Determinants of Limit Order Cancellations

Petter Dahlström, Björn Hagströmer, Lars L. Nordén*

First draft: Aug. 2, 2017

Current draft: Jun. 28, 2018

Abstract

We investigate the economic rationale behind limit order cancellations from the perspective of liquidity suppliers. We predict that an order is cancelled whenever its expected revenue no longer exceeds the expected cost and we model how order profitability variation can be determined from changes in the state of the order book and the order queue position. Our empirical evidence supports the predictions in general and for orders submitted by high-frequency trading firms in particular. Consistent with our model approach, we find that order cancellation patterns are more consistent with market making than with liquidity demand strategies.

*All authors are at Stockholm Business School, Stockholm University, SE-106 91 Stockholm, Sweden; emails petter.dahlstrom@sbs.su.se; bjorn.hagstromer@sbs.su.se; lars.norden@sbs.su.se. Petter Dahlström is affiliated with Nasdaq Stockholm. 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, Erik Hjalmarsson (discussant), Amy Kwan, Richard Philip, Gideon Saar, Patrik Sandås, Vincent van Kervel, Ran Xing (discussant), Bart Zhou Yueshen, Sarah Zhang (discussant), Marius Zoican, and Olivier David Zerbib (discussant), as well as participants at the 2017 EFMA Annual Meeting, the 1st SAFE Conference on Market Microstructure, the 11th Annual SoFiE Conference, the 2018 EFM Europe Conference, and the 6th Swedish House of Finance National PhD Workshop in Finance for useful discussions. We thank Nasdaq for providing the data. The paper does not reflect the views of Nasdaq, its officers, or directors.

Continuous auction markets feature a constant flow of limit order submissions and cancellations, potentially in response to fluctuations in security valuations, market conditions, and trading interests. It is well known that investors submit limit orders either to fulfill a trading need or to profit from intermediation, but the rationale for withdrawing an order is unclear. In this paper, we provide a theoretical framework and an empirical analysis of the determinants for limit order cancellations.

Our analytical starting point is that a limit order is a free option issued to all market participants (Copeland and Galai (1983)), offered by rational investors only when it implies a positive expected profit upon execution. For example, a market maker prices an order such that the revenue at execution (the difference between the limit price and the fundamental value) is expected to exceed the costs. By the same reasoning, a limit order should be withdrawn only when its expected profit becomes negative.

We use the reduced-form model of limit order supply with adverse selection costs of Sandås (2001), building on the work of Glosten (1994) and Seppi (1997), to identify two key determinants of the expected profit of a limit order. First, the expected revenue of a limit order depends on the fundamental value of the security. We show that, in equilibrium, the fundamental value is reflected by the quantities posted at the best bid and offer (BBO) prices of the limit order book. Second, the expected adverse selection cost of an order depends on its execution priority. Limit orders submitted at the same price level are executed on a first-come, first-served basis, referred to as time priority. We show that orders with higher time priority are less exposed to adverse selection. Taken together, these two channels imply that the expected profit of a limit order is a function of four observable order book variables: (i) the quantity on the opposite side of the order book (relative to the order in question); (ii) the quantity on the same side but with lower execution priority; (iii) the quantity on the same side but with higher execution priority; and (iv) the bid–ask spread.

We postulate that cancellations are due to shocks to the expected profit, making a change in each of the four order book variables a potential cancellation trigger. Our expression of the expected profit is an equilibrium outcome of the static model of Sandås (2001). To derive predictions about cancellations, we view them as part of the transition to a new equilibrium.[1] To understand the intuition, consider a security whose 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 security, moving the fundamental value closer to the best ask price. The expected profits of orders posted at the best ask price therefore fall. The marginal ask-side limit order is now expected to become unprofitable and the rational response is to cancel it. The larger the fundamental value shock, the more ask-side orders are canceled.

For the empirical analysis, we turn to the equity market, using a proprietary data set of limit order book events on Nasdaq Stockholm (henceforth Nasdaq). The sample includes 29 large-cap stocks and all the trading days of May 2014. The data allow us to track each limit order from submission to either cancellation or execution and to trace the order time priority, which is key to testing the model predictions. To obtain independent observations, we sample one order at a time for each stock and track it for up to 100 limit order book events.

---

[1] Sandås (2001) and van Kervel (2015) use similar approaches to form dynamic predictions from a static model. Whereas they focus on how a trade signals updates in the security value, we use order book updates to gauge changes in the fundamental value and the adverse selection costs.

The aim of our empirical analysis is to assess the determinants of the cancellation intensity over the next instant, known as the hazard rate of cancellation. We estimate the hazard rate sensitivity to shocks in the expected profit of each limit order in a stratified proportional duration hazard model (PDHM) with time-varying covariates. The stratified PDHM allows us to analyze the impact of changes in the expected profit during the limit order lifetime while controlling for executions, order-specific market conditions, and stock-level effects. Our findings lend strong support to the prediction that a limit order is more likely to be canceled following a negative shock to its expected profit.

The theoretical model assumes an economy in which a publicly known fundamental value determines the liquidity supply decisions and adverse selection risk is equal across liquidity providers. All trade is centralized at one trading venue. When investigating the model predictions empirically, we relax the model assumptions by controlling for changes in inventory (Hendershott and Menkveld (2014), Raman and Yadav (2014)) and trades at other venues that could reflect new information (van Kervel (2015)). Our results show that, even though these alternative motives have a statistically significant influence on cancellations, their inclusion does not alter the main findings and their economic significance is relatively small. For example, a standard deviation increase in the order book quantity at the same price level as the order of interest but with a lower time priority (one of our model variables) boosts the life expectancy of a limit order by almost 7 seconds, an increase by 27% for a typical sample order. A standard deviation change in any of the cancellation determinants that are not in the model leads to no more than a 0.6-second change in order life expectancy.

Our main contribution to the literature is to derive and test predictions for the determinants of limit order duration and cancellations. 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 (Parlour (1998), Foucault (1999), 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 empirically showing that the model predictions dominate competing motivations for cancellations. Furthermore, the role of the orders' time priority, also known as the queue position, as a determinant of cancellations is new to the literature.

A key aspect of our contribution to the understanding of limit order cancellations is that we take the perspective of a liquidity supplier. This is in contrast to the models of Harris (1998), Foucault, Kadan, and Kandel (2005), Kaniel and Liu (2006), and Roşu (2009), which highlight how traders with private values of trading (liquidity demanders) choose between limit orders and market orders. Prior empirical studies of limit order cancellations by Ranaldo (2004) and Hasbrouck and Saar (2009) also take a liquidity demand perspective. O'Hara (2015) emphasizes that liquidity demanders make extensive use of passive limit orders. Our empirical evidence indicates that market makers rather than liquidity demanders dominate the limit order traffic, supporting the supply-side approach. First, we find that limit order cancellations contribute to mean reversion of the inventory at the trading firm level, which is consistent with market making but not with liquidity demand. Second, we find that cancellations as part of order price revisions, which are commonly associated with liquidity demand strategies (e.g., Hasbrouck and Saar (2009)), are relatively rare. Finally, we argue that an extrapolation of our predictions for cancellations to limit order traders with private values has merit, because they adhere to the same common values as traditional market makers. Raman and Yadav (2014), van Kervel (2015), and Yueshen (2014) also take the supply-side perspective of limit order management.

Our contribution is distinct from their work. We focus on the role of changes in common values, whereas they analyze the effects of inventory management, cross-venue market making, and queuing uncertainty, respectively.

High-frequency traders (HFTs) are a group of trading firms that is frequently associated with short-lived limit orders. Given the extensive debate of such orders—known as fleeting orders (Hasbrouck and Saar (2009)), flickering quotes (Baruch and Glosten (2016)), or phantom liquidity (US Securities and Exchange Commission (2010))—an interesting question is whether HFTs are distinct from other traders in terms of order cancellations. We identify HFTs at the trading firm level and analyze their limit order activities relative to other trading firms. We confirm that HFT orders have shorter lifetimes but find that the determinants of cancellations are the same. The main difference relative to other trading firms is that HFTs respond more strongly to changes in the expected profits of their limit orders. This is consistent with the notion that HFTs invest heavily in the technology required to monitor the market in real time and the results of Brogaard, Hendershott, and Riordan (2014) show that HFTs drive price discovery by quickly updating their quotes following fundamental value changes.[2] Our evidence contributes to the debate on the role of HFTs in modern equity markets, supporting the view that market making dominates the HFT order flow (Hagströmer and Nordén (2013), Menkveld (2013)).

Overall, this paper concludes that even small indications of changes in the fundamental value can trigger cancellations of marginally profitable limit orders. Consistent with the high-

[2] The result is also consistent with HFTs, compared to non-HFTs, caring more about the tradeoff between waiting costs and the costs of immediate execution (Hagströmer, Nordén, and Zhang (2014)) and reacting more strongly to order book depth imbalances (Goldstein, Kwan, and Philip (2017)).

frequency order management described by Budish, Cramton, and Shim (2015) and Foucault, Kozhan, and Tham (2017), our results indicate that the increasing prevalence of short-lived orders is due to an increased benefit of order monitoring and order management. With continuous changes in fundamental values, manifested in frequent order book updates, it is not surprising that cancellations are also frequent.

Our results have implications for initiatives aimed at curbing the frequency of cancellations. Examples of such policies include the minimum quote life at EBS (a leading inter-dealer platform in the foreign exchange market), cancellation fees applied at Canadian stock exchanges, and order-to-trade ratio caps imposed by German regulators and at many stock exchanges. The fact that traders, both fast and slow, adhere to a supply-related economic rationale in their order management implies that curbs on cancellations are likely to restrict the ability of market makers to fine-tune their outstanding orders. Hence, policies aimed at curbing order cancellations instead risk curbing liquidity supply.

I. Theoretical Framework

The premise of this paper is that liquidity suppliers cancel limit orders in response to reductions in the profit they expect 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 of Glosten (1994) and Seppi (1997). We use Sandås' (2001) model to show that the expected profit of a limit order is a function of its execution priority, as well as the prices and volumes available at the BBO prices. The outcome is testable predictions on the determinants of limit order cancellations.

A. Model Setup

The market consists of a large number of liquidity demanders, who could 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:

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

The parameters φ_A and φ_B are the expected sizes for market orders arriving on the ask and bid sides, respectively. The liquidity demander is a buyer or 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 expect a non-negative profit. Let $\tilde{\pi}_{i,t}$ denote the expected profit upon execution of a limit order i posted on the ask side of the book and assume it takes the following form:

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

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. Let $\vec{Q}_i$ denote the market order volume required to execute the full size of order *i,* which includes the size of order i (assumed to be one unit) as well as the sizes of all orders with higher execution priority than order i. We refer to $\vec{Q}_i$ as *Front of Queue Quantity* and use the forward arrow to symbolize the front. Below, we also use the notation $\overleftarrow{Q}_i$ to describe *Back of Queue Quantity*, where the backward arrow indicates that the variable captures the order volume at the same price level but with a 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. We henceforth drop the time index t for brevity of notation.

The model assumes that liquidity suppliers know and agree on the fundamental value of the risky asset at the beginning of each period. In addition, they do not have private values influencing their order submissions. In reality, however, liquidity suppliers can 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 do not include 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 economically more important.

