhou Yueshen and Marius Zoican for useful discussions.

1 Introduction

The stylized image of a limit order market consists of patient market makers who stand ready to trade with incoming investors. In reality, liquidity suppliers are restless rather than patient. At the equity market, Conrad, Wahal, and Xiang (2015) report that the top of book quotes of each large-cap US stock are updated on average once every 50 milliseconds. That corresponds to six quote updates in the blink of an eye. The short-lived limit orders are known as “fleeting orders” (Hasbrouck and Saar, 2009), “flickering quotes” (Baruch and Glosten, 2016), or “phantom liquidity” (US Securities and Exchange Commission, SEC, 2010). In this paper, we seek to understand why submitters of limit orders cancel or revise their orders at such high frequency.

Just like road traffic congestion exerts negative externalities on all travelers, heavy limit order traffic (including submissions, modifications, and cancellations) puts a strain on the market infrastructure and causes delays and data processing costs to all market participants. In addition, fast cancellations expose traders to execution risk, as the market liquidity potentially changes from the time when a trader sends an order to when it arrives at the market. To curb the negative externalities, many market regulators attempt to control the order traffic. For example, the Canadian regulator applies a congestion charge, i.e., a fee on limit order submissions and cancellations. In Germany and at several exchanges, regulators cap the order-to-trade ratio to deter excessive order traffic. EBS, one of the leading inter-dealer platforms in the foreign exchange market, disallows cancellations within a quarter of a second after the submission of a limit order.

Congestion is a macro-level problem because the aggregate volume causes the problems. The individual vehicle or limit order is harmless if it travels alone, but during rush hour, it contributes to the problem. When regulating a macro-level problem, understanding the micro-level is key to avoid unintended consequences. Accordingly, we aim to understand the economics of cancellations at the individual order level.

A limit order to buy or sell a stock is a free option issued to all market participants (Copeland and Galai, 1983). Rational economic agents offer such options to the market only as long as the limit order price, relative the fundamental value of the stock, implies a non-negative expected profit to the issuer. We extend the model of liquidity supply in the limit order book by Sandås (2001), which builds on the models by Glosten (1994) and Seppi (1997). The key feature of the model is that the expected profit of a limit order is a function of the fundamental value and the queue position of the limit order. We show that, in equilibrium, the expected profit is a function of the prices and volumes at the best bid and ask prices in the limit order book. This outcome is important because it circumvents the empirical challenge that the fundamental value is unobservable.

We postulate that a trader cancels a limit order when the expected profit of the order turns negative by a shock to the fundamental value. The cancellation decision is not explicit in the model, but we form dynamic empirical predictions by comparing equilibrium liquidity before and after changes in the fundamental value. Thus, we view order submissions and cancellations as part of the transition to the new equilibrium. Van Kervel (2015) takes a similar approach to form dynamic predictions from a static model.

To see how a change in the fundamental value can trigger cancellations, consider a stock which order book is in equilibrium, meaning that no further limit orders with non-negative expected profit can be added. Now consider the arrival of positive news about the stock, moving the fundamental value closer to the best ask price. As a consequence, the expected profit of limit orders at the best ask price is reduced. The marginal ask-side limit order becomes unprofitable in expectation, and the rational response is to cancel it. The larger the shock to the fundamental value is, the more cancellations of ask-side orders occur.

We predict that orders with higher expected profit are more resilient to fluctuations in the fundamental value, and accordingly have longer expected duration (meaning that they remain in the order book for a longer time). In terms of order book variables, the duration increases with the bid-ask spread, the depth on the opposite side of the order book, and the execution priority of the limit order in question. The latter prediction is novel to the literature and implies that limit orders posted at the same price do not have the same probability of cancellation. Traders who submit market orders in the model may have private information, implying that the market makers face adverse selection risk. Limit orders in the back of the order queue require a large market order to execute and incur higher adverse selection costs than orders in the front of the queue.

The data we use to analyze cancellations empirically reflect limit order book events at NASDAQ Stockholm (henceforth NASDAQ). The sample includes data on thirty large-cap stocks for all trading days of May 2014. The data allow us to track each limit order from submission to either cancellation or execution. We can also trace the order execution priority, which is key to test the model predictions.

The model predictions find strong support in the data. We draw a random sample of orders posted at the best bid and ask prices, and track for how long they remain in the order book after the sampling time. We find that orders with high queue priority (that is, with many orders behind it in the order queue) have significantly longer expected time to cancellation than orders with low queue priority. Similarly, orders with a high opposite-side depth or a wide bid-ask spread conditional upon execution, have relatively long expected duration.

We perform the expected duration analysis ex ante, based on the market and order conditions at the time when we sample the order. To understand the triggers for limit order cancellations, we also analyze the influence of limit order book events for a period after the sampling. Specifically, we analyze the hazard rate of cancellation using up to 100 order book events for each limit order. The hazard rate measures the cancellation intensity over the next instant. Our prediction is that any order book event that indicates a weakening of the expected profit of the limit order increases the hazard rate of cancellation. Following Hasbrouck and Saar (2009), we estimate a proportional duration hazard model (PDHM) with time-varying covariates. The PDHM allows us to analyze cancellations while controlling for executions (an exogenous censoring process), and the time-varying covariates account for changes in the expected profit during the limit order lifetime. The PDHM estimates lend support to our model predictions.

The model assumes an economy in which a publicly known fundamental value determines the liquidity supply decisions, and the adverse selection risk is equal across liquidity providers. As a result, we predict that only common values are driving limit order cancellations. When we take the model predictions to the data, it is important to control for potential private values, such as inventory effects (Raman and Yadav, 2014) and order monitoring costs (Liu, 2009; Fong and Liu, 2010). Another model limitation is the assumption that the stocks are traded at a single venue, where a fragmented marketplace would be more consistent with the empirical setting. Van Kervel (2014) shows that trades at competing venues work as signals about changes in the fundamental value, and, thus, influence cancellations. Budish, Cramton and Shim (2015) and Foucault, Kozhan and Tham (2016) follow the same reasoning, and show that liquidity suppliers who do not immediately cancel their orders when they become stale, are exposed to toxic arbitrage.

We extend the PDHM model to control for the alternative motives to cancel limit orders. The results show that even though all of them have statistically significant effects, our predicted effects are virtually unchanged. We also analyze the relative economic importance of each variable by calculating the cancellation hazard rate change (HRC). The HRC is the effect of a one standard deviation change in the variable of interest, while holding all other variables constant. The analysis shows that variation in the fundamental value and the execution priority yield higher HRC than variation in trading firm-level inventory levels (approximated by cumulative trade flows at NASDAQ), order monitoring costs (approximated by order size, following Fong and Liu, 2010), trades at competing venues (following van Kervel, 2015), and other market conditions not included in the theoretical model. We also investigate how the hazard rate of cancellation differs when the bid-ask spread is constrained by the minimum tick size, and find that the model predictions are largely robust to such an environment.

Our main contribution to the literature is to derive and test predictions for the determinants of limit order duration and cancellations. The predictions are based on common rather than private values, and, most importantly, the order queue position is important for the expected profit of a limit order. The queue priority result reflects the adverse-selection risk of orders in the back of the queue (as in Yueshen, 2014), as well as the expected time to execution. Our predictions and empirical results with respect to the bid-ask spread and the depth on each side of the order book are consistent with prior theoretical literature (Foucault, 1999; Parlour, 1998; Handa, Schwarz and Tiwari, 2003) and empirical studies (Biais, Hillion and Spatt, 1995; Chakrabarty, Han, Tyurin and Zheng, 2006; Hollifield, Miller, Sandås and Slive, 2006). We contribute by providing a unified theoretical framework and by showing that the model predictions dominate competing motivations for cancellations.

Heavy limit order traffic is often associated with high-frequency traders (HFTs). For example, Hendershott, Jones and Menkveld (2011) propose to use the quote intensity (relative the number of trades) as a proxy for algorithmic trading activity. Brogaard, Hendershott and Riordan (2016) show that HFTs drive price discovery by quickly updating their quotes following fundamental value changes, just like the liquidity suppliers in our model. A second contribution of this article is to show that cancellations by HFTs largely depend on the same factors as for other liquidity suppliers. We identify HFTs as members of the industry organization European Principal Traders Association (FIA-EPTA), following Baron, Brogaard, Hagströmer and Kirilenko (2016). Estimating the hazard rate of cancellation on HFT orders separately, we find that HFTs are more responsive to changes in the expected profits of their limit orders, relative other liquidity suppliers. This is consistent with the notion that HFTs invest heavily in the technology required to monitor the market at a low latency.

Our third contribution is to show that cancellations are primarily part of liquidity supply activities. Hasbrouck and Saar (2009) conclude that “fleeting orders” (orders that remain in the order book for less than two seconds) behave more like demand than supply of liquidity. They hypothesize that many cancellations are part of bundled orders whereby traders aim to detect latent liquidity, improve the execution probability of an existing order, or even to replace it with a market order. Our data set allows us to investigate these scenarios through direct observation of the chains of events at the trading firm level. We find that only 1.0% of the cancellations in our sample are consistent with the liquidity-demand related scenarios described by Hasbrouck and Saar (2009). Liquidity-supply related bundled orders, such as limit price revisions, are much more common, at 10.6% of the total (cancellations followed by a new non-marketable order at the same side of the limit order book, but at a different price). Almost half of all cancellations, 43.8%, are not at all part of a bundled order.

We acknowledge that liquidity-demanding limit orders are common in modern markets (see evidence in O'Hara, 2015), but our evidence points to that limit order cancellations are part of benign liquidity supply revisions in response to fluctuations in the fundamental asset value. Even small indications of changes in the fundamental value can trigger cancellations of marginally profitable limit orders. With continuous changes in fundamental values, manifested in frequent order book updates, it is not surprising that cancellations are also frequent. This is also consistent with predictions by Budish et al. (2015) and Foucault et al. (2016). The implication is that policies aimed at curbing the aggregate volume of cancellations are likely to restrict the ability of market makers to fine-tune their outstanding orders.

2 Theoretical Framework

The premise of this paper is that a limit order cancellation is triggered by a reduction in the profit that the liquidity supplier expects to make if the limit order is executed. Sandås (2001) models liquidity provision in a limit order market with discrete prices and time priority, building on the work by Glosten (1994) and Seppi (1997). We use the model of Sandås (2001) to derive the expected profit of a limit order as a function of its execution priority as well as the prices and volumes available at the best bid and offer. The outcome is testable predictions on the determinants of limit order durations and cancellations.

2.1 Model Setup

The market consists of a large number of liquidity demanders, who may be privately informed and submit market orders, and liquidity providers, who are risk-neutral profit-maximizers submitting limit orders. The agents trade a risky asset.

Each period t consists of three stages. Shortly before t, liquidity providers post their limit orders. At t, a liquidity demander arrives and submits a market order of size m, where the order size follows an exponential distribution with the following density function:

$$f(m) = \begin{cases} f_{mb}(m) = (2\varphi_A)^{-1} exp(-m/\varphi_A) & \text{if } m > 0 \text{ (market buy orders),} \\ f_{ms}(m) = (2\varphi_B)^{-1} exp(-m/\varphi_B) & \text{if } m < 0 \text{ (market sell orders).} \end{cases} \tag{1}$$

The parameters φ_A and φ_B are the expected sizes for market orders arriving at the ask-side and the bid-side, respectively. The liquidity demander is a buyer or a seller with equal probability 50%. After the trade, the new fundamental value of the risky asset, X_{t+1}, is announced.