Sandås (2001) shows that, by taking an integral over the distribution of market order arrivals and factoring in Equation (3), the expected profit expression in Equation (2) becomes

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

Note that $P_i - \gamma - X$ represents the effective spread net of order processing costs, which could be viewed as the revenue of the limit order. The order is expected to be profitable as long as the adverse selection cost, represented by $\alpha\left(\vec{Q}_i + \varphi_A\right)$, does not exceed the revenue. We use Equation (4) as our starting point for the analysis of the determinants of limit order cancellations. In equilibrium, we express the fundamental value X in terms of observable variables.

B. 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, since 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.

We solve Equation (6) for α, insert the resulting expression in Equation (5), and solve for X to obtain an expression for the fundamental value as follows:

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

Equation (7) shows the fundamental value in equilibrium as a weighted average of the BBO prices, with the imbalance between the ask- and bid-side depths (adjusted for the expected market order size and order processing costs) determining the weights. The concept of depth-weighted prices to proxy for fundamental values is common in the financial industry and sometimes goes under the name *micro-price* or *weighted midpoint* (Harris (2013), Hagströmer (2017), Stoikov (2017)). Empirical evidence presented by Cont, Kukanov, and Stoikov (2014) and Gould and Bonart (2016) shows that the depth imbalance has predictive power for future price changes.

By inserting Equation (7) in place of 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, *Front of Queue Quantity* ($\vec{Q}_i$). The difference between the two is *Back of Queue Quantity* ($\overleftarrow{Q}_i$).

We analogously obtain an expression for the expected profit of bid-side limit orders. To set up a general expression for the expected profit of a limit order, we adopt the following terminology. We 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. Note that, at the time of execution, order *i* is typically at the best price level, making the definition similar to the conventional measure of the quoted bid-ask spread. Furthermore, we define *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 and opposite sides 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 assumed constant, all variation in the expected profit is due to fluctuations in *Bid-Ask Spread*, *Opposite-Side Quantity*, *Back of Queue Quantity*, and *Front of Queue Quantity*. We refer to these four variables as the model variables.

C. Predictions

We form dynamic predictions by comparing equilibrium liquidity before and after a change in the fundamental value. The model is static and does not convey how the limit order book reaches its equilibrium. We postulate that changes in the fundamental value trigger cancellations of existing limit orders and submissions of new limit orders. This approach to form dynamic predictions from a similar static model is akin to that of van Kervel (2015), who analyzes how a trade at one venue leads to cancellations at another venue where the same security is traded.

Our predictions concern the hazard rate of cancellation (the cancellation intensity over the next instant). We postulate that the hazard rate of a limit order is negatively related to shocks to 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 (unless they are executed). A negative shock can turn the expected profitability of a limit order negative, leading a rational agent to cancel the order.[3]

We use Equation (9) to form empirical predictions of how each of the model variables influences the hazard rate of limit order cancellation. Note that the expression in square brackets determines the sign of the expected profit. The exponential function, which captures the probability that the order is executed, is merely a scaling factor. The first term within squared brackets in Equation (9) reflects the limit order revenue at execution, known as the effective bid-ask spread, net of order processing costs. The term corresponds to $P_i - \gamma - X$ in Equation (4). Due to the equilibrium expression for the fundamental value, we are able to derive clear predictions about the effects of three of the model variables:

- The variable *Bid-Ask Spread* (BAS_i) is positively related to the limit order expected profit and thus negatively related to the hazard rate of cancellation.

[3] The spirit of the model is that the probability of cancellation is binary: zero for orders with non-negative expected profits and one otherwise. If we knew 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. Our aim here, however, is to understand the determinants of cancellations. Rather than a threshold effect, we investigate the continuous relation between shocks to the expected profit and the hazard rate of cancellation.

- The variable *Opposite-Side Quantity* (Q^{Opp}) is negatively related to the expected profit and positively related to the hazard rate of cancellation. To see this, note that the fraction in 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.[4]

- The variable *Back of Queue Quantity* ($\overleftarrow{Q}_i$) is positively related to the expected profit and negatively related to the hazard rate of cancellation. The variable enters in both the numerator and denominator of the depth imbalance ratio, but its effect on the numerator is always greater than that on the denominator.[5]

The effect from the fourth model variable, *Front of Queue Quantity* ($\vec{Q}_i$), on expected profits is unclear, because it enters both terms in squared brackets in Equation (9). The second term, $\alpha(\vec{Q}_i + \varphi)$, represents the adverse selection cost of the order. It captures the fact that an order further back in the queue requires a larger market order to be executed and larger market orders are associated with higher price impacts. The variable *Front of Queue Quantity* is hence positively related to both the adverse selection cost and the revenue or the order. The net of the two effects is an empirical question. A relative prediction is still possible: the effect of *Front of Queue Quantity* on the expected profit is lower than the corresponding effect of *Back of Queue*

---

[4] Note that the expected market order sizes φ and φ^{Opp} are strictly positive and that *Bid-Ask Spread* is always greater than twice the order processing cost.

[5] To see that the relation to expected profit is strictly positive, note again that *Bid-Ask Spread* is always greater than twice the order processing cost and that the quotient rules imply $\partial\left(\frac{(\vec{Q}_i+\overleftarrow{Q}_i+\varphi)}{\vec{Q}_i+\overleftarrow{Q}_i+\varphi+Q^{Opp}+\varphi^{Opp}}\right)/\partial\overleftarrow{Q}_i = \frac{(\vec{Q}_i+\overleftarrow{Q}_i+\varphi+Q^{Opp}+\varphi^{Opp})-(\vec{Q}_i+\overleftarrow{Q}_i+\varphi)}{(\vec{Q}_i+\overleftarrow{Q}_i+\varphi+Q^{Opp}+\varphi^{Opp})^2} > 0.$

Quantity. Similarly, we predict that the effect of *Front of Queue Quantity* on the hazard rate of cancellation is higher (less negative) than that of *Back of Queue Quantity*.

The difference in the predicted effects of volume with higher and lower execution priority illustrates 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 the empirical evidence of Chakrabarty et al. (2006) and Hollifield et al. (2006)). To our knowledge, the role of the order queue position presented here is new to the literature.

II. Data and Sample

A. 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 that 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 each trading day.

All order submissions in the Nasdaq data set contain a trading firm identifier. A trading firm is an entity connecting to Nasdaq either as an exchange member or through another exchange member as a sponsored access client. Following Baron, Brogaard, Hagströmer, and Kirilenko (2017), we identify a subset of the trading firms as HFTs based on the member list of FIA European Principal Traders Association (EPTA), an industry organization for proprietary trading firms in Europe.

We also retrieve trade records and order book data at the microsecond frequency from Thomson Reuters Tick History. We use these data to measure trading activity at other trading venues besides Nasdaq. During our sample period, Nasdaq comprises 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%).

B. Institutional Details

Nasdaq operates an electronic limit order book market open from 9:00 AM to 5:30 PM every weekday except for 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 of a Swedish krona (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.[6] The tick size rule is important because it restricts, for example, the possibility of competing for order flow by posting aggressive orders

[6] On April 30, 2014, 100 SEK corresponded to about 15 US dollars (USD).

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 on Nasdaq is set according to price, internal, visibility, and 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 to trade with an incoming market order, even if it is first in terms of time priority. The internal matching can affect traders' decision to cancel and we control for that in our empirical analyses.

The use of hidden liquidity is highly restricted on Nasdaq. For the stocks in our sample, an order must be worth at least 1 million euros (EUR) to be eligible for full non-visibility. For some stocks, the threshold for hidden orders 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, p. 2) report that, in the US equity markets, hidden liquidity is associated with "enormous order activities" related to liquidity detection strategies, which are likely to involve aggressive market orders. The restrictions on hidden liquidity on Nasdaq are likely to induce fewer liquidity detection strategies relative to, for example, activities on the US equity markets.

Do traders know their queue positions? The Nasdaq data feed contains information about order submissions, cancellations, modifications, and executions, but no explicit information about queue positions. It conveys order identification numbers that allow trading firms to track which of the existing orders is canceled or executed. The information enables trading firms to track their execution priority with high accuracy, but such real-time monitoring can be costly.

However, for trading firms that choose to incur the costs, queue position tracking can constitute a competitive advantage.

C. *Sample Construction*

For each stock-day, we randomly sample one order posted at the best prices in the order book at 9:05 AM.[7] We track that order until it is executed or canceled or until 100 order book events have passed.[8] After 100 order book events have elapsed, irrespective of outcome, we sample a new order. We repeat this procedure until the end of the continuous trading. Sample orders that are partially or fully hidden or that have been sampled before are omitted from the subsequent analysis.

We design the order sampling to correspond closely to the theoretical model. Consistent with the fundamental value being derived from break-even conditions for the BBO prices, we sample orders posted at the BBO only. Because orders in the model are always for one share, we ignore partial cancellations and executions, effectively focusing on the last share of each sample order. Finally, an important feature of the model is that orders have non-negative expected profits. By drawing from all BBO orders prevailing at a fixed point in time, rather than at the time of order submission, we avoid orders with negative expected profits.[9]

---

[7] At each sampling time, we first randomly select, with equal probability, 50%, whether the order should be drawn from the best bid price or the best ask price. We then randomly select one order, based on the uniform distribution, from the orders available on the selected price.
[8] We define order book events as changes to either the prices or volumes at the BBO. If there are multiple changes to the order book quantities within the same unit of time (1 nanosecond), we still view them as one event. Given the short time span, it is unlikely that multiple changes in the same instant are due to different orders.
[9] Yueshen (2014) shows that queuing uncertainty can generate flickering quotes. In the author's model, stochastic latency introduces randomness in the order queues, leading agents to quickly cancel orders that do not obtain a

The sampling procedure also considers assumptions underlying our empirical estimation. The use of event times ensures that we do not sample orders at the same clock time for different stocks, which is important to satisfy the requirement of independent observations in the statistical analysis. Within stocks, we ensure independence by restricting the analysis to one order at a time. Furthermore, by using event time, we can analyze cancellation hazard rates across order books that differ in the intensity of events. By drawing orders prevailing in the order book, we obtain a sample dispersed in terms of expected profits and queue positions, which is important for the identification of all variables.

The sampling procedure implies that we obtain more observations from order books with more intense BBO updates. Hence, one way to interpret our results is that they are order-activity weighted across stock-days. Our sample selection is also more likely to select orders of longer duration, because such orders are eligible for selection in more draws. In robustness tests presented below, we show that restrictions to orders sampled at the time of submission or to short-lived orders do not alter the conclusions of our analysis.

D. Descriptive Statistics

The sampling procedure provides us with 138,487 limit orders, of which 38,656 orders (27.9%) are submitted by HFTs. Within the period of 100 order book events, 53% of all orders are canceled, 21% are executed, and 26% remain in the order book. Orders submitted by HFTs

---

good enough queue position. Our model is silent on how equilibrium is reached and hence does not explain this type of cancellation.

are canceled in 63% of cases, executed in 11% of cases, and, similar to the entire sample of orders, remains in the order book after 100 events in 26% of the cases.

Table I presents descriptive statistics for the model variables at the time of sampling of each order. Panel A shows the statistics for the sample of all orders, while Panel B reports the corresponding statistics for the subsample of orders submitted by HFTs. For comparability across orders, we normalize the model variables as follows. We divide the variable *Bid-Ask Spread* by the spread midpoint. 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 BBO. In addition, we define the *Order Quantity* as the volume of the sampled order divided by the total BBO depth.