In the first phase of each period, liquidity providers submit orders as long as they have a non-negative profit in expectation. Let $\tilde{\pi}_{i,t}$ denote the expected profit at execution of a limit order i posted on the ask-side of the book, and assume it takes the following form:

$$\tilde{\pi}_{i,t} = P_i - \gamma - E(X_{i,t+1} | m_t \geq \vec{Q}_{i,t}), \tag{2}$$

where P_i is the limit order price, γ is the order processing cost, and the expected fundamental value is conditional on the arrival of a market order of quantity m, large enough to execute order i. We denote the market order volume required to execute the full size of order i $\vec{Q}_i$, which includes the size of order i as well as the sizes of all orders with higher queue priority than order i. We refer to $\vec{Q}_i$ as the *Front of Queue Quantity*, and use the forward arrow notation to symbolize "the front". Below we also use the notation $\overleftarrow{Q}_i$ to describe the *Back of Queue Quantity*, where the backward arrow indicates that the variable captures order volume at the same price level but with lower time priority than order i.

Liquidity providers account for potential private information by applying the following linear function to the market order size when forming their expectation of the fundamental value:

$$E(X_{t+1}|X_t, m_t) = X_t + \mu + \alpha m_t, \tag{3}$$

where μ is the expected change in the fundamental value and α is the per unit price impact of the market order size. We henceforth drop the time index t for brevity of notation.

The model assumes that liquidity suppliers agree about the fundamental value of the risky asset. In addition, they do not have private values influencing their order submissions. In reality, however, liquidity suppliers may be heterogeneous in terms of liquidity needs, inventory costs, and monitoring costs. In addition, trading venues frequently charge their members different fees depending on their commitments and aggregate trading volume. It is possible to include private values in the model, for example by allowing the order processing cost to vary in the cross-section of liquidity suppliers. Following Sandås (2001), we abstract from private values in the theoretical model. In the empirical investigation, we show that private values are significant determinants of order cancellations, but that common values are dwarfing the effect.

Sandås (2001) shows that by using Equations (1) and (3), Equation (2) becomes:

$$\tilde{\pi}_i = \tfrac{1}{2}\left[P_i - \gamma - X - \alpha\left(\vec{Q}_i + \varphi_A\right)\right]e^{-\vec{Q}_i/\varphi_A}. \tag{4}$$

We use the expression for the expected profit of a limit order in Equation (4) as our starting point for the analysis of the determinants of limit order cancellations. In equilibrium, we express the fundamental value in terms of observable variables.

2.2 Equilibrium

Sandås (2001) derives the following break-even condition for limit orders on the ask-side of the order book:

$$Q_A = \frac{P_A - X - \gamma}{\alpha} - \varphi_A, \tag{5}$$

where P_A is the ask-side price, and Q_A is the depth beyond which liquidity providers have no incentive to post additional orders at that price. The corresponding condition for the bid-side is:

$$Q_B = \frac{X - P_B - \gamma}{\alpha} - \varphi_B. \tag{6}$$

Up until this point, we follow Sandås' (2001) model exactly. We now take the model in a different direction, as we seek to express the expected profit of a limit order as a function of observable quantities. Although we do not alter any of the assumptions of the original model, the following results are new.

To obtain an expression for the fundamental value, we solve Equation (6) for α, insert the resulting expression in Equation (5), and solve for X to get:

$$X = \frac{(P_A - \gamma)(Q_B + \varphi_B) + (P_B + \gamma)(Q_A + \varphi_A)}{Q_B + \varphi_B + Q_A + \varphi_A}. \tag{7}$$

Equation (7) shows that the fundamental value in equilibrium is a weighted average of the best bid and ask prices, with the imbalance between ask-side and bid-side depth (adjusted for expected market order size and order processing costs) determining the weights. The concept of depth-weighted prices to proxy fundamental value is common in the financial industry and sometimes goes under the name "the microprice" (Harris, 2013; Hagströmer, 2017).

By inserting Equation (7) in place for X in Equation (4), and noting that $P_A = P_i$ holds at the execution of the ask-side order i, we obtain:

$$\tilde{\pi}_i = \frac{1}{2}\left[\frac{(P_i - P_B - 2\gamma)(Q_A + \varphi_A)}{Q_B + \varphi_B + Q_A + \varphi_A} - \alpha(\vec{Q}_i + \varphi_A)\right] e^{-\vec{Q}_i/\varphi_A}. \tag{8}$$

Equation (8) captures the role of the queue position of order i by distinguishing between the full depth Q_A (at the price level in question) and the market order size required to execute the order, the *Front of Queue Quantity* ($\vec{Q}_i$). The difference between the two is the *Back of Queue Quantity* ($\overleftarrow{Q}_i$).

We obtain an expected profit expression for bid-side limit orders analogously. To set up a general expression for the expected profit of a limit order, we adopt the following terminology. Define *Bid-Ask Spread* (BAS_i) as the signed difference between the limit price and the best price on the opposite side of the book, conditional upon execution of order i. Furthermore, define the *Opposite-Side Quantity* (Q^{Opp}) as the volume of the best price level on the opposite side of the order book (from the perspective of order *i*), and let the parameters φ and φ^{Opp} represent the expected sizes of market orders arriving on the same side and opposite side of the book. Using this notation, the expected profit for order *i* is:

$$\tilde{\pi}_i = \frac{1}{2}\left[\frac{(BAS_i - 2\gamma)\left(\vec{Q}_i + \overleftarrow{Q}_i + \varphi\right)}{\vec{Q}_i + \overleftarrow{Q}_i + \varphi + Q^{Opp} + \varphi^{Opp}} - \alpha\left(\vec{Q}_i + \varphi\right)\right] e^{-\vec{Q}_i/\varphi}. \tag{9}$$

Because the model parameters are constant, all variation in the expected profit is due to fluctuations in the *Bid-Ask Spread*, the *Opposite-Side Quantity*, the *Back of Queue Quantity*, and the *Front of Queue Quantity*. We refer to these four variables as the “model variables”.

2.3 Empirical Predictions

We form dynamic empirical predictions by comparing equilibrium liquidity before and after an order book event. The model is static and does not convey how the limit order book reaches its equilibrium. Our approach to the dynamics of the limit order, similar to that of van Kervel (2014), is based on that changes in the fundamental value trigger cancellations of existing limit orders and submissions of new limit orders. Our predictions concern the limit order duration and the hazard rate of cancellation (the cancellation intensity over the next instant).

We postulate that the duration of a limit order (controlling for executions) depends positively on its expected profit. Our reasoning relies on the model assumption that limit orders remain in the order book as long as the expected profit is non-negative. A limit order with a high, expected, profit is more resilient to fundamental value changes than one that is only marginally profitable in expectation. Along the same line of reasoning, we assume that changes in the expected profit of a limit order negatively affect the hazard rate of cancellation.[1]

We use Equation (9) to form empirical predictions on how each of the model variables influences the duration and the hazard rate of cancellation for a limit order.

- The variable *Bid-Ask Spread* (BAS_i) is positively related to the limit order duration, and changes therein (ΔBAS_i) are negatively related to the hazard rate of cancellation.

- The variable *Opposite-Side Quantity* (Q^{Opp}) is negatively related to the limit order duration, and changes therein (ΔQ^{Opp}) are positively related to the hazard rate of cancellation. The fraction within squared brackets in Equation (9), which is a depth imbalance ratio, is strictly positive and the variable Q^{Opp} enters only in the denominator of that fraction.[2]

- The variable *Back of Queue Quantity* ($\overleftarrow{Q}_i$) is positively related to the limit order duration, and changes therein ($\Delta\overleftarrow{Q}_i$) are negatively related to the hazard rate of

[1] The spirit of the model is that the probability of cancellation is binary; zero for orders with non-negative expected profit and one otherwise. If we would know the model parameters, we could use Equation (9) to identify orders with negative expected profits following a change in the state of the order book. The frequency of cancellation of such orders would potentially indicate the merit of the model, or the prevalence of private values in liquidity supply. Sandås (2001) shows how to estimate the parameters of the model. Our aim here, however, is to understand the determinants of cancellations. We do not investigate a threshold effect, whether a trader cancels a limit order with negative expected profits. Instead, we investigate a continuous relation, whether changes in the expected profit negatively affects the hazard rate of cancellation.
[2] This holds because the expected market order sizes φ and φ^{Opp} are strictly positive, and the bid-ask spread is always greater than twice the order processing cost (otherwise, liquidity suppliers would quote zero volumes at those prices).

cancellation. The intuition is that a longer queue behind the order of interest implies a lower adverse selection cost for the order.

- The effect of the *Front of Queue Quantity* ($\vec{Q}_i$) on the limit order duration is lower than that of the *Back of Queue Quantity*. The net effect is however unclear. The *Front of Queue Quantity* influences the expected profit positively through its inclusion in the depth imbalance ratio, but it has a negative influence through the price impact term $\alpha\left(\vec{Q}_i + \varphi\right)$, and the exponential function $e^{-\vec{Q}_i/\varphi}$. The price impact term captures adverse selection costs, which are increasing in market order size. The exponential function captures the expected time to execution, as the ratio $\vec{Q}_i/\varphi$ is a proxy for the number of (same-side) market order arrivals required to execute order i. By the same reasoning, the effect of changes in the *Front of Queue Quantity* ($\Delta\vec{Q}_i$) on the hazard rate of cancellation is higher (less negative) than that of changes in the *Back of Queue Quantity*.

The difference in predicted effects of volume with higher and lower execution priority shows the importance of considering an order's queue position. Previous studies frequently consider total same-side depth ($\vec{Q}_i + \overleftarrow{Q}_i$) to have a negative influence on the probability of limit order cancellations (see empirical evidence by Chakrabarty et al., 2006; and Hollifield et al., 2006). To our knowledge, the role of the order queue position is new to the literature.

3 Data and Sample

3.1 Data

We conduct the study with proprietary data from NASDAQ. The data set from NASDAQ consists of all messages entered into the trading system INET, including limit order submissions, cancellations, modifications and executions. We use data from all 19 trading days in May 2014, a month without changes to the trading system. For each limit order submission, we observe the time of entry, quantity, limit price, visibility conditions, and time in force. Through an order sequence number, which connects order submissions to cancellations, modifications and/or executions, we follow the life of every limit order. We reconstruct the state of the order book for each stock throughout the trading day.

All order submissions in the NASDAQ data set contain a trading firm identifier. A trading firm is an entity that connects to NASDAQ either as an exchange member or through another exchange member as a sponsored access client. We identify a subset of the trading firms as HFTs, based on the member list of FIA-EPTA, following Baron et al. (2016).

We also retrieve trade records and order book data at the microsecond frequency from Thomson Reuters' Tick History (TRTH) database, maintained by the Securities Research Centre of Asia-Pacific (SIRCA). We use the TRTH data to measure trading activity at other trading venues than NASDAQ. During our sample period, NASDAQ has roughly 60% of the lit trading volume in the sample stocks. The main competitors are BATS Chi-X CXE (19%), Turquoise (10%), and BATS Chi-X BXE (8%).

3.2 Institutional Detail

NASDAQ operates an electronic limit order book market which is open from 9:00 AM to 5:30 PM every weekday except Swedish bank holidays. Trading closes at 1:00 PM on trading days followed by a Swedish bank holiday. Call auctions determine opening and closing prices.

Trading is allowed at any price on a grid determined by the minimum tick size, which depends on the stock price level. Specifically, the tick size is 0.01 SEK for stocks priced below 50 SEK, 0.05 SEK for stocks priced between 50 SEK and 100 SEK, and 0.10 SEK for stocks priced between 100 SEK and 500 SEK. The tick size rule is important as it puts restrictions on, e.g., the possibilities to compete for order flow by posting aggressive orders inside the bid-ask spread. When that is not possible, we say that the tick size is binding. In our sample, the tick size is binding 69% of the time.