INSERT TABLE I ABOUT HERE

The first column of Panel A in Table I presents the mean of each variable for the entire sample of orders. At the time of sampling, the average *Bid-Ask Spread* is 7.29 basis points. The limit order book is, on average, symmetric, since *Opposite-Side Quantity* accounts for 50% of the total depth and the quantity on the same side accounts for the remaining 50% of the total depth (*Back of Queue Quantity* plus *Front of Queue Quantity*). The sampled *Order Quantity* accounts for 8% of the total depth. Since *Front of Queue Quantity* includes *Order Quantity,* it constitutes a higher share of the total depth (29%) than *Back of Queue Quantity* (21%). From Panel B, we note that the quantity variables for the subsample of HFT orders follow very similar distributions as for the entire sample. Moreover, the average bid-ask spread is somewhat lower for the HFT order sample relative to the entire sample, which could be an indication of HFTs supplying liquidity relatively often in more liquid stocks.

Table I also presents descriptive statistics for *Order Duration* conditional on outcomes. We define *Order Duration* as the number of seconds from the time of sampling to each of the outcomes cancellation, execution, and orders remaining after 100 events, respectively represented by *Cancellation*, *Execution*, and *Surviving 100 Events*. For the entire sample in Panel A, the mean of *Cancellation* (26.05 seconds) is shorter than the mean of *Execution* (38.77 seconds). The mean of *Surviving 100 Events* is 101.96 seconds. This result implies that we track the sampled orders for, on average, no more than approximately one and a half minutes. The percentile statistics show that many orders are short-lived. For example, 25% of the cancellations occur within 0.05 second after the sampling time and 25% of the executions take place within 0.85 second.

From Panel B of Table I, we note that orders submitted by HFTs also have a shorter mean time to cancellation (19.72 seconds) than to execution (28.72 seconds). Moreover, irrespective of outcome, HFT orders are relatively short-lived, which is in line with previous evidence showing that HFT orders have a shorter duration than the orders of non-HFTs (Hagströmer and Nordén (2013)).

Figure 1 displays the cumulative probabilities of the cancellation and execution of all sample orders and HFT sample orders, respectively, expressed in event time (Panel A) and clock time (Panel B). Our sampling procedure implies a cap at 100 events on the *x*-axis of Panel A, whereas the *x*-axis of Panel B ranges up to 10 minutes. The curves show that both the probability of cancellation and the probability of execution are decreasing functions of time (measured in event time and clock time). Evidently, the cumulative cancellation probability is diminishing fast. Figure 1 also illustrates the result that HFTs, on average, cancel their orders faster and more often than non-HFTs.

INSERT FIGURE 1 ABOUT HERE

We present stock-level descriptive statistics in Table A1 of the Appendix. In short, the 29 large-cap stocks in our sample have market capitalizations from 5.4 billion SEK (0.8 billion USD) to 386.8 billion SEK (58.8 billion USD). The quoted bid–ask spread ranges from 1.64 basis points (bps) to 6.53 bps. As a point of comparison, Brogaard et al. (2014) report that US large- and mid-cap equities, on average, have quoted bid–ask spreads of 2.36 bps and 7.31 bps, respectively.

III. Analysis of Limit Order Cancellations

In this section, we develop an empirical model of limit order cancellations. The model controls for market and order conditions when we sample each order, as well as stock-specific effects. The key focus of the model is what happens *after* the time of sampling and how that influences the hazard rate of cancellation. We estimate the model for the full sample as well as for subsamples based on market conditions and the cross-section of trading firms.

A. Empirical Model

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 $\lambda_i(t)$, with t indicating the event time.[10] The time of sampling for each order is set to $t = 0$.

---

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

The predictions presented in Section I.C are based on the fact that the hazard rate of cancellation is a function of changes in the expected profit, which can be written as

$$\lambda_i(t) = f\big(\Delta\ln(\tilde{\pi}_{it})\big), \tag{10}$$

where $\tilde{\pi}_i$ is the level of the expected profit of order *i* at the time of sampling, and $\Delta\ln(\tilde{\pi}_{it})$ is the logarithmic change in expected profit from the time of sampling to time t.

Because the parameters of Equation (9) are unknown, we are unable to measure the level of expected profits, but (logarithmic) changes in expected profit can be approximated with a linear Taylor expansion. Consistent with the theoretical model, 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), respectively, are constant by definition. The Taylor expansion is a function of changes in the four model variables in Equation (9)[11]

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

We analyze the relation in Equation (10) by estimating a stratified PDHM with time-varying covariates and executions as censored observations.[12] Specifically, the model we estimate is:

---

[11] We acknowledge that high-order non-linear relations between the expected profit and the model variables could exist. In the interest of tractability, the Taylor expansion ignores such effects.
[12] We assume that the censoring is independent, meaning that orders censored at time t are representative of all orders that remain at time t with respect to survival, conditional on each level of the covariates.

$$\lambda_{ij}(t) = \lambda_{0j}(t)\exp\left[Level_{ij}\theta + \Delta\Pi_{ijt}\beta + Control_{ijt}\delta\right], \quad (12)$$

where $\lambda_{0j}(t)$ is an unspecified baseline hazard rate unique for stock j. The stock-specific baseline hazard rates are akin to stock fixed effects in a panel regression. The vector $Level_{ij}$ controls for order and market conditions at the time of sampling that potentially influence the resiliency of the order to shocks in the expected profitability. The variables included in $Level_{ij}$ are presented below. The term $\Delta\Pi_{ijt}\beta$ corresponds to $\Delta\ln(\tilde{\pi}_{it})$ in Equation (11). Our main variables of interest are included in the vector of time-varying covariates, defined as $\Delta\Pi_{ijt} = \left[\Delta BAS_{ijt}, \Delta Q_{it}^{Opp}, \Delta\overleftarrow{Q}_{it}, \Delta\vec{Q}_{it}\right]$, where changes in each variable for order i in stock j are defined relative to the time of sampling, $t = 0$. The vector $Control_{ijt}$, includes time-varying covariates that are not part of our theoretical model: $Firm\ Inventory$, $Same\text{-}Side\ MTF\ Trades$, and $Opposite\text{-}Side\ MTF\ Trade$. We denote the coefficient vectors corresponding to $Level_{ij}$, $\Delta\Pi_{ijt}$, and $Control_{ijt}$, respectively, as θ, β, and 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. We define *Firm Inventory* as the net trading volume (buy volume minus sell volume, expressed in numbers of shares) of the trading firm that submitted the order in question, calculated for each event during which the order is tracked. To account for the direction of trade, the net trading volume is multiplied by -1 if the order of interest is submitted on the ask side of the order book. The measure is based on trades executed on Nasdaq. If *Firm Inventory* influences order cancellation, we expect the variable to have a positive relation to the hazard rate of cancellation.

Van Kervel (2015) integrates the model of Sandås (2001) with the two-venue setting of Foucault and Menkveld (2008) and shows that trades at a competing venue work as signals about changes in the fundamental value. The author predicts, for example, that a buy trade at a competing venue signals an increase in the fundamental value and that it should, accordingly, lead to cancellations of ask-side orders and submissions of bid-side orders. We measure the trading volumes at the three multilateral trading facilities (MTFs) that are Nasdaq's main competitors 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 trading volume recorded at the three MTFs from the time of sampling until time t on the same side and the opposite side as the order in question, respectively. We expect *Same-Side MTF Trades* to have a positive relation and *Opposite-Side MTF Trades* a negative relation with the hazard rate of cancellation.

The vector $Level_{ij}$ includes covariates measured at the time of sampling and that therefore do not vary in the time dimension. It includes the model variables, defined as in Section II.D, to control for orders with higher expected profits being more resilient to shocks, and thus having a lower hazard rate of cancellation.[13] Similar to Hasbrouck and Saar (2009), we include *Lagged Absolute Return* (the absolute five-minute stock return prior to the sampling time) and *Lagged Volume* (the number of shares traded during the five minutes prior to the sampling time, relative to the total depth) to control for market conditions at the time of sampling. Following

[13] Because the three quantity variables Q_{ij}^{Opp}, $\overleftarrow{Q}_{ij}$, and $\vec{Q}_{ij}$ sum to unity, the model cannot be estimated if all of them are included. For this reason, we exclude Q_{ij}^{Opp} from all the estimations. The same issue does not apply to the variables in $\Delta\Pi_{ijt}$.

the prediction of Liu (2009), stating that the incentives to monitor a limit order increase with the order size, we also control for *Order Quantity* (defined as in Section II.D). Finally, we control for the propensity of the trading firm behind each order to benefit from internal matching. We define the variable *Internal Match Rate* as the proportion of passively executed limit orders that are internally matched during April 2014, the month preceding our sample month, for the trading firm in question. We expect firms with a higher propensity for internal matching to be less inclined to have longer order durations.

All the variables in the vector $Control_{ijt}$, as well as the level variables *Lagged Absolute Return* and *Lagged Volume*, are standardized within each stock to have zero mean and unit variance.

Our empirical model relates closely to the PDHM with time-varying covariates suggested by Hasbrouck and Saar (2009). It also relates to the parametric model of time to execution of Lo, MacKinlay, and Zhang (2002) and the competing risk model of cancellation and execution times of Chakrabarty et al. (2006). In contrast to Hasbrouck and Saar (2009) but similar to Lo et al. (2002), we estimate a pooled model for all stocks and orders in our sample.

Table A.II in the Appendix presents a correlation table for all the explanatory variables in the empirical model. Among the time-varying covariates, changes in *Front of Queue Quantity* and changes in *Back of Queue Quantity* have a relatively high correlation (0.37). Furthermore, changes in *Bid-Ask Spread* are positively related to changes in all the order book quantity variables (with correlations ranging from 0.34 to 0.36). None of these correlations is high enough to cause multicollinearity problems.

B. Main Results

We report the PDHM coefficient estimates in Table II. To facilitate the analysis of economic significance, we also include the hazard rate change (HRC) of each variable. Similar to a marginal effect 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.[14] In addition, we report the survival time change (STC), which is the corresponding effect expressed in event time.[15]

INSERT TABLE II ABOUT HERE

As a useful benchmark for our subsequent analysis of the determinants of cancellations, we first estimate the empirical model using only the level variables, that is, restricting the β and δ coefficient vectors to zero. The results are presented to the left in Table II, under the heading of restricted model. Overall, the estimates are in line with our expectations. The hazard rate of cancellation is negatively related to the expected profit as indicated by the negative coefficients of *Bid-Ask Spread* and *Back of Queue Quantity*. The fact that the *Front of Queue Quantity* coefficient estimate is significantly higher than that of *Back of Queue Quantity* shows that adverse selection

---

[14] To obtain the HRC of an explanatory variable X_{ijt} 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.

[15] Specifically, the STC is the change in the restricted mean survival time following a one standard deviation increase in the explanatory variable in question. The restricted mean survival time is based on survival functions for a median order in each stock *j*, $S_j(t) = exp[-\widehat{\Lambda}(t, \hat{\theta}, \hat{\beta}, \hat{\delta}, \hat{\lambda}_{0j}, X^m)]$, where $\widehat{\Lambda}$ denotes the cumulative hazard rate based on the original estimation of $\hat{\theta}$, $\hat{\beta}$, and $\hat{\delta}$. The variable X^m shows the median values of the covariates in the empirical model. For details on the calculation of the baseline hazard rate, $\hat{\lambda}_{0j}$, and the cumulative hazard rate (see Therneau and Grambsch (2000), pp. 266-267). The restricted mean survival time, $\hat{\mu}_j$, is given by the area under the survival curve, from time zero to 100. The reported STC is the mean over all stocks *j*.

costs have a significant influence on the hazard rate of cancellations. In economic terms, it is clear from the HRCs and STCs that the model variables are more important than the other level variables. A standard deviation increase in *Bid-Ask Spread* and *Back of Queue Quantity*, respectively, lowers the hazard rate of cancellation by 40.9% and 25.4% (increases the order survival time by 14.3 and 8.3 events), respectively, compared to the much lower results for all the other level variables.

We present the results of the model including both level and time-varying variables on the right side of Table II (under the full model heading). The coefficient estimates for the time-varying covariates of the model variables are all consistent with our predictions. The hazard rate is decreasing with changes in *Bid-Ask Spread* and *Back of Queue Quantity* and increasing with *Opposite-Side Quantity Changes*. The coefficients are statistically significant at the 1% level. This shows that negative shocks to the expected revenue of a limit order increase the probability of the order being canceled in the next instant.

The coefficient for changes in *Front of Queue Quantity* is not significantly different from zero, but more importantly, it is significantly higher than that for changes in *Back of Queue Quantity.* A Wald test of the null hypothesis that the coefficients for $\Delta\overleftarrow{Q}_{ijt}$ and $\Delta\vec{Q}_{ijt}$ are equal results in a *t*-value of 34.11, implying that the null can be rejected with a high level of confidence.[16] According to our model, the coefficient for changes in *Front of Queue Quantity* is associated with liquidity suppliers factoring in the adverse selection costs associated with their

[16] The corresponding test in levels, with the null hypothesis that the coefficients of $\overleftarrow{Q}_{ij0}$ and $\vec{Q}_{ij}$ are equal, yields a *t*-value of 75.66.

orders. For example, if *Back of Queue Quantity* shrinks, ceteris paribus, the liquidity supplier infers that the fundamental value is moving closer to the limit order price, undermining the expected revenue of the order. If, instead, *Front of Queue Quantity* shrinks, the same revenue reduction effect applies, but it is offset by a reduction in expected adverse selection costs. The latter effect is due to the limit order in question now requiring a smaller market order to be executed and the liquidity supplier then expects the market order to have less price impact.

The economic significance of the model variables is high. For example, a standard deviation increase in the changes of *Back of Queue Quantity* lowers the hazard rate by 49.4%, which corresponds to an STC of 17.5 events. An event corresponds to about 0.4 second on average, meaning that 17.5 events lasts about 7 seconds (17.5 events × 0.4 second/event). As the average order lifetime is 26 seconds (see Table I, Panel A), the effect corresponds to about a 27% longer order duration (7 seconds/26 seconds). The low HRC for changes in *Front of Queue Quantity* should not be interpreted as the variable is unimportant. On the contrary, it implies that the adverse selection cost effect is high enough to neutralize the revenue effect seen for changes in *Back of Queue Quantity*.

The full model coefficients for the model variables in levels have the same signs as the restricted model estimates. A notable difference is that the level effect of the *Front of Queue Quantity* variable is not statistically significant in the model setup with time-varying covariates, while it is significantly positive in the restricted model.

The results for the other time-varying covariates are in line with the expectations and statistically significant, but their economic significance is relatively low. The sign of the coefficient of *Firm Inventory* is consistent with limit order traders applying price pressure to mean-revert their inventory (see also the discussion in Section IV). The small economic effect of this private value variable lends support to the theoretical model's focus on common values. The signs of the coefficients for variables capturing common values indicated at the competing trading venues (*Same-Side MTF Trades* and *Opposite-Side MTF Trades*) are consistent with the predictions of van Kervel (2015). However, the HRC for these variables is relatively low, showing 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.

We conclude that all our predictions with respect to the model variables find empirical support in the stratified PDHM with time-varying covariates. The results are robust to alternative explanations for cancellation, such as private values, monitoring costs, trading venue competition, and internal matching. In Table A.III of the Appendix, we repeat the analysis on more restricted order samples. We show there that our results are also robust to (i) a restriction to orders that ex post have a lifetime of less than 2 seconds (adhering to the definition of fleeting orders of Hasbrouck and Saar (2009)) and (ii) a restriction to orders that are sampled at the time of submission.

We proceed by analyzing cross-sectional variation in our sample of orders. We split the sample first by market conditions (price discreteness) and then by trading firms (comparing HFTs' activities to those of others).

C. Results for Subsamples Based on Price Discreteness

Price discreteness constrains price competition and elevates the importance of time priority. O'Hara, Saar and Zhong (2015) show that the relative tick size influences the biodiversity of traders, as well as limit order behavior. Their empirical investigation shows that a larger relative tick size leads to greater quoted depth and the greater profitability of HFT market making. Moreover, Yao and Ye (2017) document that a large relative tick size benefits HFTs over non-HTFs in achieving time priority for limit orders. Conversely, a small relative tick size benefits non-HFTs in establishing price priority. These results with respect to fast traders are consistent with those of Budish et al. (2015), who argue that price discreteness leads to competition in latency.

To investigate if the determinants of 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 subsequent analyses, we also split the sample by the trading firm 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 sizes. Table III reports the coefficient estimates and associated HRC and STC values.

INSERT TABLE III ABOUT HERE

When the tick size is binding (on the left in Table III), changes in the *Back of Queue Quantity* and *Front of Queue Quantity* variables have a more negative influence on the hazard rate of cancellations than they have in the full sample. A binding tick size restricts liquidity suppliers to setting their quotes symmetrically around the fundamental value (as shown by Anshuman and Kalay (1998)). Our results are consistent with the depth imbalance then becoming more influential in gauging where the fundamental value is inside the bid–ask spread. Furthermore, the result for changes in *Front of Queue Quantity* shows that its effect on the order revenue dominates that of adverse selection costs.

Changes in *Bid-Ask Spread* have no significant influence on the hazard rate when the tick size is binding. In this setting, variation in *Bid-Ask Spread* is by definition restricted to cases in which the opposite-side price moves away from the limit order price. An increasing spread generally implies greater expected revenue, but when the tick size was binding the previous instant, liquidity suppliers could consider the prospect of earning a two-tick spread unrealistic, leaving the hazard rate unchanged.

For sample orders where the tick size is not binding, the results (reported in the three rightmost columns of Table III) are more consistent with the full-sample results. Notably, the effects of the *Front of Queue Quantity* variables are positive and significant. This finding is consistent with a more important *Front of Queue Quantity* effect on the adverse selection cost when the tick size is not binding.

Whether the tick size is binding or not, the HRCs and STCs show that the economic significance for the model variables is higher than for the variables outside the model. In addition, the Wald test for equality between the coefficients for *Front of Queue Quantity* and *Back of Queue Quantity* is rejected at the 1% level for both levels and changes, in both subsamples.[17] Hence, the order queue position is important for cancellations irrespective of whether the tick size is binding or not.

[17] When the tick size is binding, the Wald test results in a *t*-value of 57.08 for the level variables and a *t*-value of 18.70 for the change variables. When the tick size is not binding, the corresponding *t*-values are 47.24 for the level variables and 20.87 for the change variables.

D. 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 largely 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 have. Brogaard et al. (2014) find that HFTs lead the price discovery process by updating their quotes at high speed when new information emerges in the marketplace. Brogaard, Hagströmer, Nordén, and Riordan (2015) show that market-making HFTs have a strong tendency to invest in colocation services, allowing them to monitor and respond to changes in fundamental values with low latency. Hagströmer, Nordén, and Zhang (2014) find that HFTs are more concerned than non-HFTs with the tradeoff between waiting costs and the cost of immediate execution. Goldstein, Kwan and Philip (2017) find that HFTs are more likely than non-HFTs to adjust their trading behavior to order book depth imbalances.

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 a binding tick size submitted by HFTs and by other liquidity suppliers (non-HFTs) and orders with a non-binding tick size submitted by HFTs and non-HFTs. We estimate the stratified PDHM with time-varying covariates for each subsample and present the results in Table IV. Panel A shows the results for the subsamples with a binding tick size and Panel B reports the findings for the subsamples where the tick size is non-binding.

INSERT TABLE IV 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, whether the tick size is binding or not. Second, the economic influence of the model variables (in terms of HRCs and STCs) 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 changes in *Back of Queue Quantity*. For this variable, when the tick size is binding, the HFT coefficient is −4.277, whereas that for non-HFTs is −0.650. The corresponding HRC values are −88.7% and −29.1%. The strong emphasis of *Back of Queue Quantity* among HFTs signals that they use the depth imbalance to track fundamental values and that they do so to a larger extent than other liquidity suppliers. The coefficient estimates for changes in *Front of Queue Quantity* for HFTs and non-HFTs convey the same story. When the tick size is not binding, the same story holds but is less pronounced.

Our results are consistent with HFTs investing heavily in infrastructure, allowing them to monitor the limit order book and their queue positions at a high frequency. According to the model view of the expected profit of a limit order, HFTs respond relatively strongly to factors that could undermine the profitability of their limit orders. That HFTs pay more attention to order book information than non-HFTs is consistent with the findings of Hagströmer et al. (2014) and Goldstein et al. (2017).

We also find support for the common notion that HFTs manage their inventory risk at high frequency, demonstrated empirically by Brogaard et al. (2015). The positive coefficient estimate for *Firm Inventory* shows that, when an HFT increases a long position in a stock, the

HFT hazard rate of canceling outstanding buy-side limit orders also increases. The hazard rate of canceling sell-side limit orders falls. This result 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 inventory levels. We do not find any effect of *Firm Inventory* for non-HFTs.

IV. Supply- and Demand-Driven Limit Order Strategies

The model we analyze in this paper is based on the view that the limit order is a tool used exclusively by market makers with no intrinsic motive for trading other than to profit from the trading process itself. Investors who demand liquidity are assumed to be impatient. Though our empirical results lend support to the theoretical model, this view is not uncontroversial. 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. Consistent with the author's empirical evidence, ample theoretical models show how liquidity demanders optimize the choice between limit orders and market orders (e.g., Harris (1998), Foucault et al. (2005), Kaniel and Liu (2006), Roşu (2009)). Ranaldo (2004), for example, presents empirical evidence in support of the theoretical models. Specific to cancellations, in an empirical study of Nasdaq stocks in October 2004, Hasbrouck and Saar (2009) find that cancellations of short-lived limit orders show patterns consistent with liquidity demand strategies. In light of this theoretical and empirical evidence, how can our liquidity supply-related explanations of cancellations be reconciled with the literature?

We present three arguments to support our supply-side approach. First, any limit order trader, regardless of trading motive, is arguably exposed to the common values captured by the model. For example, an institutional investor seeking to sell a block of shares can derive a private value of trading, due to the opportunity cost of holding the block. Nevertheless, the institutional investor also shares the common values captured by the model. If the institution's limit order executes against a large incoming market order, the institution incurs the same adverse selection costs as a market maker would. The price impact of the market order decreases the institution's profit margin for the trade. O'Hara (2015) supports the view that liquidity demanders actively manage their passive orders. Describing a block order executed in current markets, the author writes that an "algorithm turns the parent order into scores of limit orders placed in layers on the book (or across many books), with orders canceled and updated as trading progresses" (p. 264). We argue that, relative to market makers, traders with private values place less emphasis on the model variables, but in principle, the model predictions should hold for all limit order traders alike.

Second, our empirical analysis controls for the inventory management of individual trading firms. The coefficient reported for *Firm Inventory* in Table II indicates that limit order traders, on average, apply price pressure to avoid adding to their current inventory. The evidence of price pressures is consistent with limit order traders being predominantly market makers trying to mean-revert their inventory (as described by Hendershott and Menkveld (2014)), rather than position-building liquidity demanders. The evidence of mean-reverting inventory for the full sample is likely driven by HFTs, for which we find a strongly significant positive coefficient in both panels of Table IV. We do not find significant evidence of position-building behavior (which would correspond to a negative coefficient for *Firm Inventory*) in any of our subsample analyses.