The execution priority of limit orders at NASDAQ is set according to price-internal-visibility-time. Internal priority means that market orders are matched primarily to limit orders posted by the same trading firm, if such orders exist at the best price level. That is, a limit order posted at the best price by another trading firm may not get the trade even if it is first in line in terms of time priority. The internal matching may affect traders' decision to cancel, and we control for that in our empirical analyses.

The use of hidden liquidity is highly restricted at NASDAQ. For the stocks in our sample, an order must exceed at least one million EUR to be eligible for full non-visibility (for some of the stocks the threshold is even higher). This restriction is important not only for order management (to hide or not to hide orders) but also for order aggressiveness. Hautsch and Huang (2012) report that, in the U.S. equity markets, hidden liquidity is associated with *"enormous order activities"* (p.2) related to liquidity-detection strategies, which are likely to involve aggressive market orders. The restrictions on hidden liquidity at NASDAQ is likely to induce less liquidity-detection strategies relative, e.g., activities on the U.S. equity markets.

3.3 Sample Construction

We analyze orders at the best bid and ask prices. Both Hasbrouck and Saar (2009) and Chakrabarty and Tyurin (2011) report such orders to be relatively short-lived. Our sample procedure is however not restricted to short-lived orders. The analysis is in this sense more general than that of Hasbrouck and Saar (2009), who restrict much of their analyses to orders cancelled within two seconds.

At any given moment, we track at most one order in each stock. Specifically, for each stock-day, we randomly sample one order posted at the best prices in the order book at 9:05 AM. We track that order over at most 100 order book events, where a limit order book event is a change to either the prices or volumes at the best price levels of the order book. After 100 order book events have elapsed, we sample a new order. We repeat this procedure until the end of the continuous trading.[3] We allow sampling of hidden orders, but we do not consider them in the subsequent analyses.

The sampled orders have three possible outcomes: cancelled, executed or still alive after 100 events. In the model, an order is always for one unit of a stock. When we take the model to

[3] We allow for orders to be re-sampled, but we omit these orders in the analyses.

the data, it is important to note that the predictions are only valid for the last share of the order.[4] In reality, orders are typically for larger quantities, and cancellations and executions may be partial. To be true to the model predictions, we ignore partial cancellations and executions and record the outcome for the last unit of each order.

There are a few more notes to make about the sampling procedure. The use of event time in our sampling procedure implies that we do not sample orders at the same clock-time in different stocks, which is important in order to satisfy the requirement of independent observations in the statistical analysis. Furthermore, by sampling from existing orders, we do not typically analyze orders at the time of their submission. In this way, we avoid orders labeled "immediate-or-cancel", that by definition have zero duration. Finally, the sampling implies that stock-days with more order book activity have more observations in our sample. Hence, one way to interpret our results is that they are order activity-weighted across stock-days.

3.4 Descriptive statistics

Table 1 presents descriptive statistics for our model variables at the time of sampling.[5] For comparability across orders, we normalize the model variables. We divide the variable *Bid-Ask Spread* by the spread midpoint. In addition, we divide each of the three quantity variables, *Opposite-Side Quantity*, *Back of Queue Quantity*, and *Front of Queue Quantity*, by the total depth at the best bid and ask prices. The sampling procedure provides us with 138,487 limit orders.

[4] Note that we define the *Front of Queue Quantity* as the volume ahead of the order in question, plus the quantity of the order itself.
[5] Appendix A presents summary statistics for the individual stocks in our sample.

The outcomes are 53% cancellations, 21% executions, and 26% of the orders are still alive after 100 events.

INSERT TABLE 1 ABOUT HERE

The first column of Table 1 presents means of the variables. At the time of sampling, the average *Bid-Ask Spread* is 7.29 basis points. The limit order book is on average symmetric in that *Opposite-Side Quantity* accounts for 50% of total depth and the quantity on same side accounts for the remaining 50% of total depth (*Back of Queue Quantity* plus *Front of Queue Quantity*). The sampled *Order Quantity* accounts for 8% of total depth. Since *Front of Queue Quantity* includes the *Order Quantity* it constitutes a higher share of total depth (29%) than the *Back of Queue Quantity* (21%).

Table 1 also presents descriptive statistics on *Order Duration* conditional on outcomes. We define *Order Duration* as the time in seconds to each of the outcomes *Cancellation*, *Execution*, or *Orders Surviving 100 Events*. The mean time to *Cancellation* (26.05 seconds) is shorter than the mean time to *Execution* (38.77 seconds). The mean time for the *Orders Surviving 100 Events* is 101.96 seconds. This implies that we track the sampled orders for, on average, a little more than one and a half minutes. The percentile statistics shows that many orders are short-lived. Of the cancellations, 25% have a shorter duration than 0.05 seconds, and 25% of the executions have a shorter duration than 0.85 seconds.

4 Analysis of Order Durations

In this section, we analyze limit order durations in relation to the predictions presented in Section 2.3. In general, the prediction is that orders with higher expected profit have longer expected duration, because they are more resilient to changes in the fundamental value.

Figure 1 displays the cumulative probability of time to cancellation and time to execution of the orders, expressed in the number of events (Panel A) and time (Panel B). The curves show that both the probability of cancellation and the probability of execution are decreasing functions of time (measured by the number of events and by clock time). Evidently, the cumulative cancellation probability is diminishing fast.

INSERT FIGURE 1 ABOUT HERE

Table 2 presents limit order durations of the sampled orders. We report both unconditional durations, and durations conditional on whether the model variables and selected control variables are "Large" (above the median) or "Small" (below the median). We report durations separately for cancellations and executions.

INSERT TABLE 2 ABOUT HERE

The first row of Table 2 shows the unconditional order durations. Cancelled orders have an average duration amounting to around 24 events or about 26 seconds. For executions, the corresponding average durations are 26 events or almost 39 seconds.

The conditional duration analysis lends support to our empirical predictions regarding the model variables. In Table 2, the duration of cancelled orders is increasing with the *Bid-Ask*

Spread and the *Back of Queue Quantity*, and decreasing with the *Opposite-Side Quantity*. The differences between the "Large" and "Small" values for these categories are all statistically significant at the 1% confidence level, according to a *t*-test of the hypothesis that each duration measure is the same across subsamples. The fourth variable, the *Front of Queue Quantity*, also shows an increasing relation to the limit order duration. Our model does not provide a prediction of the direction of the effect of the *Front of Queue Quantity*, but it predicts that the effect is weaker than that of the *Back of Queue Quantity*.

In addition, from Table 2 we note that average durations (in both seconds and number of events) to both cancellation and execution are significantly shorter when the *Order Quantity* is larger than the median rather than smaller than the median. We also investigate the effect of the *Internal Match Rate*, measured for each trading firm as the proportion of passively executed limit orders that are internally matched during April 2014, the month preceding our sample month. We find that trading firms with higher *Internal Match Rate* have significantly longer order duration before cancellation and execution. These effects are, however, significant only when measuring order duration in seconds, not in events. Moreover, the duration to cancellation (and execution) is significantly shorter when the tick size is binding (at the time of sampling the order) compared to when it is not binding. Finally, HFTs cancel, and get execution of, their limit orders significantly quicker than non-HFTs.

5 Analysis of Order Cancellations

So far, we perform the expected duration analysis ex ante, based on market and order conditions holding at the time when we sample each order. To understand the triggers for limit order cancellations, we now turn to an analysis of what happens *after* the time of sampling and how that influences the hazard rate of cancellation. Specifically, for each limit order we analyze the time series of up to 100 limit order book events. The hazard rate for order i is the intensity of cancellation over the next instant, which we denote as $\lambda_i(s)$, with s indicating the event time.[6] The time of sampling for each order is set to $s = 0$, and the maximum event time for each order is $s = 100$.

In addition to our expectation that orders with higher expected profits have lower hazard rates of cancellation, we incorporate the prediction from Section 2 that the hazard rate of cancellation relates negatively to changes in expected profit. The predictions are summarized as:

$$\lambda_i(s) = f\left(\tilde{\pi}_i, \Delta\ln\left(\tilde{\pi}_{i,s}\right)\right), \tag{10}$$

where $\Delta\ln(\tilde{\pi}_{i,s})$ is the log change in expected profit from time 0 to time s.

We approximate the log change in expected profit with a linear Taylor expansion of Equation (9), in which we assume that the parameters γ, φ, and φ_{Opp} remain constant over the lifetime of the order. The order price and side (P_i and D_i) are constant by definition. The Taylor expansion is then a function of changes in the four model variables in Equation (9):

[6] According to Lee and Wang (2003), the hazard rate relates to the survival function of the order, which for a cancellation time T is $S(s) = Pr(T > s)$. The hazard rate is $\lambda_i(s) = -\mathrm{d}\log\big(S(s)\big)/\mathrm{d}s = S(s)^{-1}(-S'(s))$, where $S'(s)$ is the derivative of S with respect to s.

$$\Delta\ln(\tilde{\pi}_{i,s}) = \underbrace{\frac{\partial\ln\left(\tilde{\pi}_{i,s}\right)}{\partial(BAS_{i,s})}(\Delta BAS_{i,s})}_{\substack{\textit{Bid-Ask Spread} \\ \textit{Change}\text{ effect}}} + \underbrace{\frac{\partial\ln\left(\tilde{\pi}_{i,s}\right)}{\partial Q_s^{Opp}}\Delta Q_s^{Opp}}_{\substack{\textit{Opposite-Side} \\ \textit{Quantity Change} \\ \text{effect}}} + \underbrace{\frac{\partial\ln\left(\tilde{\pi}_{i,s}\right)}{\partial\overleftarrow{Q}_{i,s}}\Delta\overleftarrow{Q}_{i,s}}_{\substack{\textit{Back of Queue} \\ \textit{Quantity Change} \\ \text{effect}}} + \underbrace{\frac{\partial\ln\left(\tilde{\pi}_{i,s}\right)}{\partial\overrightarrow{Q}_{i,s}}\Delta\overrightarrow{Q}_{i,s}}_{\substack{\textit{Front of Queue} \\ \textit{Quantity Change} \\ \text{effect}}}. \qquad (11)$$

5.1 Empirical Model

We analyze the relation in Equation (10) by estimating a PDHM with time-varying covariates. This type of model allows us to analyze the hazard rate of cancellation in relation to the model variables in levels as well as changes, while controlling for executions (which we assume to be an exogenous censoring process). The model is:

$$\lambda_i(s) = \lambda_0(s)\exp\left[Model_{i,0}\theta + \Delta Model_{i,s}\beta + \Delta Other_{i,s}\delta + Control_{i,0}\gamma\right], \qquad (12)$$

where $\lambda_0(s)$ is an unspecified baseline hazard rate. The vector $Model_{i,0}$ contains the model variables in levels, measured at the time of sampling and constant in the time dimension. The purpose of this vector is to capture the level effects of expected order durations. The vector $\Delta Model_{i,s} = Model_{i,s} - Model_{i,0}$ contains changes in the model variables. They are time-varying covariates, meaning that they vary both across orders and across time; from the time when the order is sampled until it is either cancelled or executed, or until 100 time periods have elapsed. The coefficient vectors corresponding to $Model_{i,s}$ and $\Delta Model_{i,s}$ are θ and β.

Because the three quantity variables in $Model_{i,0}$ sum to unity, the model cannot be estimated if all of them are included. For this reason, we exclude the *Opposite-Side Quantity* in all estimations. The same issue does not apply to the $\Delta Model_{i,s}$ variables.

The vector $\Delta Other_{i,s}$ contains three time-varying covariates that are not part of our theoretical model but that may influence cancellation: *Firm Inventory, Same-Side MTF Trades,* and *Opposite-Side MTF Trades*, each presented below. The coefficients corresponding to $\Delta Other_{i,s}$ are represented by δ.