Third, in models of order choice, traders adapt their trading aggressiveness in response to changes in market conditions. For example, Roşu (2009) allows the traders to cancel an order only to submit a more aggressive, potentially marketable, limit order. Harris (1998) shows that an optimal strategy in the presence of a trading deadline can be to gradually increase trading aggressiveness. The level of detail in our data set allows us to observe the prevalence of the different types of order price revisions described by the theoretical literature. In the analysis presented below, we show that only 15.6% of the limit order cancellations are associated with order price revisions. Most of the revisions make the order less aggressive and no more than 1.1% of all cancellations correspond to the specific liquidity demand strategies outlined by Hasbrouck and Saar (2009).

We define order price revisions as cases in which we observe, within half a second after a cancellation, a new limit order submission on the same side of the market but at a different price, posted by the same trading firm for the same stock. On Nasdaq, and many other exchanges, traders can automate order price revisions by submitting a so-called replace order. The replace order bundles the cancellation of an existing limit order with the simultaneous submission of a new order of the same quantity at a different limit price. By allowing a delay in the subsequent order, we allow order price revisions to be either automatic or manual.

We observe order price revisions in the data and report the frequencies of the different types of revisions in Table V. Panel A shows that 15.6% of all cancellations in our sample are part of order price revisions. The other cancellations (84.4%) either are standalone messages

(43.3%) or are followed by another cancellation (20.5%), by a submission on the opposite side of the book (10.8%), or by a submission on the same side and at the same price (8.8%). None of these other cancellation types is consistent with the liquidity demand strategies described by the order choice literature. Thus, the subsequent discussion focuses on order price revisions.

INSERT TABLE V ABOUT HERE

Order price revisions can be categorized by the aggressiveness of the new limit order relative to the canceled order. We label revisions that lead to immediate execution as marketable. We mark all other price revisions as either more aggressive or less aggressive. For example, if the limit price of a buy order is adjusted upward, we mark it as more aggressive. In our sample, 10.6% of all cancellations (roughly two thirds of the order price revisions) decrease the aggressiveness of the limit order. The order price revisions conforming to the order choice literature discussed above (Harris (1998), Roşu (2009)) are only 5.1% of all cancellations (0.5% marketable and 4.6% more aggressive). However, Raman and Yadav (2014) show that all types of price revisions may be rational from a supply-side point of view, as part of inventory management strategies.

Hasbrouck and Saar (2009) hypothesize that many cancellations are due to liquidity-demanding strategies. Consider a liquidity demander with a limit order to buy posted in the limit order book. Hasbrouck and Saar's (2009) "chasing hypothesis" states that, if the best ask price increases, the buy order price can be revised upwards to maintain the same execution probability. According to their "cost-of-immediacy hypothesis", if the best ask decreases, the cost of crossing the spread shrinks, which can trigger a marketable order price revision. We assess these hypotheses by conditioning the order price revisions on changes in the opposite-side price in the half-second preceding the cancellation (or during the order lifetime, if it is shorter than half a second). In the three rightmost columns of Table V, Panel A, we report the fraction of all cancellations occurring when the opposite price moves away (as in the chasing hypothesis), moves closer (as in the cost-of-immediacy hypothesis), or stays the same. We find that 1.0% of all cancellations correspond to the chasing hypothesis (more aggressive conditional on the opposite price moving away) and only 0.1% are consistent with the cost-of-immediacy hypothesis (marketable when the opposite price moves closer).[18]

The low incidence of the strategies outlined by Hasbrouck and Saar (2009) can be due to differences in market structure between the United States and Sweden, and between their sample period (October 2004) and ours (May 2014). The sample analyzed in this paper is also broader in the sense that we do not restrict orders by their duration. Hasbrouck and Saar analyze "fleeting orders," of order duration shorter than 2 seconds. To control for differences in this dimension, we consider a subsample that meets their definition of fleeting orders. In Table V, Panel B, we repeat the analysis of order price revisions for fleeting orders, which is almost one-fifth of all cancellations (14,107 out of 73,442). We find that only 0.4% of the fleeting orders correspond to the chasing hypothesis, and virtually none (0.0%) are consistent with the cost-of-immediacy hypothesis. We run the same analysis on HFT orders, which constitute roughly one-third of all canceled sample orders (24,324 out of 73,442; see Table V, Panel C). The results with respect to Hasbrouck and Saar's hypotheses are the same as for fleeting orders: only 0.3% of all

[18] Hasbrouck and Saar (2009) also present a third liquidity demand-related strategy. The so-called "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 in the order book, this order type is mechanically excluded.

orders are explained by the liquidity demand strategies they discuss. We also note that order price revisions are less prevalent for fleeting orders and HFT orders than for the full sample (13.4% and 10.3%, respectively, compared to 15.6% overall).

The conclusion of this section is that passive liquidity demand strategies exist but do not dominate the sample. An interesting direction for future theoretical research would be to model how liquidity supply depends on both common and private values. However, the distinction between liquidity suppliers and demanders in Sandås' (2001) model delivers predictions that are consistent with the data and exhibits stronger economic significance than competing explanations for cancellations.

V. Conclusions

We derive empirical predictions for the determinants of limit order cancellations. The predictions rely on limit orders being free options offered to the market in the expectation of intermediation profits. We postulate that negative shocks to expected profits can trigger cancellations. A novel prediction, driven by adverse selection costs, is that changes in the order queue position are an important determinant of cancellations. Our predictions find strong support in the data.

Overall, our study points to frequent order cancellations being a benign feature of modern market making. Competitive liquidity suppliers closely monitor even marginal fluctuations in fundamental values and cancel orders that are not expected to be profitable. HFTs, known to invest in low-latency technology, adhere particularly strongly to the economics of limit order cancellations. Our results show that even though liquidity demanders rely extensively on limit orders in modern markets, the economics of liquidity supply are powerful in explaining the triggers of limit order cancellations. Our findings imply that policies banning, discouraging, or taxing limit order cancellations are likely to increase the costs of market making, which in turn would lead to lower market liquidity.

Appendix

A1. Summary Statistics for Sample Stocks

Table A.I presents summary statistics for the sample stocks. The variable *Market capitalization* is in millions of Swedish krona (MSEK), using the closing prices of April 30, 2014. On that date, 1 MSEK was worth roughly 110,000 EUR, or 152,000 USD. From the second column in Table A.I, we note that market capitalization varies among stocks, the highest value being 386,786 MSEK (HM B) and the lowest being 5,429 MSEK (NOKI SEK).

INSERT TABLE A.I ABOUT HERE

The other variables in Table A.I use data from all trading days during April 2014. The daily trading volume and daily turnover statistics include continuous trading and the opening and closing call auctions on Nasdaq but exclude trading at other venues. The daily turnover is the average fraction of market capitalization traded on a daily basis. The daily trading volume varies between 75 MSEK (MTG B) and 860 MSEK (ERIC B), while the daily turnover is between 0.16% (SHB A) and 4.27% (NOKI SEK) of the market capitalization.[19]

We present three spread measures, all measured at the time of trade and denoted in basis points (bps). When calculating the spread measures, we exclude trades in the opening and closing call auctions, block trades worth more than 1 MSEK, and trades on venues other than

[19] Data on market capitalization are obtained from Nasdaq's website.

Nasdaq. The variable *Quoted Spread* is half the difference between the BBO 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 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).[20] The term *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). The variable *Effective Spread* is the trade value-weighted average absolute difference between the transaction price and the spread midpoint, relative to 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).

The variable *Depth at BBO* is the average of the volumes available at the BBO prices, or the average volume required to move the price in either direction. Typically, the average depth at BBO is around 1 MSEK but varies substantially across stocks, in the range between 0.18 MSEK (MTG B) and 2.76 MSEK (ERIC B). As for the spread measures, the depth at the BBO is observed at times of trades. The measure *Volatility* is the average 10-second squared basis point returns and *Order-to-Trade Ratio* is the number of trades divided by the number of quote updates at the BBO.

[20] Brogaard et al. (2015) compare the quoted spread levels in their sample to spreads for US stocks studied by Brogaard et al. (2014). Accordingly, a 2-bp spread is comparable to US large-cap stocks, while a 6-bp spread is slightly higher than for an average US mid-cap stock.

A2. Correlations

Table A.II displays the sample correlations between the explanatory variables from the PDHM according to Equation (12). Specifically, we report the pooled Pearson correlation coefficient, estimated across stocks, trading days, orders, and time. For the variable definitions, see the captions of Tables II and A.II.

INSERT TABLE A.II ABOUT HERE

The variable *Front of Queue Quantity* has a mechanically positive relation to *Order Quantity*, which could potentially introduce multicollinearity problems. However, the sample correlation is only 0.29, making it unlikely that multicollinearity decreases the significance of the *Front of Queue Quantity* coefficient estimates.

A3. Subsample Analysis

Table A.III presents the results for the same stratified PDHM, according to Equation (12), as in Table II for a subsample of orders with lifetimes shorter than 2 seconds (fleeting orders) and a subsample of orders sampled at the time of submission (orders at submission). Overall, for our model variables, the results for each subsample are similar to the full sample results in Table II (full model).

INSERT TABLE A.III ABOUT HERE

For the fleeting orders, on the left in Table A.III, the coefficients for *Bid-Ask Spread* and *Back of Queue Quantity* in levels are both significantly negative. However, the two coefficients are lower and the associated HRCs and STCs are closer to zero for the fleeting orders than for the corresponding coefficients for all orders. Apparently, shocks in levels of *Bid-Ask Spread* and *Back of Queue Quantity* have smaller impacts on the hazard rate of fleeting orders than the corresponding hazard rates of all orders. The coefficient for *Front of Queue Quantity* is significantly positive in the fleeting order sample, while it is not significantly different from zero in the full sample. Although the associated HRC (and STC) is rather small, the results imply that adverse selection costs are relatively more important for fleeting orders than for all orders. Consistent with our theoretical model, the *Front of Queue Quantity* coefficient is significantly larger than the *Back of Queue Quantity* coefficient for fleeting orders as well.

The coefficients for changes in *Bid-Ask Spread*, *Opposite-Side Quantity*, and *Back of Queue Quantity* have the same signs and significance levels for fleeting orders as for all sample orders. In addition, the coefficient for changes in *Front of Queue Quantity* is not significantly different from zero in any specification.

For the orders sampled at the time of submission, the coefficients for the model variables, in both levels and time-varying changes, have similar signs and significance levels as for all orders. Interestingly, while the results for our model variables persist irrespective of the time of sampling (at submission or using our sampling technique outlined in Section II.C), the corresponding results for the control variables do not. On the right in Table A.III, we note that the coefficients for the level variable *Order Quantity* and the time-varying covariates *Firm Inventory*, *Same-Side MTF Trades,* and *Opposite-Side MTF Trades* are all not significantly different from zero. We interpret these results as strong evidence in favor of our theoretical model for explaining cancellations, irrespective of sampling technique; relative alternative explanations, such as private values; monitoring costs; and trading venue competition.

References

Anshuman, V. R., A. Kalay, 1998, Market making with discrete prices, *Review of Financial Studies* 11(1) 81–109.

Baron, M., J. Brogaard, B. Hagströmer, and A. Kirilenko, 2017, Risk and return in high-frequency trading, forthcoming in *Journal of Financial and Quantitative Analysis*.

Baruch, S., and L. R. Glosten, 2016, Tail expectations, imperfect competition, and the phenomenon of flickering quotes in limit order book markets, Working Paper.

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

Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan, 2015, Trading fast and slow: Colocation and market liquidity, *Review of Financial Studies* 28, 3407–3443.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, High frequency trading and price discovery, *Review of Financial Studies* 27, 2267–2306.

Budish, E., P. Cramton, and J. Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics,* 130(4) 1547–1621.

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

Cont, R., Kukanov, A. and Stoikov, S., 2014. The price impact of order book events. *Journal of Financial Econometrics* 12(1), 47–88.

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

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., O. Kadan, and E. Kandel, 2005, Limit order book as a market for liquidity, *Review of Financial Studies* 18(4), 1171-1217.

Foucault, T., R. Kozhan, and W. W. Tham, 2017, Toxic arbitrage, *Review of Financial Studies* 30(4), 1053-1094. .

Foucault, T., and A. J. Menkveld, 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? *Journal of Finance* 49(4), 1127–1161.

Goldstein, M., A. Kwan, and R. Philip, 2017, High-frequency trading strategies, Working Paper.

Gould, M.D. and Bonart, J., 2016. Queue imbalance as a one-tick-ahead price predictor in a limit order book. *Market Microstructure and Liquidity* 2(02), doi: 10.1142/S2382626616500064.

Hagströmer, B., 2017, Overestimated effective spreads: Implications for investors, Working Paper.

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

Hagströmer, B., L. Nordén, and D. Zhang, 2014, How aggressive are high-frequency traders? *Financial Review* 49, 395–419

Handa, P., R. Schwarz, and A. Tiwari, 2003, Quote setting and price formation in an order driven market, *Journal of Financial Market* 6, 461–489.

Harris, L., 1998, Optimal dynamic order submission strategies in some stylized trading problems, *Financial Markets, Institutions & Instruments* 7(2), 1–76.

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

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

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

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

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

Kaniel, R., and H. Liu. (2006), So what orders do informed traders use? *Journal of Business* 79(4), 1867–1913.

Lee, E. T., and J. W. Wang, 2003, *Statistical Methods for Survival Data Analysis,* 3rd edition (Hoboken, NJ: Wiley).

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

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

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

O'Hara, M., 2015, High frequency market microstructure, *Journal of Financial Economics,* 116(2), 257–270.

O'Hara, M., G. Saar, and Z. Zhong, 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., and P. Yadav, 2014, Liquidity provision, information and inventory management in limit order markets: An analysis of order revisions, Working Paper.

Ranaldo, A., 2004, Order aggressiveness in limit order book markets, *Journal of Financial Markets* 7, 53–74.

Roşu, I., 2009, A dynamic model of the limit order book, *Review of Financial Studies* 22(11), 4601–4641.

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

Seppi, D. J., 1997, Liquidity provision with limit orders and a strategic specialist, *Review of Financial Studies* 10(1), 103–150.

Stoikov, S., 2017, The micro-price: A high frequency estimator of future prices, Working Paper.

Therneau, T. M., and P. M. Grambsch, 2000, *Modeling Survival Data: Extending the Cox Model,* 1st edition. (Springer-Verlag, New York, NY).

US Securities and Exchange Commission, 2010, Concept release on equity market structure; proposed rule, *Federal Register* 75(13), 3593–3614.

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

Yao, C., and M. Ye, 2017, Why trading speed matters: A tale of queue rationing under price controls, forthcoming, *Review of Financial Studies*.

Yueshen, B. Z., 2014, Queuing uncertainty in limit order market, Working Paper.

Figure 1: Limit Order Durations

This figure shows the fraction of sample orders that are canceled (solid line) or executed (dashed line) after a fixed amount of time after the time of sampling, with time expressed in event time in Panel A and clock time in Panel B. A unit of event time elapses whenever the prices or volumes at the BBO of the stock in question change. The full sample includes 138,486 (all) limit orders, and a subset holds 38,656 limit orders submitted by HFTs (HFT), sampled from the OMXS 30 stocks in May 2014 on Nasdaq. The HFTs are trading firms that are members of FIA EPTA. For details on the sampling procedure, see Section II.C.

Panel A: Event time

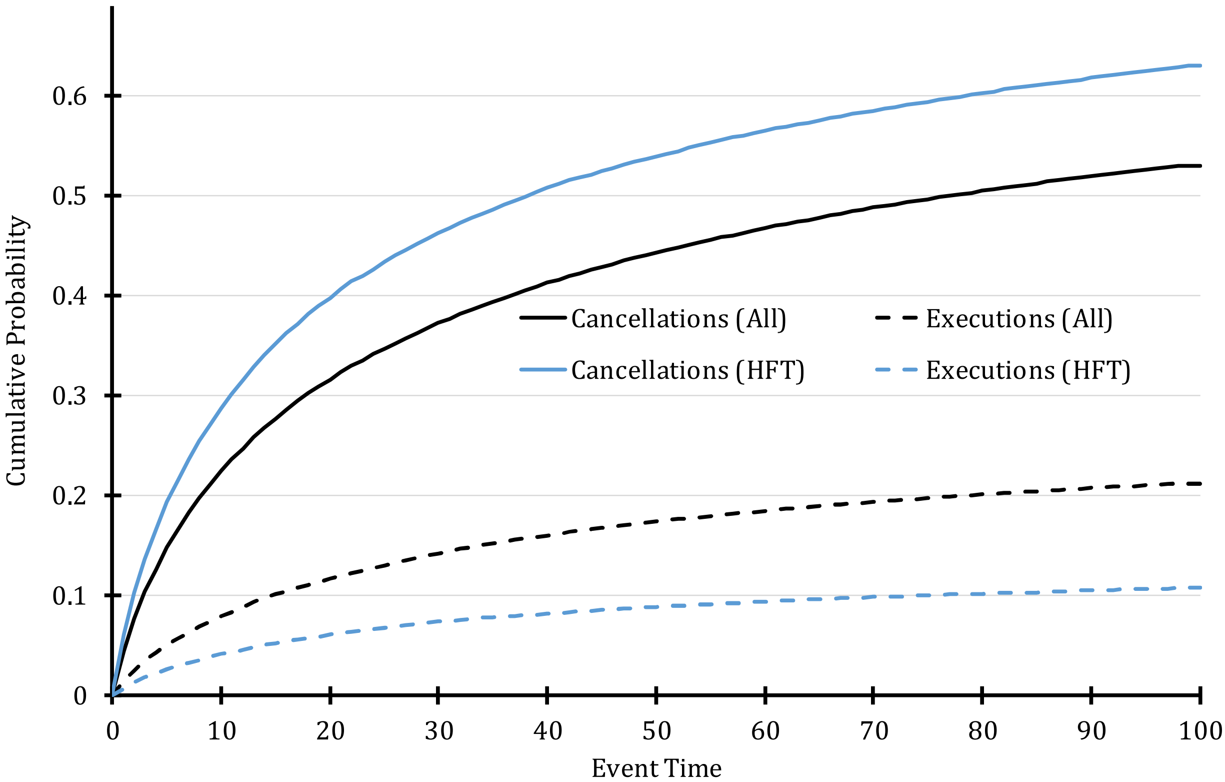

Panel B: Clock time

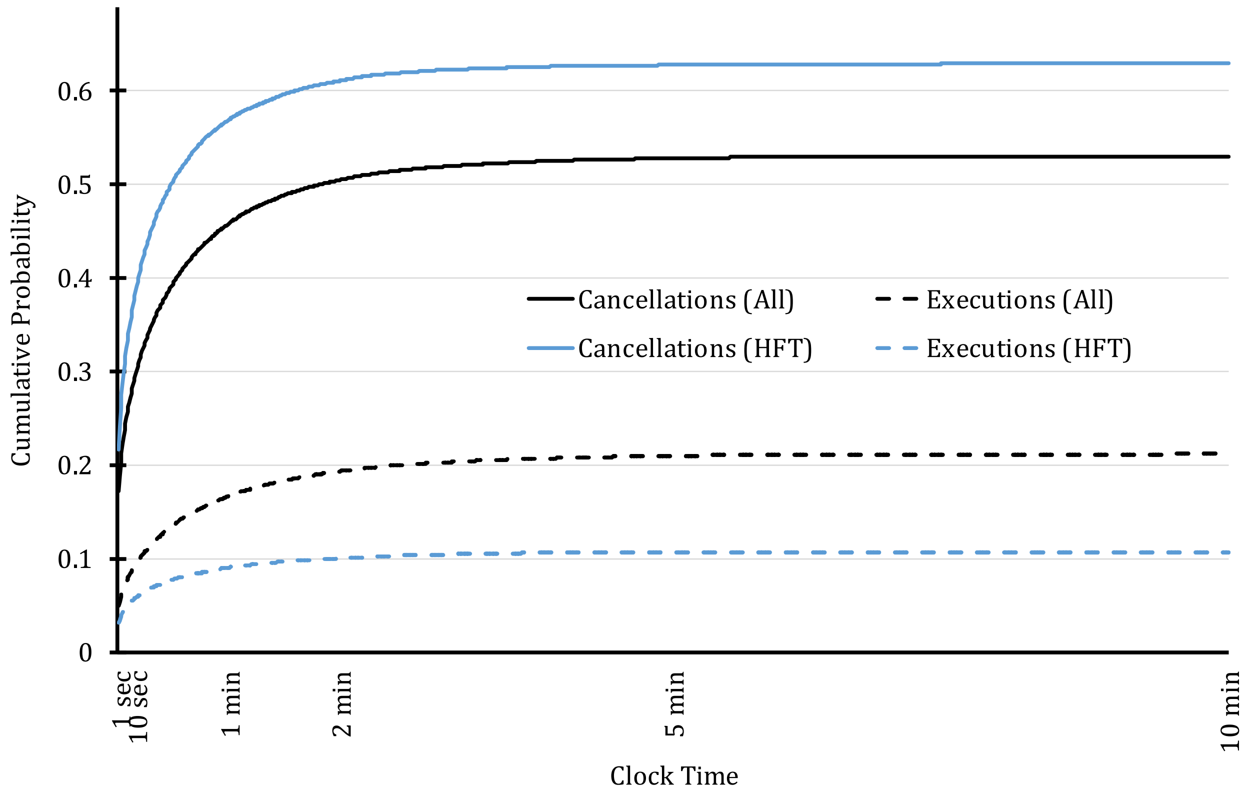

Table I: Sample Order-Level Descriptive Statistics

This table presents order-level descriptive statistics for our model variables at the time of sampling. The full sample includes 138,486 limit orders (all orders in Panel A) and a subset holds 38,656 limit orders submitted by HFTs (HFT orders in Panel B), sampled from the OMXS 30 stocks in May 2014 on Nasdaq. The HFTs are trading firms that are members of FIA EPTA. For details on the sampling procedure, see Section II.C. The variable *Bid-Ask Spread* is the difference between the limit order price and the best price on the opposite side of the book, divided by the midpoint, measured at the time of sampling and expressed in basis points. The following variables are measured at the time of sampling and expressed as fractions of the volumes posted at BBO: *Back of Queue Quantity* is the volume posted at the same price but with a lower time priority than the order of interest; *Front of Queue Quantity* is the volume posted at the same price but with a higher time priority than the order of interest, plus the order's own volume; *Opposite-Side Quantity* is the volume posted at the best price level on the opposite side of the order in question; *Order Quantity* is the volume of the order in question; and *Order Duration* is the event time from order sampling to order expiry, where expiry is defined as the earliest of cancellation, execution, and surviving 100 events, conditional on the expiry reason. A unit of event time elapses whenever the prices or volumes at the BBO of the stock in question change.