Hendershott and Menkveld (2014) show how inventory constraints introduce private values in the liquidity supply. Raman and Yadav (2014) find support for inventory levels influencing limit order management. The variable *Firm Inventory* captures changes in inventory at the trading firm-level, measured from the time of sampling until time *s*. We approximate inventory changes with the cumulative trading volume at NASDAQ, in the stock for which order *i* is posted, by the trading firm that posts order *i*. The cumulative trading volume is the number of shares and, if order *i* is on the ask-side, is positive for buy trades and negative for sell trades. If order *i* is on the bid-side, the trading volume is signed in the opposite direction. Should *Firm Inventory* be important for order cancellation, the variable should have a positive relation to the hazard rate of cancellation.

The model by Sandås (2001) assumes an economy where the risky security is traded in a centralized limit order book. However, our empirical context is decentralized, and the trading activity in our sample stocks is fragmented across both lit and dark trading venues. Van Kervel (2014) integrates the model by Sandås (2001) with the two-venue setting by Foucault and Menkveld (2008), and shows that trades at a competing venue work as signals about changes in the fundamental value. For example, Van Kervel (2014) predicts that a buy trade at a competing venue signals an increase in the fundamental value, and should accordingly lead to cancellations of ask-side orders and submissions of bid-side orders. We measure trading volume at the three multilateral trading platforms (MTFs) that are the main competitors to NASDAQ for lit order flow in Swedish stocks, BATS Chi-X CXE, Turquoise, and BATS Chi-X BXE. The variables *Same-Side MTF Trades* and *Opposite-Side MTF Trades* measure the cumulative trading volume recorded at the three MTFs from the time of sampling until time s, on the same side and the opposite side as the order in question. We expect the *Same-Side MTF Trades* to have a positive relation, and the *Opposite-Side MTF Trades* to have a negative relation, to the hazard rate of cancellation.

Finally, the vector $Control_i$ contains four control variables observed at the time of sampling, and not time-varying, with coefficients γ. It includes *Lagged Absolute Return*, *Lagged Volume*, *Order Quantity*, and *Internal Match Rate*. The variable *Order Quantity* captures the effect of order monitoring costs on cancellations. Liu (2009) predicts that because the pick-off risk is greater for larger orders, the incentives to monitor a limit order increases with the order size. We expect *Order Quantity* to have a positive relation to the hazard rate of cancellation. The *Internal Match Rate* controls for the effect of internal matching on the cancellation decisions. We expect firms with a higher propensity for internal matching to have longer order duration, implying a negative relation between the hazard rate of cancellation and the *Internal Match Rate*.

All the variables in the vectors $\Delta Other_{i,s}$ and $Control_i$ are standardized within each stock to have zero mean and unit variance.

5.2 Results for the Full Sample

We report the coefficient estimates of the PDHM in Table 3. To facilitate the analysis of the economic significance, we also include the HRC of each variable. Similar to marginal effects in probability models, the HRC captures the change in the cancellation hazard rate following a one-standard deviation change in the explanatory variable in question, while holding all other explanatory variables constant.[7]

INSERT TABLE 3 ABOUT HERE

The results for the model variables in levels are comparable to the order duration analysis above. However, the PDHM model considers the effect of all model and control variables simultaneously, and controls for the censoring process of executions. The results show that the hazard rate of cancellation has a negative relation to the *Bid-Ask Spread* and the *Back of Queue Quantity*. This is consistent with our predictions and the order duration results, and the results are statistically significant at the 1% level. The effect of the *Front of Queue Quantity* variable is also negative, which is consistent with the results in Table 2, but it is only significant at the 5% level. The difference in results indicate the importance of controlling for other variables and the censoring process when analyzing the hazard rate.

For the time-varying changes, the results are consistent with our empirical predictions for all the model variables. The hazard rate is decreasing with the *Bid-Ask Spread Changes* and the *Back of Queue Quantity Changes*, and it is increasing with *Opposite-Side Quantity Changes*. The coefficients are statistically significant at the 1% level. The variable *Front of Queue Quantity Changes* has a marginally significant positive coefficient.

[7] To obtain the HRC of an explanatory variable $X_{i,s}$ with an estimated coefficient β_X and standard deviation σ_X, we first calculate the hazard ratio, which is $e^{\beta_X \sigma_X}$. The hazard rate change is the hazard ratio minus one, reported as a percentage.

The most important result is that the *Front of Queue Quantity* and the *Back of Queue Quantity* coefficients differ significantly, both in economic and statistical terms, and both in levels and in changes. Economically, the hazard rate change for the *Back of Queue Quantity* is −35.7%, while for the *Front of Queue Quantity* it is only −0.8%. The corresponding numbers for the time-varying variables are −48.3% and 3.0%. To determine the statistical significance of the difference, we perform Wald tests. The null hypothesis for the level variables, that the coefficients of $\overleftarrow{Q}_{i,s}$ and $\vec{Q}_{i,s}$ are equal, is rejected at any level of confidence. The same results hold for the hypothesis stating that the coefficients of $\Delta\overleftarrow{Q}_{i,s}$ and $\Delta\vec{Q}_{i,s}$ are equal.[8] The results highlight the importance of the order queue position.

Notably, the *Front of Queue Quantity* variables contribute to the depth imbalance (which is positively related to the expected profit of the order, see Equation (9)), and also to the time to execution and the adverse selection costs (that are negatively related to the expected profit). Our finding that the *Front of Queue Quantity* coefficients are close to zero shows that the effects of these elements are in balance for the overall sample. The subsample analysis below, however, indicates that this result depends on whether the tick size is binding and that it differs across trader groups.

The PDHM includes time-varying covariates and level control variables that are not present in the theoretical model. The coefficient estimates for those variables are all consistent

[8] The test statistics are available from the authors upon request. We do not report the tests in the tables, but the null hypothesis of equal coefficients is rejected for both levels and changes in all specifications of the PDHM considered below.

with our expectations. The variables *Firm Inventory* and *Lagged Volume* are however not statistically significant.

We find the *Bid-Ask Spread* and the *Back of Queue Quantity*, expressed in both levels and in changes, to be the dominant variables in terms of HRC. For example, a one standard deviation increase in the *Back of Queue Quantity Changes* decreases the hazard rate of cancellation by 48.3%. Among the variables that are not part of our theoretical model, the most important in terms of HRC is the *Order Quantity*, which according to Liu (2009) captures the cost of order monitoring.[9]

Private values as captured by *Firm Inventory* have a relatively small effect on the hazard rate, lending support to our focus on common values in the theoretical model. Furthermore, the low HRC associated with the variables capturing common values indicated at the competing trading venues (*Same-Side MTF Trades* and *Opposite-Side MTF Trades*) shows that the theoretical assumption of a central limit order book is empirically relatively benign, at least in the context of analyzing the hazard rate of cancellation. The negative coefficient on the *Internal Match Rate* indicates that trading firms with a high propensity to internal matching have a low hazard rate of cancellation, which is consistent with them putting less weight on the time priority of their limit orders.

[9] The *Front of Queue Quantity* has a mechanical positive relation to *Order Quantity*, which could potentially introduce multicollinearity problems. However, the sample correlation is only 0.21, making it unlikely that multicollinearity drives down the significance of the *Front of Queue Quantity* coefficient.

We conclude that all our predictions with respect to the model variables find empirical support in the PDHM. The results are robust to alternative explanations for cancellation, such as private values, monitoring costs, trading venue competition, and internal matching.

5.3 Results for Sub-Samples Based on Tick Size

Price discreteness constrains price competition. O'Hara, Saar and Zhong (2015) show that the relative tick size influences the biodiversity of traders as well as the limit order behavior. Their empirical investigation shows that a larger relative tick size leads to greater quoted depth, and greater profitability of HFT market making. The result with respect to fast traders is consistent with Budish et al. (2015), who argue that price discreteness lead to competition on latency.

To see how the cancellation hazard rates depend on price discreteness, we split the sample by whether a stock is trading under a binding tick size or not. In the subsequent analysis we also split the sample by the limit order supplier type, comparing HFTs to Non-HFTs.

We determine whether the tick size is binding for each order at the time of sampling. If at that time the bid-ask spread is equal to the minimum tick size, we consider the tick size to be binding. We rerun the PDHM analysis separately for sample orders with binding and non-binding tick size. Table 4 holds the coefficient estimates and associated HRCs .

INSERT TABLE 4 ABOUT HERE

When the tick size is binding (presented to the left in Table 4), the results for the *Back of Queue Quantity* (both level and changes) and *Opposite-Side Quantity Changes* are consistent with our findings for the entire sample in Table 3. For the *Bid-Ask Spread* the result is negative in the level, which is consistent with the prediction, but positive for the time-varying changes. The *Front of Queue Quantity* is significantly negative, both in the level and in the time-varying changes.

Our interpretation of the deviations from the full-sample results is that when the tick size is binding, variation in the expected time to execution is reflected more by the changes in the spread than by the *Front of Queue Quantity*. To see this, note that a widening spread may be indicating that the price is moving away from the limit order in question, making its expected time to execution higher. When the tick size is binding, this effect is likely to be stronger, because a one-tick change implies a doubling of the spread. The positive coefficient on the *Bid-Ask Spread Changes* implies that the expected time to execution effect is stronger than the effect of increasing expected revenue at execution. The negative net effects of the *Front of Queue Quantity* variables are consistent with that the effect of the expected time to execution, when the tick size is binding, is reflected more by the spread variation than by the *Front of Queue Quantity*.

For the sample orders where the tick size is not binding, the results (reported to the right in Table 4) are more consistent with the full-sample results. Notably, the effects of the *Front of Queue Quantity* variables are positive, providing further support for the notion that the time to execution is a more important part of the front of queue effect when the tick size is not binding.

Overall, regardless of whether the tick size is binding or not, the hazard rate changes show that the economic significance for the model variables is higher than for the variables outside the model.

5.4 Differences in Results between HFTs and Non-HFTs

Frequent limit order revisions and cancellations are commonly associated with HFTs. Hagströmer and Nordén (2013) and Menkveld (2013) show that HFTs to a large extent engage in market making, consistent with the liquidity suppliers in our model, and that the market-making HFTs have higher order-to-trade ratios than other HFTs. Brogaard, Hendershott and Riordan (2016) find that HFTs lead the price discovery process by updating their quotes at high speed when new information emerges in the marketplace. Brogaard et al. (2015) show that market-making HFTs have a strong tendency to invest in colocation services allowing them to monitor and respond to changes in the fundamental value with low latency.

The role of HFTs in the modern equity market motivates us to study their quoting behavior relative to that of other liquidity suppliers. We proceed by splitting the sample into four parts: orders with binding tick size submitted by HFTs and by other liquidity suppliers (Non-HFTs), respectively; and orders with a non-binding tick size submitted by HFTs and Non-HFTs. We estimate the PDHM for each subsample and present the results in Table 5. Panel A holds results for the subsamples with the binding tick size and Panel B reports the findings for the subsamples where the tick size is non-binding.

INSERT TABLE 5 ABOUT HERE

Two overall observations emerge from this analysis. First, the signs of the model variable coefficients are largely the same for HFTs and Non-HFTs, regardless of whether the tick size is binding or not. Second, the economic influence of the model variables (in terms of HRCs) tends to be greater for HFTs than for other liquidity suppliers. Hence, HFTs adhere more strongly to the economics conveyed by the model.

Perhaps the most important difference between HFTs and Non-HFTs is observed for the *Back of Queue Quantity* variables. For example, when the tick size is binding, the HFT coefficient on *Back of Queue Quantity Changes* is −4.100, whereas for Non-HFTs it is −0.575. The corresponding HRCs are −87.6% versus −26.2%. The difference persists for the levels of *Back of Queue Quantity* as well as for the subsamples where the tick size is not binding.