Panel A: All orders			Percentile				
Variable	Mean	St. Dev.	5th	25th	50th	75th	95th
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 (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

Panel B: HFT orders			Percentile				
Variable	Mean	St. Dev.	5th	25th	50th	75th	95th
Bid-Ask Spread (bps)	7.04	3.13	3.56	5.28	5.99	8.96	12.52
Quantity variables expressed as fractions of total depth							
Back of Queue Quantity	0.21	0.21	0.00	0.02	0.15	0.34	0.65
Front of Queue Quantity	0.26	0.20	0.04	0.10	0.21	0.37	0.68
Opposite-Side Quantity	0.52	0.25	0.10	0.34	0.54	0.72	0.91
Order Quantity	0.08	0.09	0.02	0.04	0.06	0.10	0.25
Order duration conditional on outcome (seconds)							
Cancellation	19.72	39.47	0.00	0.03	3.89	21.73	91.70
Execution	28.72	49.64	0.00	0.11	7.09	33.66	134.21
Surviving 100 events	79.62	76.67	8.22	29.37	56.66	103.84	228.41

Table II: PDHM of Order Cancellations

This table presents estimates of the following stratified PDHM with time-varying covariates for the full model, according to Equation (12), and a restricted version without time-varying covariates:

$$\lambda_{ij}(t) = \lambda_{0j}(t)\exp\left[Level_{ij}\theta + \Delta\Pi_{ijt}\beta + Control_{ijt}\delta\right],$$

where $\lambda_{ij}(t)$ is the hazard rate of cancellation of order i over the next instant at time t in stock j and $\lambda_{0j}(t)$ is an unspecified baseline hazard rate for stock j. The vector $Level_{ij}$ includes the following variables at the time of sampling: *Bid-Ask Spread*, *Back of Queue Quantity*, *Front of Queue Quantity*, and *Opposite-Side Quantity* (defined as in Table I), *Order Quantity* (the volume of the sampled order, relative to the total depth), *Lagged Absolute Return* (the absolute five-minute stock return prior to the sampling time), *Lagged Volume* (the number of shares traded during the five minutes prior to the sampling time, relative to the total depth), and *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). The vector $\Delta\Pi_{ijt}$ includes changes in the model variables relative to the time of sampling. The vector $Control_{ijt}$ includes the following time-varying covariates: *Firm Inventory* is the stock-specific net trading volume (buy volume minus sell volume, expressed in numbers of shares) of the trading firm that submitted the order in question. To account for the direction of trade, the net trading volume is multiplied by -1 if the order of interest is submitted on the ask side of the order book. The variables *Same-Side MTF Trades* and *Opposite-Side MTF Trades* measure the trading volume recorded at the three MTFs from the time of sampling until time s, on the same and opposite sides as the order in question. The sample includes 138,486 limit orders sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. Each sample order is tracked over up to 100 events, resulting in 6,141,702 observations. For each estimate, the t-values are reported in parentheses and ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered on order i. Here, HRC is the change in the cancellation hazard rate following a one standard deviation change in the explanatory variable in question, holding all other explanatory variables constant, and STC is the change in the restricted mean survival time following a one standard deviation increase in the explanatory variable in question.

Variable		Restricted Model Coefficient	HRC	STC	Full Model Coefficient	HRC	STC
Level covariates: $Level_{ij}$							
Bid-Ask Spread	BAS_{ij}	-0.113*** (-100.4)	-40.9%	14.3	-0.069*** (-44.51)	-27.3%	8.9
Back of Queue Quantity	$\overleftarrow{Q}_{ij}$	-1.323*** (-62.32)	-25.4%	8.3	-1.931*** (-71.51)	-34.8%	11.6
Front of Queue Quantity	$\vec{Q}_{ij}$	0.206*** (11.26)	4.6%	-1.3	0.011 (0.55)	0.2%	-0.1
Order Quantity		0.246*** (6.00)	2.2%	-0.7	0.613*** (14.03)	5.7%	-1.6
Lagged Absolute Return		0.022*** (7.08)	2.3%	-0.7	0.019*** (6.45)	1.9%	-0.5
Lagged Volume		0.008* (2.18)	0.8%	-0.2	0.003 (0.86)	0.3%	-0.1
Internal Match Rate		-0.737*** (-15.62)	-5.9%	1.8	-0.744*** (-15.43)	-6.0%	1.8
Time-varying covariates: $\Delta\Pi_{ijt}$							
Bid-Ask Spread	ΔBAS_{ijt}				-0.022*** (-9.25)	-8.2%	2.5
Opposite-Side Quantity	ΔQ_{ijt}^{Opp}				0.166*** (8.40)	10.1%	-2.9
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{ijt}$				-1.292*** (-39.01)	-49.4%	17.5
Front of Queue Quantity	$\Delta\vec{Q}_{ijt}$				0.023 (1.16)	1.9%	-0.6
Time-varying covariates: $Control_{ijt}$							
Firm Inventory					0.016** (3.40)	1.6%	-0.5
Same-Side MTF Trades					0.051*** (14.74)	5.2%	-1.5
Opposite-Side MTF Trades					-0.016*** (-3.51)	-1.6%	0.5

Table III: PDHM of Order Cancellations Conditional on the Binding Tick Size

This table presents the results for the same stratified PDHM as in Table II (full model), with different estimates depending on whether the tick size is binding or not. The sample orders are grouped by the indicator variable *Binding Tick Size,* which equals one if the tick size is binding at the time of sampling and zero otherwise. All variable definitions are the same as in Table II. The sample includes 89,559 orders for when the tick size is binding and 48,927 orders for when it is not, sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. Each sample order is tracked over up to 100 events. For each estimate, the *t*-values are reported in parentheses and ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered on order *i*. Here, HRC is the change in the cancellation hazard rate following a one standard deviation change in the explanatory variable in question, holding all other explanatory variables constant, and STC is the change in the restricted mean survival time following a one standard deviation increase in the explanatory variable in question.

Variable		Binding Tick Size: Coefficient	HRC	STC	Non-Binding Tick Size: Coefficient	HRC	STC
Level covariates: $Level_{ij}$							
Bid-Ask Spread	BAS_{ij}	-0.061*** (-10.14)	-21.6%	7.0	-0.066*** (-27.89)	-28.6%	8.4
Back of Queue Quantity	$\overleftarrow{Q}_{ij}$	-2.078*** (-65.09)	-38.7%	13.4	-1.646*** (-36.28)	-27.9%	8.2
Front of Queue Quantity	$\vec{Q}_{ij}$	-0.229*** (-9.1)	-5.1%	1.5	0.411*** (11.55)	8.4%	-2.2
Order Quantity		0.609*** (10.56)	5.5%	-1.6	0.480*** (6.85)	4.6%	-1.2
Lagged Absolute Return		0.018*** (3.62)	1.6%	-0.5	0.018*** (4.66)	2.1%	-0.6
Lagged Volume		0.001 (0.13)	0.1%	0.0	0.007 (1.14)	0.7%	-0.2
Internal Match Rate		-0.638*** (-10.31)	-5.1%	1.5	-0.848*** (-10.9)	-6.8%	1.9
Time-varying covariates: $\Delta\Pi_{ijt}$							
Bid-Ask Spread	ΔBAS_{ijt}	0.006 (0.83)	2.0%	-0.6	-0.058*** (-18.57)	-20.3%	5.8
Opposite-Side Quantity	ΔQ_{ijt}^{Opp}	0.185*** (12.3)	12.1%	-3.4	0.201*** (11.66)	10.2%	-2.7
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{ijt}$	-1.434*** (-34.45)	-52.8%	19.4	-1.153*** (-21.03)	-45.8%	14.3
Front of Queue Quantity	$\Delta\vec{Q}_{ijt}$	-0.175** (-3.01)	-10.6%	3.3	0.067*** (4.29)	7.1%	-1.9
Time-varying covariates: $Control_{ijt}$							
Firm Inventory		0.009 (1.56)	0.9%	-0.3	0.021* (2.28)	2.0%	-0.5
Same-Side MTF Trades		0.051*** (12.22)	5.5%	-1.6	0.030*** (4.60)	2.7%	-0.7
Opposite-Side MTF Trades		-0.044*** (-7.46)	-4.6%	1.4	0.024** (3.14)	2.1%	-0.6

Table IV: PDHM of Order Cancellations for HFTs and Non-HFTs, Conditional on Binding and Non-Binding Tick Sizes

This table presents the results for the same stratified PDHM as in Table II, with different estimates depending on whether the tick size is binding or not and whether the sample order is posted by a trading firm defined as an HFT or not. The sample orders are grouped by the indicator variables *Binding Tick Size* (which equals one if the tick size is binding at the time of sampling and zero otherwise) and *HFT* (which equals one if the order is submitted by a trading firm defined as an HFT and zero otherwise). HFTs are trading firms that are members of FIA EPTA. The results for a binding (non-binding) tick size are shown in Panel A (Panel B). All variable definitions are the same as in Table II. The sample includes 27,334 orders submitted by HFTs when the tick size is binding, 62,225 orders submitted by non-HFTs when the tick size is binding, 11,322 orders submitted by HFTs when the tick size is not binding, and 37,605 orders submitted by non-HFTs when the tick size is not binding. The orders are sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. Each sample order is tracked over up to 100 events. For each estimate, the t-values are reported in parentheses and ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered on order i. Here, HRC is the change in the cancellation hazard rate following a one standard deviation change in the explanatory variable in question, holding all other explanatory variables constant, and STC is the change in the restricted mean survival time following a one standard deviation increase in the explanatory variable in question.

Panel A: Binding Tick Size		HFTs			Non-HFTs		
Variable		Coefficient	HRC	STC	Coefficient	HRC	STC
Level covariates: $Level_{ij}$							
Bid-Ask Spread	BAS_{ij}	-0.102*** (-9.20)	-32.5%	10.7	-0.014 (-1.70)	-5.7%	1.7
Back of Queue Quantity	$\overleftarrow{Q}_{ij}$	-3.919*** (-56.44)	-61.0%	23.5	-1.480*** (-39.71)	-29.0%	9.5
Front of Queue Quantity	$\vec{Q}_{ij}$	-1.610*** (-23.80)	-28.3%	9.1	0.266*** (9.11)	6.4%	-1.8
Order Quantity		2.301*** (14.77)	16.2%	-4.3	0.198** (3.10)	1.9%	-0.6
Lagged Absolute Return		0.009 (0.85)	0.8%	-0.2	0.017** (3.06)	1.6%	-0.5
Lagged Volume		-0.015 (-1.53)	-1.5%	0.4	-0.014* (-2.09)	-1.3%	0.4
Internal Match Rate		-1.143*** (-11.20)	-9.6%	2.8	0.956*** (11.15)	7.2%	-2.1
Time-varying covariates: $\Delta\Pi_{ijt}$							
Bid-Ask Spread	ΔBAS_{ijt}	-0.029* (-2.23)	-9.5%	2.8	-0.024** (-2.63)	-8.1%	2.4
Opposite-Side Quantity	ΔQ_{ijt}^{Opp}	0.345*** (11.50)	22.5%	-5.8	0.145*** (11.16)	9.6%	-2.7
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{ijt}$	-4.277*** (-41.77)	-88.7%	40.9	-0.650*** (-18.87)	-29.1%	9.5
Front of Queue Quantity	$\Delta\vec{Q}_{ijt}$	-2.892*** (-11.88)	-79.4%	34.3	-0.022 (-0.57)	-1.4%	0.4
Time-varying covariates: $Control_{ijt}$							
Firm Inventory		0.141*** (6.66)	12.3%	-3.3	-0.008 (-1.23)	-0.9%	0.3
Same-Side MTF Trades		0.007 (0.93)	0.7%	-0.2	0.065*** (12.76)	7.0%	-2.0
Opposite-Side MTF Trades		-0.080*** (-6.01)	-9.2%	2.7	-0.025*** (-3.66)	-2.5%	0.7