The strong emphasis of *Back of Queue Quantity* among HFTs signals that they use the depth imbalance to track the fundamental value, and that they do so to a larger extent than other liquidity suppliers. As discussed above, the depth imbalance effect is also reflected by the *Front of Queue Quantity*. Accordingly, HFTs have lower coefficients (more negative or less positive) than Non-HFTs for the *Front of Queue Quantity* variables, both for the level and changes, and regardless of the tick size.

Our results are consistent with HFTs investing heavily in infrastructure allowing them to monitor the limit order book at high frequency. According to the model view of the expected profit of a limit order, HFTs respond relatively strongly to factors that may undermine the profitability of their limit orders.

We also find support for the common notion that HFTs manage their inventory risk at high frequency, demonstrated empirically by Brogaard et al. (2015). For example, when an HFT increases his long position in a stock, he increases the hazard rate of cancelling buy-side limit orders, and decreases the hazard rate of cancelling sell-side limit orders (see the coefficient estimate for *Firm Inventory*). This is consistent with the price pressure behavior described by Hendershott and Menkveld (2014), who show that market makers adjust their limit order execution probabilities in response to their level of inventory. We do not find any effect of *Firm Inventory* for Non-HFTs.

6 Cancellations as Part of Bundled Orders

The starting point of this paper is that the orders in the limit order book originate from liquidity suppliers who maximize profits earned as intermediaries and have no trading needs per se. In contrast, Hasbrouck and Saar (2009) find that short-lived orders have characteristics consistent with the demand of liquidity. In line with the view that limit orders may represent liquidity demand, O'Hara (2015) shows that in a data set of institutional trades executed through the brokerage firm ITG, almost two thirds of the executions are passive. In this section, we seek to determine the prevalence of liquidity demand strategies relying on passive limit orders in our data.

Hasbrouck and Saar (2009) hypothesize that many cancellations come from liquidity demanders who try to detect latent liquidity, improve the execution probability of an existing order, or even to replace it with a market order. They suggest that such cancellations are part of bundled orders, that is, a combination of an order submission and an order cancellation. For example, at NASDAQ, the price of an existing order can be revised by submitting a "replace" order. The replace order bundles a cancellation of the existing limit order with a simultaneous submission of a new order of the same quantity at a different limit price.

Bundled orders may be part of either liquidity supply or liquidity demand strategies. Hasbrouck and Saar (2009) present three types of liquidity demand related bundled order types (presented below). They present indirect empirical evidence consistent with such order types.

Our data allows direct observation of bundled orders. We investigate such orders as a way to determine whether liquidity demanders or suppliers populate the limit order book.

To identify bundled orders we match each limit order cancellation in our sample to the next message sent to the market in the same stock, by the same trading firm. We postulate that when there is a matching message within half a second after the cancellation, it is potentially a bundled order. By allowing a delay in the subsequent order we allow bundled orders to be either automatic, such as the replace message described above, or manual. In Table 6, we show how common different types of messages bundled with cancellations are.

INSERT TABLE 6 ABOUT HERE

We categorize messages as *Nothing* (when there is no message matched to the cancellation), *Cancel* (when the cancellation is followed by another cancellation), *Insert* (when the cancellation is followed by a limit order submission), or *Execute* (when the next event for the trading firm in question is a passive execution). Arguably, out of these categories, only *Insert* is consistent with bundled orders, and we focus the discussion on that category. We find that 35.4% of the 73,458 cancellations in our sample are followed by the insertion of a new limit order, that is, potentially part of bundled orders.

To determine whether the bundled orders relate to liquidity demand or supply, we break down the *Insert* category in two dimensions. First, we subcategorize the subsequent limit order submissions by their relation to the state of the order book; whether they are on the same or the opposite side relative the cancelled order, and whether they are more, less or equally aggressively priced. In the second dimension, we condition the frequencies of bundled orders on changes in opposite price over the half second preceding the cancellation (*Moves Away*, *Moves*

Closer, and *Unchanged*). Hasbrouck and Saar (2009) hypothesize that if the fundamental value is increasing, a liquidity demander may revise the price of his limit buy orders upwards, in order to maintain the same execution probability as before the value change (the "chasing hypothesis"). Furthermore, if the bid-ask spread shrinks, they hypothesize that liquidity demanders may replace their passive limit orders with marketable limit orders, because the price of immediacy is then lower (the "cost-of-immediacy hypothesis"). The breakdown of the *Insert* category enables a direct investigation of how likely it is that the demand-related hypotheses according to Hasbrouck and Saar (2009) can explain the order cancellations.

The results in Table 6 show that bundled orders consistent with the "chasing" hypothesis, i.e., when subsequent to a cancellation and conditional on that the opposite price moves away, the trading firm inserts a new limit order at a better price, represent 0.9% of all cancellations in the sample. Moreover, bundled orders in line with the "cost-of-immediacy" hypotheses, when a trading firm cancels an order and, conditional on that the opposite price moves closer, the trading firm inserts a same-side marketable limit order, represent only 0.1% of all cancellations. Hence, out of the total cancellations in the sample, only about 1.0% are related to the two liquidity-demand related hypotheses.[10]

The results presented in this section point to that the liquidity demand-related strategies discussed by Hasbrouck and Saar (2009) are relatively rare in our sample. There may of course

[10] Hasbrouck and Saar (2009) also present a third liquidity demand-related strategy. The "search hypothesis" involves orders that search for hidden liquidity through the submission of orders at prices inside the prevailing bid-ask spread. These orders are typically "immediate-or-cancel", meaning that if they do not strike hidden liquidity, they are immediately withdrawn. Given that our sampling procedure randomly selects orders that rest in the order book, this order type is mechanically excluded.

be other demand-related strategies influencing cancellations that the bundled order analysis cannot capture. However, the empirical results in Section 4 and 5 strongly support the predictions coming out of a model with liquidity suppliers behaving as market makers. We regard that as a strong indication of such agents dominating the liquidity supply dynamics.

7 Conclusions

We derive empirical predictions for the determinants of limit order durations and cancellations. The predictions rely on limit orders being free options offered to the market in the expectation of intermediation profits. We postulate that limit orders with high, expected, profit have long expected duration, and that a reduction in the expected profit triggers cancellations. A novel prediction, driven by the presence of adverse selection risk, is that the order queue position (its time priority) is an important determinant of the expected duration, and that changes therein can trigger cancellations. The predictions find strong support in the data.

Overall, our study points to that frequent order cancellations are a benign feature of modern market making. Competitive liquidity suppliers monitor even marginal fluctuations in the fundamental value closely and these can trigger order revisions. HFTs, who are notorious investors in low-latency technology, adhere particularly strongly to the economics of limit order cancellations. We do not find support for the conjecture that short-lived limit orders originate from liquidity demanders. Our findings imply that policies banning, discouraging, or taxing limit order cancellations are likely to increase the costs of market making.

Appendix A: Summary Statistics for the Sample Stocks

Table A1 presents summary statistics for the sample stocks. *Market capitalization* is in millions of Swedish Krona (MSEK) using the closing prices of April 30, 2014. On that date, one MSEK was worth roughly 110,000 EUR or 152,000 USD. From Column 2 in Table A1, we note that market capitalization varies among the stocks with the highest being 386,786 MSEK (HM B) and the lowest being 5,429 MSEK (NOKI SEK).

INSERT TABLE A1 ABOUT HERE

The other variables in Table A1 use data from all trading days during April 2014. The daily trading volume and the daily turnover statistics include the continuous trading and the opening and closing call auctions at NASDAQ, but exclude trading at other venues. Daily turnover is the average fraction of market capitalization traded on a daily basis. The daily trading volume varies between 75 MSEK (MTG B) to 860 MSEK (ERIC B), while the daily turnover is between 0.16% (SHB A) and 4.27% (NOKI SEK) of the market capitalization.[11]

We present three spread measures, all measured at the time of trades and denoted in basis points (bps). When calculating the spread measures we exclude trades in the opening and closing call auctions, block trades each amounting to more than one MSEK, and trades on other venues than NASDAQ. The *Quoted Spread* is half the difference between the best offer and best bid prices available in the limit order book, divided by the midpoint. The quoted spread varies between 1.64 bps (TLSN) and 6.53 bps (GETI B). The quoted spread range corresponds closely

[11] Data on market capitalization are from NASDAQ Stockholm's website.

to the range Brogaard et al. (2015) report (between 2 bps and 6 bps for OMXS 30 stocks between mid-September and mid-October 2012).[12] Tick Spread is the nominal quoted spread divided by the tick size, and it ranges from 1.02 bps (VOLV B) to 3.82 bps (MTG B). *Effective Spread* is the trade value-weighted average absolute difference between the transaction price and the spread midpoint, relative the midpoint. We calculate the average effective spread first across trades and then across days, and it ranges between 2.02 bps (ASSA B) and 6.24 bps (ELUX B).

Depth at BBO is the average of the volumes available at the best bid and offer prices, or the average volume required to move the price in either direction. Typically, the average depth at BBO is around one MSEK, but varying substantially across stocks in the range between 0.18 MSEK (MTG B) and 2.76 MSEK (ERIC B). As for the spread measures, depth at BBO is observed at times of trades. *Volatility* is the average ten-second squared basis point returns. *Order to Trade Ratio* is the number of trades divided by the number of quote updates at the BBO.

References

Baron, M., Brogaard, J., Hagströmer, B., & Kirilenko, A. (2016). Risk and return in high-frequency trading. Working Paper.

Baruch, S., & Glosten, L. R. (2016). Tail Expectations, Imperfect Competition, and the Phenomenon of Flickering Quotes in Limit Order Book Markets. Working Paper.

[12] Brogaard et al. (2015) compare the quoted spread levels in their sample to spreads for the US stocks studied by Brogaard, Hendershott and Riordan (2014). Accordingly, a two bps spread is comparable to US large cap stocks, while a six bps spread is slightly higher than for an average US mid-cap stock.

Biais, B., Hillion, P., & Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the Paris Bourse. *The Journal of Finance*, 50(5), 1655-1689.

Brogaard, J., Hagströmer, B., Nordén, L., & Riordan, R. (2015). Trading Fast and Slow: Colocation and Market Liquidity. *Review of Financial Studies*, 28, 3407-3443.

Brogaard, J., Hendershott, T., & Riordan, R. (2014). High frequency trading and price discovery. *Review of Financial Studies*, 27, 2267–306.

Budish, E., Cramton, P, & Shim, J. (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics,* 130(4) 1547-1621.

Chakrabarty, B., Han, Z., Tyurin, K., & Zheng, X. (2006). A competing risk analysis of executions and cancellations in a limit order market. CAEPR Working Paper 2006-015.

Chakrabarty, B., & Tyurin, K. (2011). Book chapter: Market liquidity, Stock characteristics and order cancellations: The case of fleeting orders.

Conrad, J. S., Wahal, S., & Xiang, J. (2015). High Frequency Quoting, Trading, and the Efficiency of Prices. *Journal of Financial Economics,* 116(2), 217-291.

Copeland, T. E., & Galai, D. (1983). Information Effects on the Bid-Ask Spread. *The Journal of Finance*, 38, 1457-1469.

Fong, K. Y., & Liu, W. M. (2010). Limit order revisions. *Journal of Banking & Finance*, 34(8), 1873-1885.

Foucault, T. (1999). Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets,* 2(2), 99-134.

Foucault, T., Kozhan, R., & Tham, W. W. (2016). Toxic Arbitrage. *Review of Financial Studies* (forthcoming).

Foucault, T. & Menkveld, A. J. (2008). Competition for Order Flow and Smart Order Routing Systems. *The Journal of Finance,* 63, 119-158.

Glosten, L. R. (1994). Is the electronic order book inevitable? *The Journal of Finance,* 49(4), 1127-1161.

Hagströmer, B. (2017). Overestimated Effective Spreads: Implications for Investors. Working Paper.

Hagströmer, B., & Nordén, L. (2013). The diversity of high-frequency traders. *Journal of Financial Markets,* 16(4), 741-770.

Handa, P., Schwarz, R., & Tiwari, A. (2003). Quote setting and price formation in an order driven market. *Journal of Financial Markets*, 6, 2003, 461-489.

Harris, L. (2013). Maker - taker pricing effects on market quotations. Working Paper.

Hasbrouck, J., & Saar, G. (2009). Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12(2), 143-172.

Hautsch, N., & Huang, R. (2012). On the dark side of the market: Identifying and analyzing hidden order placements. Working paper.

Hendershott, T., Jones, C. M., & Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *The Journal of Finance*, 66(1), 1-33.

Hendershott, T., & Menkveld, A. J. (2014). Price pressures. *Journal of Financial Economics*, 114(3), 405-423.

Hollifield, B., Miller, R., Sandås, P., & Slive, J. (2006). Estimating the gains from trade in limit-order markets. *The Journal of Finance*, 61(6), 2753-2804.

Lee, E. T., & Wang J. W. (2003). Statistical Methods for Survival Data Analysis, 3rd edition. *Wiley Series in Probability and Statistics.*

Liu, W. M. (2009). Monitoring and limit order submission risks. *Journal of Financial Markets*, 12(1), 107-141.

Menkveld, A. J. (2013). High frequency trading and the new market makers*. Journal of Financial Markets,* 16(4), 712-740.

O'Hara (2015). High Frequency Market Microstructure. *Journal of Financial Economics,* 116(2) 257-270.

O'Hara, M., Saar, G., & Zhong, Z. (2015). Relative tick size and the trading environment. Working paper.

Parlour, C. A. (1998). Price dynamics in limit order markets. *Review of Financial Studies,* 11(4), 789-816.

Raman, V., & Yadav, P. (2014). Liquidity provision, information and inventory management in limit order markets: an analysis of order revisions. Working paper.

Sandås, P. (2001). Adverse selection and competitive market making: empirical evidence from a limit order market. *Review of Financial Studies,* 14(3), 705-734.

Securities and Exchange Commission (2010). Concept Release on Equity Market Structure; Proposed Rule. Federal Register, 75(13), 3593-3614.

Seppi, D. J. (1997). Liquidity Provision with Limit Orders and a Strategic Specialist. *Review of Financial Studies,* 10(1), 103-150.

Van Kervel, V. (2015). Competition for order flow with fast and slow traders. *Review of Financial Studies,* 28(7), 2094-2127.

Yueshen, B. Z. (2014). Queuing Uncertainty in Limit Order Market. Working paper.

Figure 1: Limit Order Duration

The figure shows the respective observed distribution function for the number of events (Panel A) and the time (Panel B) to cancellation and execution for the orders in our sample. We sample non-marketable, visible, orders randomly by choosing an order, in any one of the stocks, on either the best bid or best ask side in the stock's order book. We then track the order until cancellation, execution, or, if the order stays in the order book, until a maximum of 100 order book events occurs. We define events as any changes on the best levels in the stock's order book. We repeat the sampling procedure throughout the trading day. We use data from all trading days in May 2014 in the OMXS 30 stocks.

Panel A: Event time

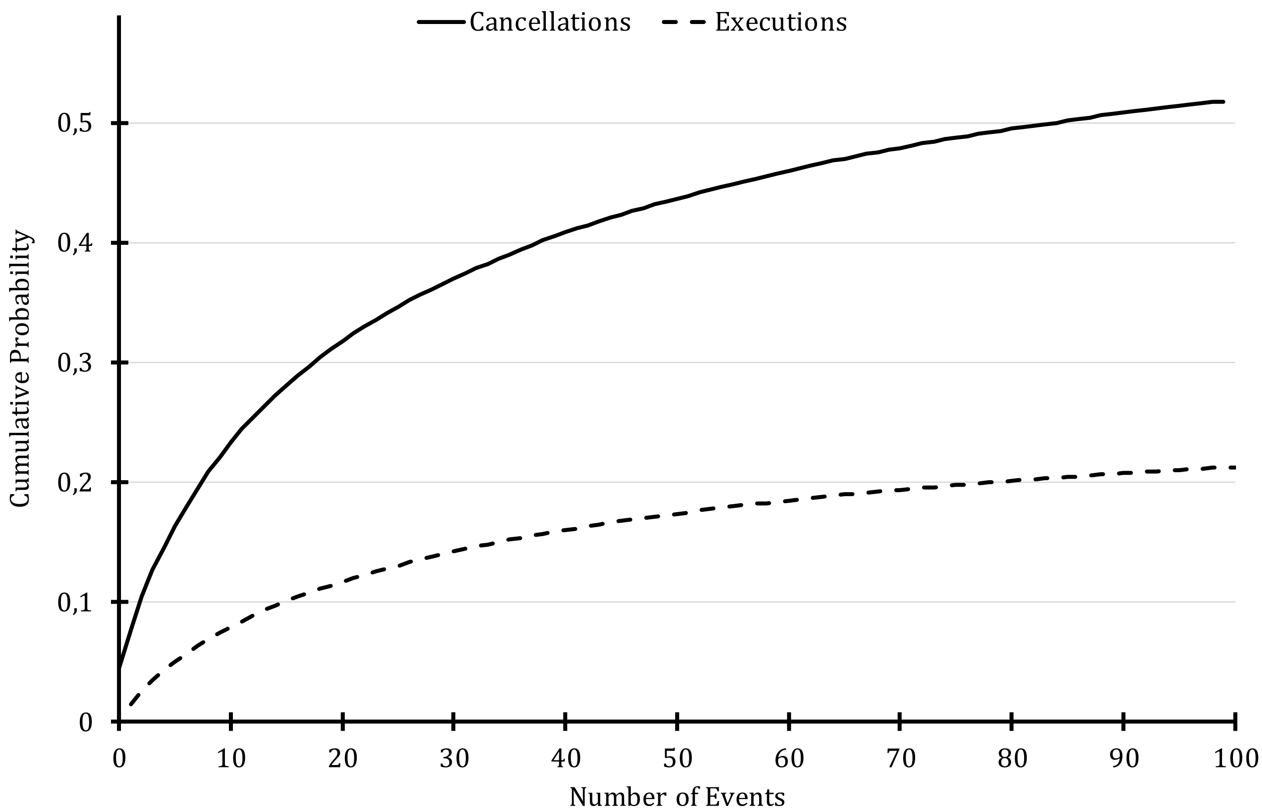

Panel B: Clock time

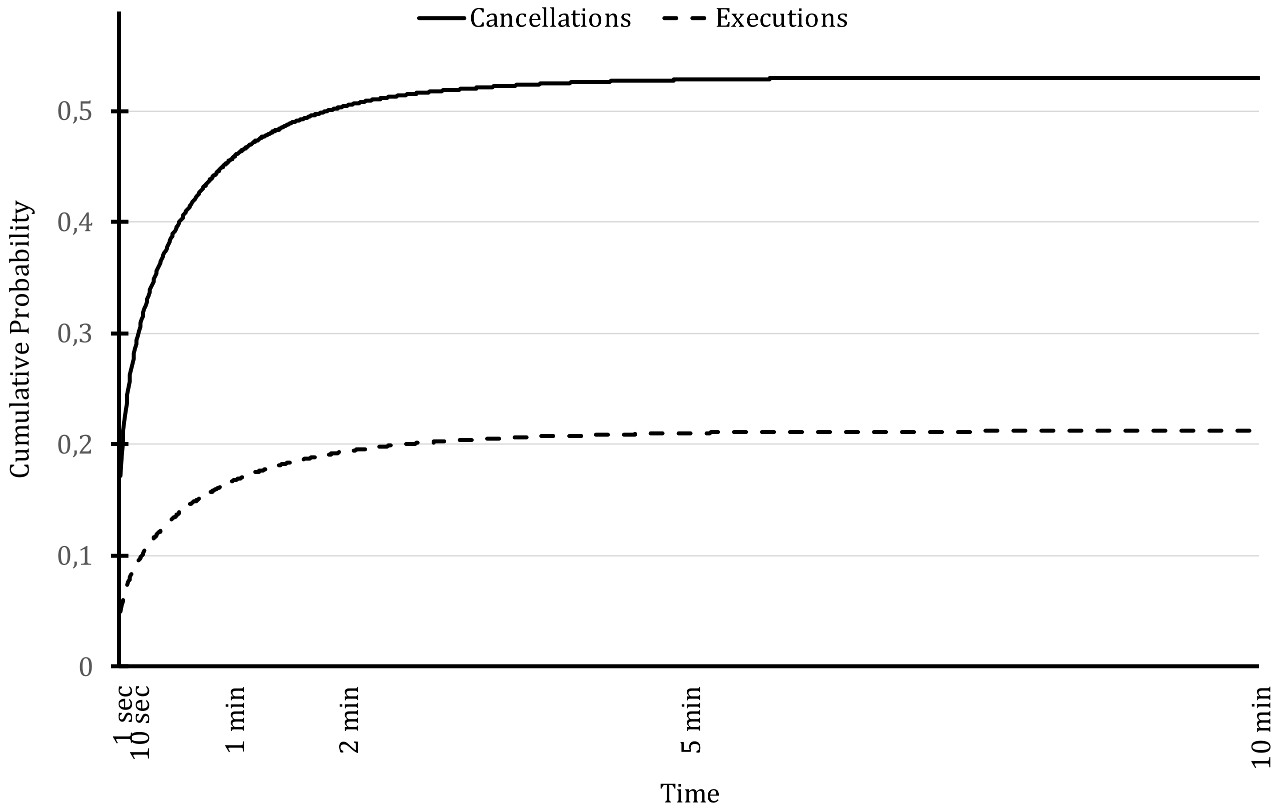

Table 1: Sample Order Descriptive Statistics

The table presents order level descriptive statistics for our model variables at the time of sampling. We sample non-marketable, visible, orders randomly by choosing an order, in any one of the stocks, on either the best bid or best ask side in the stock's order book. We then track the order until cancellation, execution, or, if the order stays in the order book, until a maximum of 100 order book events occurs. We define events as any changes on the best levels in the stock's order book. We repeat the sampling procedure throughout the trading day. We use data from all trading days in May 2014 in the OMXS 30 stocks. *Bid-Ask Spread* as the difference between the prevailing ask price and the bid price, divided by the midpoint, prevailing when sampling each order. *Back of Queue Quantity* is the fraction of volume with lower priority than each order available at the same price. *Front of Queue Quantity* is the fraction of volume with higher priority than each order, including the order's volume, available at the same price. *Opposite-Side Quantity* is the fraction of volume at the best price level on the opposite side of the order book (from the perspective of each order). *Order Quantity* is fraction of the number of stocks in the order. Fractions are relative to total depth (same side plus opposite side). *Order Duration* is the time in seconds to cancellation, execution, or until a maximum of 100 order book events occurs, of each order, given cancellation, execution, or survival of the order.

			Percentile				
Variable	Mean	St. Dev.	5^{th}	25^{th}	50^{th}	75^{th}	95^{th}
Bid-Ask Spread (bps)	7.29	3.50	3.10	5.29	6.02	9.61	13.24
Quantity variables expressed as fractions of total depth							
Back of Queue Quantity	0.21	0.21	0.00	0.02	0.15	0.34	0.64
Front of Queue Quantity	0.29	0.22	0.03	0.11	0.24	0.43	0.73
Opposite-Side Quantity	0.50	0.25	0.09	0.31	0.50	0.69	0.91
Order Quantity	0.08	0.10	0.01	0.03	0.05	0.10	0.26
Order duration conditional on outcome, expressed in seconds							
Cancellation	26.05	49.22	0.00	0.05	5.80	30.77	117.40
Execution	38.77	62.75	0.00	0.85	13.90	50.15	160.61
Surviving 100 Events	101.96	97.90	10.12	35.93	72.35	135.47	294.00

Table 2: Order Durations

The table presents order durations with respect to our model variables and selected control variables. *Order Duration* is the time in seconds, or number of events, to cancellation or execution of each order, given cancellation or execution of the order. For each sampled order, model variables are defined as in Table 1. *Lagged Absolute Return* is the absolute five-minute stock return prior to the sampling time. *Lagged Volume* is the number of traded stocks, relative to total depth (same side plus opposite side), during the five minutes prior to the sampling time. *Order Quantity* is the number of stocks in the order, relative to total depth (same side plus opposite side). *Internal Match Rate* is the proportion of each trading firm's limit orders that executes with internal priority during April 2014. *Binding Tick Size* refers to the case when, at the sampling time, the bid-ask spread equals its minimum. We identify high-frequency trading firms (*HFTs*) as those who are members in the proprietary trading lobby organization European Principal Traders Association (FIA-EPTA). For each variable, the sample orders are split into two groups: *Small* and *Large*, based on whether the variable in question is below or above its median. The table shows for each order group the mean time until order cancellation or execution. For each variable, *t*-values from a test of equality across the two groups (*Small* and *Large*) are reported for each order duration measure. *, **, and *** denote statistical significance on the 10%, 5%, and 1% level respectively.

	Order Duration			
Category	Seconds to Cancellation	Number of Events to Cancellation	Seconds to Execution	Number of Events to Execution
All orders	26.05	24.24	38.77	26.21
Small Bid-Ask Spread	21.67	22.76	30.06	24.60
Large Bid-Ask Spread	30.42	25.71	47.49	27.83
t-value	-24.17***	-16.03***	-24.02***	-10.98***
Small Back of Queue Quantity	21.06	19.92	30.10	20.98
Large Back of Queue Quantity	31.03	28.55	47.44	31.45
t-value	-27.61***	-47.46***	-23.90***	-36.32***
Small Opposite-Side Quantity	29.36	27.45	48.98	32.05
Large Opposite-Side Quantity	22.74	21.02	28.56	20.37
t-value	18.27***	35.07***	28.25***	40.72***
Small Front of Queue Quantity	25.56	23.99	33.44	23.53
Large Front of Queue Quantity	26.53	24.48	44.09	28.89
t-value	-2.69***	-2.63***	-14.59***	-18.29***
Small Order Quantity	28.28	26.69	41.19	28.17
Large Order Quantity	23.81	21.78	36.36	24.25
t-value	12.30***	26.70***	6.60***	13.34***
Small Internal Match Rate	22.94	23.95	37.12	25.72
Large Internal Match Rate	27.91	23.98	39.19	26.09
t-value	-13.12***	-0.11	-2.64***	-1.18
Binding Tick Size	23.39	23.49	36.25	25.62
Non-Binding Tick Size	30.25	25.41	50.60	30.14
t-value	-17.41***	-10.07***	-13.75***	-12.34***
HFT	19.72	22.25	28.72	25.78
Non-HFT	29.18	25.22	40.43	26.28
t-value	-27.14***	-15.41***	-13.43***	-1.18

Table 3: Proportional Duration Hazard Model of Order Cancellation

The table presents results from running the proportional duration hazard model with time varying covariates using data from all 19 trading days in May 2014, for the OMXS 30 stocks. We estimate the model:

$$\lambda_i(s) = \lambda_0(s)\,exp\left[Model_{i,0}\theta + \Delta Model_{i,s}\beta + \Delta Other_{i,s}\delta + Control_{i,0}\gamma\right],$$

where $\lambda_i(s)$ is the hazard rate of cancellation of order *i* over the next instant at time *s*, and $\lambda_0(s)$ is an unspecified baseline hazard rate. The vector $Model_{i,0}$ holds the model variables in levels, measured at the time of sampling as defined in Table 1. $BAS_{i,0}$ is the relative *Bid-Ask Spread*, $\overleftarrow{Q}_{i,0}$ is the *Back of Queue Quantity*, $\vec{Q}_{i,0}$ is the *Front of Queue Quantity*, and $Q_{i,s}^{Opp}$ is the *Opposite-Side Quantity*. The vector $\Delta Model_{i,s}$ contains changes in the model variables, defined as $\Delta Model_{i,s} = Model_{i,s} - Model_{i,0}$. The vector $\Delta Other_{i,s}$ contains the following time-varying covariates that are not part of the model. *Firm Inventory* is the net change in the number of bought shares minus the number of sold shares by the same trading entity that submitted order *i*. The *Firm Inventory* variable is multiplied with an indicator taking the value +1 if order *i* is on the bid side and -1 if the order is on the ask side. *Same-Side MTF Trades* is the cumulative number of trades on the same side as order *i*, since the time of sampling. *Opposite-Side MTF Trades* is the cumulative number of trades on the opposite side as order *i*, since the time of sampling. The vector $Control_i$ contains control variables, including *Order Quantity*, *Lagged Absolute Return* and *Lagged Volume*, all defined as in Table 2. Each control variable is measured at the time of sampling of order *i* and standardized (for each stock) to have zero mean and unit variance. The estimation uses 138,486 orders, which over time result in 6,141,702 observations and 73,458 cancellation events. ***, **, and * denote significance at the 1%, 5% and 10% level respectively. *t*-values are in parentheses. Standard errors are clustered on order *i*. For each variable in model specification (3), e.g., for a change in $\Delta BAS_{i,s}$ such that $\Delta BAS_{i,s}^* - \Delta BAS_{i,s}$ is a one standard deviation change, the *Hazard Ratio* is $\exp[\beta_1(\Delta BAS_{i,s}^* - \Delta BAS_{i,s})]$, where β_1 is the coefficient for $\Delta BAS_{i,s}^*$. The *Hazard Rate Change* is the hazard ratio minus one, expressed as a percentage.

Variable		Coefficient (t-value)	Hazard Rate Change
Model Variables: Levels			
Bid-Ask Spread	$BAS_{i,0}$	-0.046*** (-37.15)	-19.2%
Back of Queue Quantity	$\overleftarrow{Q}_{i,0}$	-1.991*** (-77.33)	-35.7%
Front of Queue Quantity	$\overrightarrow{Q}_{i,0}$	-0.036* (-1.92)	-0.8%
Model Variables: Time-Varying Changes			
Bid-Ask Spread	$\Delta BAS_{i,s}$	-0.036*** (-17.10)	-12.8%
Opposite-Side Quantity	$\Delta Q_{i,s}^{Opp}$	0.167*** (14.40)	10.2%
Back of Queue Quantity	$\Delta \overleftarrow{Q}_{i,s}$	-1.249*** (-36.72)	-48.3%
Front of Queue Quantity	$\Delta \overrightarrow{Q}_{i,s}$	0.036** (2.07)	3.0%
Other Variables: Time-Varying Changes			
Firm Inventory		0.011*** (2.69)	1.1%
Same-Side MTF Trades		0.047*** (16.27)	4.8%
Opposite-Side MTF Trades		-0.010** (-2.38)	-1.0%
Control Variables: Levels			
Order Quantity		1.137*** (38.63)	10.8%
Lagged Absolute Return		0.015*** (5.51)	1.6%
Lagged Volume		-0.003 (-0.72)	-0.3%
Internal Match Rate		-0.499*** (-10.65)	-4.0%

Table 4: Proportional Duration Hazard Model of Order Cancellation when the Tick Size is Binding or Not

The table presents results from running the pooled proportional duration hazard model with time varying covariates using data from all 19 trading days in May 2014, for the OMXS 30 stocks. We estimate the model:

$$\lambda_i(s) = \lambda_0(s)\, exp\left[Model_{i,0}\theta + \Delta Model_{i,s}\beta + \Delta Other_{i,s}\delta + Control_{i,0}\gamma\right],$$

separately for two subsamples of orders: those for which the tick size is binding, and those for which the tick size is not binding, at the time of sampling. All variable definitions are the same as in Table 3. The estimation uses 138,486 orders, which over time result in 6,141,702 observations and 73,458 cancellation events. ***, **, and * denote significance at the 1%, 5% and 10% level respectively. *t*-values are in parentheses. Standard errors are clustered on order *i*. The *Hazard Rate Change* is defined as in Table 3.

Variable		Binding Tick Size		Non- Binding Tick Size	
		Coefficient	Hazard Rate Change	Coefficient	Hazard Rate Change
Model Variables: Levels					
Bid-Ask Spread	$BAS_{i,0}$	-0.062*** (-21.30)	-21.7%	-0.059*** (-32.32)	-26.2%
Back of Queue Quantity	$\overleftarrow{Q}_{i,0}$	-2.085*** (-67.43)	-38.7%	-1.775*** (-38.43)	-29.7%
Front of Queue Quantity	$\overrightarrow{Q}_{i,0}$	-0.236*** (-9.94)	-5.2%	0.320*** (9.23)	6.5%
Model Variables: Time-Varying Changes					
Bid-Ask Spread	$\Delta BAS_{i,s}$	0.018*** (3.90)	6.6%	-0.047*** (-17.28)	-16.7%
Opposite-Side Quantity	$\Delta Q_{i,s}^{Opp}$	0.159*** (13.32)	10.4%	0.212*** (32.48)	10.8%
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{i,s}$	-1.363*** (-32.40)	-51.0%	-1.225*** (-19.54)	-45.0%
Front of Queue Quantity	$\Delta\overrightarrow{Q}_{i,s}$	-0.186*** (-3.01)	-11.3%	0.053*** (3.24)	5.6%
Other Variables: Time-Varying Changes					
Firm Inventory		0.005 (0.95)	0.5%	0.019** (2.43)	1.8%
Same-Side MTF Trades		0.049*** (14.87)	5.2%	0.041*** (7.06)	3.9%
Opposite-Side MTF Trades		-0.039*** (-6.81)	-4.0%	0.037*** (6.78)	3.3%
Control Variables: Levels					
Order Quantity		0.867*** (21.65)	7.9%	1.183*** (26.72)	11.7%
Lagged Absolute Return		0.015*** (3.92)	1.4%	0.014*** (3.11)	1.6%
Lagged Volume		-0.001 (-0.18)	-0.1%	-0.012** (-2.07)	-1.3%
Internal Match Rate		-0.440*** (-7.13)	-3.5%	-0.594*** (-8.17)	-4.8%

Table 5: Proportional Duration Hazard Model of Order Cancellation for HFTs and non-HFTs, Conditional on whether the Tick Size is Binding or Not

The table presents results from running the pooled proportional duration hazard model with time varying covariates using data from all 19 trading days in May 2014, for the OMXS 30 stocks. We identify high-frequency trading firms (*HFTs*) as described in Table 2. We estimate the model:

$$\lambda_i(s) = \lambda_0(s)\, exp\left[Model_{i,0}\theta + \Delta Model_{i,s}\beta + \Delta Other_{i,s}\delta + Control_{i,0}\gamma\right],$$

separately for four subsamples of orders: those submitted by HFTs and those submitted by others (Non-HFTs), conditional on whether the tick size is binding (Panel A) or not (Panel B), at the time of sampling.. All variable definitions are the same as in Table 3. The estimation uses 138,486 orders, which over time result in 6,141,702 observations and 73,458 cancellation events. ***, **, and * denote significance at the 1%, 5% and 10% level respectively. *t*-values are in parentheses. Standard errors are clustered on order *i*. The *Hazard Rate Change* is defined as in Table 3.

Panel A: Binding Tick Size		HFTs		Non-HFTs	
Variable		Coefficient	Hazard Rate Change	Coefficient	Hazard Rate Change
Model Variables: Levels					
Bid-Ask Spread	$BAS_{i,0}$	-0.154*** (-31.49)	-44.5%	-0.038*** (-10.37)	-14.4%
Back of Queue Quantity	$\overleftarrow{Q}_{i,0}$	-3.845*** (-59.18)	-60.3%	-1.490*** (-40.54)	-29.1%
Front of Queue Quantity	$\overrightarrow{Q}_{i,0}$	-1.560*** (-24.84)	-27.5%	0.242*** (8.81)	5.8%
Model Variables: Time-Varying Changes					
Bid-Ask Spread	$\Delta BAS_{i,s}$	0.056*** (6.82)	21.2%	0.011** (2.12)	3.9%
Opposite-Side Quantity	$\Delta Q_{i,s}^{Opp}$	0.274*** (18.84)	17.5%	0.134*** (17.97)	8.9%
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{i,s}$	-4.100*** (-40.77)	-87.6%	-0.575*** (-16.82)	-26.2%
Front of Queue Quantity	$\Delta\overrightarrow{Q}_{i,s}$	-3.001*** (-13.35)	-80.6%	-0.011 (-0.29)	-0.7%
Other Variables: Time-Varying Changes					
Firm Inventory		0.117*** (10.47)	10.1%	-0.014** (-2.16)	-1.5%
Same-Side MTF Trades		0.011* (1.79)	1.2%	0.063*** (14.03)	6.8%
Opposite-Side MTF Trades		-0.068*** (-6.05)	-7.8%	-0.023*** (-3.38)	-2.3%
Control Variables: Levels					
Order Quantity		2.480*** (29.23)	17.5%	0.479*** (9.76)	4.7%
Lagged Absolute Return		-0.003 (-0.29)	-0.3%	0.015*** (2.94)	1.3%
Lagged Volume		-0.028*** (-3.18)	-2.7%	-0.013** (-2.13)	-1.3%
Internal Match Rate		-0.807*** (-8.85)	-6.8%	1.104*** (13.68)	8.4%

Panel B: Non-Binding Tick Size		HFTs		Non-HFTs	
Variable		Coefficient	Hazard Rate Change	Coefficient	Hazard Rate Change
Model Variables: Levels					
Bid-Ask Spread	$BAS_{i,0}$	-0.140*** (-27.23)	-47.3%	-0.040*** (-20.71)	-19.1%
Back of Queue Quantity	$\overleftarrow{Q}_{i,0}$	-2.526*** (-23.33)	-38.7%	-1.556*** (-30.52)	-26.7%
Front of Queue Quantity	$\overrightarrow{Q}_{i,0}$	-0.191** (-2.48)	-3.4%	0.410*** (10.49)	8.5%
Model Variables: Time-Varying Changes					
Bid-Ask Spread	$\Delta BAS_{i,s}$	-0.034*** (-5.63)	-12.0%	-0.053*** (-17.68)	-19.1%
Opposite-Side Quantity	$\Delta Q_{i,s}^{Opp}$	0.262*** (28.48)	13.6%	0.200*** (26.02)	10.1%
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{i,s}$	-2.042*** (-12.16)	-66.6%	-0.880*** (-15.13)	-37.2%
Front of Queue Quantity	$\Delta\overrightarrow{Q}_{i,s}$	0.065 (1.45)	8.6%	0.066*** (3.58)	6.3%
Other Variables: Time-Varying Changes					
Firm Inventory		0.104*** (6.06)	8.6%	0.001 (0.13)	0.1%
Same-Side MTF Trades		0.030** (2.46)	2.6%	0.042*** (6.12)	4.0%
Opposite-Side MTF Trades		0.010 (0.75)	0.9%	0.045*** (7.27)	4.0%
Control Variables: Levels					
Order Quantity		1.560*** (17.87)	17.6%	0.912*** (16.87)	8.5%
Lagged Absolute Return		0.020*** (3.12)	2.3%	0.014*** (2.94)	1.7%
Lagged Volume		0.044*** (4.31)	4.8%	-0.035*** (-4.98)	-3.6%
Internal Match Rate		-2.090*** (-12.01)	-15.5%	-0.116 (-1.20)	-0.8

Table 6: Bundled Orders

This table reports frequencies of bundled orders. We define a bundled order as a trading firm's first subsequent action immediately following a cancellation. We present frequencies for each subsequent action as percentages of total. The subsequent *Action Types* we consider are to *Cancel* a resting limit order, to *Insert* a new limit order, to *Execute* a resting limit order, as well as to do *Nothing*. Do *Nothing* is when a trading firm does not act in the subsequent 500 milliseconds. For each *Action Type* we define *Action Details* as follows*. Same Side* is cancelling/inserting/executing a bid order (an ask order) in response to cancellation of a bid order (an ask order). *Opposite Side* is cancelling/inserting/executing a bid order (an ask order) in response to cancellation of an ask order (a bid order). Each category of *Action Details* is contingent on the price level of the subsequent cancellation/insertion/execution, relative the price level of the preceding cancelled order. If a trading firm inserts a bid order (an ask order) at a higher (lower) price than the preceding cancelled order it is labeled *Better Price.* The special case when such an order executes is *Marketable*. If a trading firm inserts a bid order (an ask order) at a lower (higher) price than the preceding cancelled order it is labeled *Worse Price.* If a trader inserts an order with the same price as the preceding cancelled order, we label it *Same Price*. The two rightmost columns show frequencies (as percentages of total) conditional on if the opposite side price moved away from the preceding cancelled order's limit price (bid-ask spread increase) or moved closer to it (bid-ask spread decrease) during the 500 milliseconds preceding the cancellation. The sample consists of the 73,458 cancelled orders from the analyses in Table 1 and Table 2.

			Conditional on Preceding 500 Millisecond Price Change		
Action Type	Action Details	Percentage of Total	Opposite Side Price Moves Away	Opposite Side Price Moves Closer	Opposite Side Price Unchanged
Nothing		43.8	1.9	1.9	40.1
Cancel	Cancel - All	19.8	2.1	2.0	15.7
	Same Side - Same Price	10.7	0.8	0.8	9.1
	Same Side - Other Prices	4.8	0.5	0.5	3.8
	Opposite Side - All	4.3	0.8	0.7	2.8
Insert	Insert - All	35.4	3.6	3.8	28.0
	Same Side - All	23.9	2.4	2.4	19.1
	Same Side - Marketable	0.8	0.1	0.1	0.6
	Same Side - Better Price	4.4	0.9	0.9	2.6
	Same Side - Same Price	8.5	0.6	0.6	7.4
	Same Side - Worse Price	10.2	0.9	0.8	8.6
	Opposite Side - All	11.5	1.2	1.3	8.9
	Opposite Side - Marketable	1.8	0.1	0.1	1.7
	Opposite Side - Better Price	9.2	1.1	1.2	6.8
	Opposite Side - Same Price	0.5	0.0	0.0	0.5
	Opposite Side - Worse Price	0.0	0.0	0.0	0.0
Execute	Execute - All	1.0	0.1	0.1	0.8
	Same Side - Same Price	0.6	0.0	0.0	0.5
	Same Side – Other Prices	0.1	0.0	0.0	0.1
	Opposite Side - All	0.3	0.1	0.0	0.2
All	All	100.0	7.7	7.8	84.6

Table A1: Stock Characteristics

The table lists properties of the OMXS 30 sample stocks. All except NOKI SEK are stocks traded primarily at NASDAQ Stockholm. NOKI SEK is a Swedish Depositary Receipt issued by the Finnish firm Nokia Oyj. We calculate *Market Capitalization* based on the closing price of April 30, 2014 (expressed in million Swedish Krona, *MSEK*). The market capitalization for NOKI SEK only represents its Depositary Receipt size. We obtain all other statistics as averages across trading days during April 2014. *Daily Trading Volume* is in MSEK. *Daily turnover* is daily trading volume divided by the market capitalization. *Quoted Spread* is half the bid-ask spread divided by its midpoint, averaged across seconds and expressed in basis points. *Tick Spread* is the nominal quoted spread divided by the minimum price change. *Effective Spread* is the trade value-weighted average absolute difference between the trade price and the bid-ask midpoint. *Depth at BBO* is the MSEK trade volume required to change the stock price, averaged across seconds and the bid and offer sides of the book. *Volatility* is the average 10-second squared basis point returns, calculated from bid-ask midpoints. *Order to Trade Ratio* is the number of trades divided by the number of quote updates.

Stock	Market Cap. (MSEK)	Daily Trading Vol. (MSEK)	Daily Turnover (%)	Quoted Spread (bps)	Tick Spread (bps)	Effective Spread (bps)	Depth at BBO (MSEK)	Volatility (sq. bps)	Order to Trade Ratio
ABB	90,985	230	0.25	3.51	1.19	3.45	1.46	3.98	13.04
ALFA	72,356	253	0.35	3.95	1.41	4.30	0.82	6.74	4.22
ASSA B	121,049	244	0.20	3.41	2.30	2.02	0.56	4.28	3.39
ATCO A	157,974	565	0.36	3.49	1.33	2.99	1.80	5.59	3.72
ATCO B	68,913	170	0.25	4.93	1.77	3.49	0.94	4.18	5.60
AZN	76,983	437	0.57	2.47	1.72	3.44	0.93	6.52	9.60
BOL	27,023	215	0.80	4.31	1.68	2.97	0.61	3.92	4.35
ELUX B	54,191	396	0.73	4.71	1.42	6.24	0.81	6.20	2.75
ERIC B	238,899	860	0.36	3.16	1.07	3.28	2.76	4.55	2.25
GETI B	42,497	162	0.38	6.53	2.37	3.33	0.62	4.24	2.72
HM B	386,786	739	0.19	2.05	1.13	2.56	1.18	2.10	3.06
INVE B	114,554	282	0.25	2.90	1.39	5.16	0.58	2.76	3.03
LUPE	44,158	120	0.27	5.43	1.46	3.72	0.72	3.47	3.76
MTG B	17,834	75	0.42	6.41	3.82	4.02	0.18	7.04	3.37
NDA SEK	380,493	630	0.17	2.91	1.07	3.65	1.40	3.67	4.21
NOKI SEK	5,429	232	4.27	4.59	3.71	4.72	0.55	8.03	5.57
SAND	115,090	522	0.45	3.26	1.21	3.80	1.29	5.24	2.93
SCA B	112,482	386	0.34	4.06	1.50	2.82	0.81	5.13	3.20
SEB A	194,325	512	0.26	3.10	1.10	2.81	1.10	4.06	3.88
SECU B	27,294	95	0.35	5.00	1.54	3.62	0.37	3.60	3.35
SHB A	203,584	327	0.16	2.35	1.52	3.94	1.19	2.12	4.76
SKA B	59,479	228	0.38	4.22	1.25	3.59	0.99	3.62	3.67
SKF B	70,021	364	0.52	4.02	1.34	3.95	0.84	4.88	2.90
SSAB A	13,591	117	0.86	6.07	1.21	6.07	0.37	5.67	2.80
SWED A	195,497	487	0.25	3.27	1.12	2.92	1.67	4.52	3.82
SWMA	45,026	200	0.44	3.94	1.69	2.90	0.95	3.93	3.48
TEL2 B	35,318	227	0.64	5.19	1.66	3.37	0.63	3.25	3.05
TLSN	204,207	493	0.24	1.64	1.50	2.07	1.10	1.86	2.44
VOLV B	165,472	721	0.44	4.97	1.02	6.02	2.05	5.20	2.01
Max	386,786	860	4.27	6.53	3.82	6.24	2.76	8.03	13.04
Min	5,429	75	0.16	1.64	1.02	2.02	0.18	1.86	2.01