Panel B: Non-Binding Tick Size		HFTs			Non-HFTs		
Variable		Coefficient	HRC	STC	Coefficient	HRC	STC
Level covariates: $Level_{ij}$							
Bid-Ask Spread	BAS_{ij}	-0.140*** (-24.00)	-47.2%	13.4	-0.045*** (-16.97)	-20.9%	6.2
Back of Queue Quantity	$\overleftarrow{Q}_{ij}$	-2.135*** (-20.68)	-33.8%	9.1	-1.501*** (-29.68)	-25.9%	7.8
Front of Queue Quantity	$\vec{Q}_{ij}$	0.058 (0.71)	1.1%	-0.3	0.470*** (11.78)	9.8%	-2.6
Order Quantity		0.779*** (6.27)	8.4%	-2.0	0.309*** (3.68)	2.8%	-0.8
Lagged Absolute Return		0.029 (1.95)	3.0%	-0.7	0.018*** (4.74)	2.1%	-0.6
Lagged Volume		0.064*** (5.90)	7.1%	-1.7	-0.018* (-2.54)	-1.9%	0.5
Internal Match Rate		-1.015*** (-5.03)	-7.8%	1.9	-0.394*** (-3.91)	-2.7%	0.8
Time-varying covariates: $\Delta\Pi_{ijt}$							
Bid-Ask Spread	ΔBAS_{ijt}	-0.064*** (-10.17)	-21.2%	5.5	-0.064*** (-17.85)	-22.3%	6.6
Opposite-Side Quantity	ΔQ_{ijt}^{Opp}	0.264*** (8.99)	13.7%	-3.2	0.191*** (10.02)	9.6%	-2.6
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{ijt}$	-2.006*** (-14.77)	-66.0%	20.2	-0.931*** (-16.42)	-38.9%	12.2
Front of Queue Quantity	$\Delta\vec{Q}_{ijt}$	0.073 (1.87)	9.8%	-2.3	0.073*** (3.66)	7.0%	-1.9
Time-varying covariates: $Control_{ijt}$							
Firm Inventory		0.114*** (4.43)	9.4%	-2.2	0.005 (0.46)	0.4%	-0.1
Same-Side MTF Trades		0.017 (1.36)	1.4%	-0.3	0.033*** (4.32)	3.1%	-0.8
Opposite-Side MTF Trades		-0.019 (-1.20)	-1.7%	0.4	0.039*** (4.91)	3.5%	-0.9

Table V: Order Price Revisions

This table reports the frequencies of different types of order price revisions. A canceled sample order is defined as part of an order price revision if the same trading firm submits a new limit or market order in the same stock on the same side of the market but at a different price within half a second of the cancellation. Panel A reports information on all canceled orders in the sample (73,442 orders). Panels B and C contain the corresponding information on fleeting orders (14,107) and HFT orders (24,324). Fleeting orders are defined as having a duration of less than 2 seconds. The orders are sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. The order price revisions are subcategorized as follows*:* if the new order leads to immediate execution, it is considered marketable. If the new limit order does not lead to immediate execution, the price revision is subcategorized as either more aggressive or less aggressive relative to the price of the canceled limit order. All frequencies are conditioned on an opposite-side price change in the half-second preceding the cancellation or during the order lifetime, whichever is shorter (see the three rightmost columns). For example, for a bid-side order, an opposite-side price change is a change in the best ask price. If the spread between the order price and the opposite price increases, decreases, or remains unchanged, we say that the opposite price moves away, moves closer, or stays the same, respectively.

Panel A: Full Sample		Opposite-side price change in the half-second preceding the cancellation		
	% of Total	Moves Away	Moves Closer	Stays the Same
Order Price Revisions	15.6	1.9	1.9	11.8
Marketable	0.5	0.0	0.1	0.3
More Aggressive	4.6	1.0	1.0	2.6
Less Aggressive	10.6	0.9	0.8	8.8
Other Cancellations	84.4	5.9	6.0	72.5
All Cancellations	100.0	7.8	7.9	84.4

Panel B: Fleeting Orders		Opposite-side price change in the half-second preceding the cancellation		
	% of Total	Moves Away	Moves Closer	Stays the Same
Order Price Revisions	13.4	0.9	0.9	11.6
Marketable	0.4	0.0	0.0	0.4
More Aggressive	1.6	0.3	0.4	0.9
Less Aggressive	11.4	0.6	0.6	10.3
Other Cancellations	86.6	3.6	3.4	79.6
All Cancellations	100.0	4.4	4.4	91.2

Panel C: HFT Orders		Opposite-side price change in the half-second preceding the cancellation		
	% of Total	Moves Away	Moves Closer	Stays the Same
Order Price Revisions	10.3	0.9	0.9	8.4
Marketable	0.2	0.0	0.0	0.1
More Aggressive	1.1	0.3	0.3	0.6
Less Aggressive	9.0	0.6	0.6	7.7
Other Cancellations	89.7	4.8	5.1	79.8
All Cancellations	100.0	5.7	6.0	88.3

Table A.I: Stock Characteristics

The table lists the properties of the OMXS 30 sample stocks. All stocks except NOKI SEK are traded primarily on Nasdaq Stockholm, whereas NOKI SEK, which is a Swedish depositary receipt issued by the Finnish firm Nokia Oyj, has its primary listing on Nasdaq Helsinki. We calculate *Market Capitalization* based on the closing price on April 30, 2014 (expressed in millions of Swedish krona, MSEK). The market capitalization for NOKI SEK represents only its depositary receipt size. We obtain all other statistics as averages across trading days during April 2014. The variable *Daily Trading Volume* is in millions of Swedish krona; *Daily Turnover* is the daily trading volume divided by 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 tick size; *Effective Spread* is the trade value-weighted average absolute difference between the trade price and the bid–ask midpoint; *Depth at BBO* is the trade volume in millions of Swedish krona required to change the stock price, averaged across seconds and the sides of the order book; *Volatility* is the average 10-second squared basis point return, calculated from bid–ask midpoints; and *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

Table A.II: Correlation Matrix

The table displays the sample correlation coefficients between the explanatory variables in the stratified PDHM from Table II. The variable *Bid-Ask Spread* is the difference between the limit order price and the best price on the opposite side of the book, divided by the midpoint, measured at the time of sampling and expressed in basis points. The following variables are measured at the time of sampling and expressed as fractions of the volumes posted at BBO: *Back of Queue Quantity* is the volume posted at the same price but with a lower time priority than the order of interest; *Front of Queue Quantity* is the volume posted at the same price but with a higher time priority than the order of interest, plus the order's own volume; *Opposite-Side Quantity* is the volume posted at the best price level on the opposite side of the order in question; *Order Quantity* is the volume of the order in question; *Lagged Absolute Return* is the absolute five-minute stock return prior to the sampling time; *Lagged Volume* is the number of shares traded during the five minutes prior to the sampling time, relative to the total depth; *Internal Match Rate* is 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; and *Firm Inventory* is the stock-specific net trading volume (buy volume minus sell volume, expressed in numbers of shares) of the trading firm that submitted the order in question. To account for the direction of trade, the net trading volume is multiplied by -1 if the order of interest is submitted on the ask side of the order book. The variables *Same-Side MTF Trades* and *Opposite-Side MTF Trades* measure the trading volume recorded at the three MTFs from the time of sampling until time *t*, on the same and opposite sides as the order in question, respectively. The sample includes 138,486 limit orders sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. Each sample order is tracked over up to 100 events, resulting in 6,141,702 observations.

	Levels							Time-Varying Covariates						
	Bid-Ask Spread	*Back of Queue Quantity*	*Front of Queue Quantity*	*Order Quantity*	*Lagged Absolute Return*	*Lagged Volume*	*Internal Match Rate*	*Bid-Ask Spread*	*Opposite-Side Quantity*	*Back of Queue Quantity*	*Front of Queue Quantity*	*Firm Inventory*	*Same-Side MTF Trades*	*Opposite-Side MTF Trades*
Levels														
Bid-Ask Spread	1.00	0.05	0.03	0.01	0.06	0.04	0.03	0.64	0.17	0.18	0.28	0.00	-0.01	0.08
Back of Queue Quantity		1.00	-0.41	-0.15	-0.02	-0.02	-0.01	0.06	0.14	-0.19	0.00	-0.01	-0.09	0.07
Front of Queue Quantity			1.00	0.29	0.01	0.03	0.05	0.14	0.19	0.08	0.02	0.01	-0.03	0.06
Order Quantity				1.00	0.04	0.08	-0.05	0.13	0.17	0.22	0.17	0.05	0.05	0.03
Lagged Absolute Return					1.00	0.20	0.00	0.03	0.00	0.02	0.03	-0.01	0.03	0.02
Lagged Volume						1.00	0.00	0.03	0.02	0.03	0.04	0.01	0.11	0.12
Internal Match Rate							1.00	0.02	0.01	0.01	0.01	0.02	-0.01	0.01
Time-Varying Covariates														
Bid-Ask Spread								1.00	0.36	0.35	0.34	0.02	0.03	0.15
Opposite-Side Quantity									1.00	0.18	0.09	-0.01	-0.02	0.07
Back of Queue Quantity										1.00	0.37	0.01	0.07	0.05
Front of Queue Quantity											1.00	0.01	0.02	0.08
Firm Inventory												1.00	0.02	0.00
Same-Side MTF Trades													1.00	0.01
Opposite-Side MTF Trades														1.00

Table A.III: PDHM of Order Cancellations for Fleeting Orders and Orders Sampled at the Time of Submission

This table presents estimates of the following stratified PDHM with time-varying covariates for a subsample of orders with a lifetime shorter than 2 seconds (fleeting orders) and a subsample of orders that are sampled at the time of submission (orders at submission):

$$\lambda_{ij}(t) = \lambda_{0j}(t)\exp\left[Level_{ij}\theta + \Delta\Pi_{ijt}\beta + Control_{ijt}\delta\right],$$

where $\lambda_{ij}(t)$ is the hazard rate of cancellation of order *i* over the next instant at time *t* in stock *j* and $\lambda_{0j}(t)$ is an unspecified baseline hazard rate for stock *j*. The vector $Level_{ij}$ includes the following variables at the time of sampling: *Bid-Ask Spread*, *Back of Queue Quantity*, *Front of Queue Quantity*, and *Opposite-Side Quantity* (defined as in Table I), *Order Quantity* (the volume of the sampled order, relative to the total depth), *Lagged Absolute Return* (the absolute five-minute stock return prior to the sampling time), *Lagged Volume* (the number of shares traded during the five minutes prior to the sampling time, relative to the total depth), and *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). The vector $\Delta\Pi_{ijt}$ includes changes in the model variables, relative the time of sampling. The vector $Control_{ijt}$ includes the following time-varying covariates: *Firm Inventory* is the stock-specific net trading volume (buy volume minus sell volume, expressed in numbers of shares) of the trading firm that submitted the order in question. To account for the direction of trade, the net trading volume is multiplied by -1 if the order of interest is submitted on the ask side of the order book. The variables *Same-Side MTF Trades* and *Opposite-Side MTF Trades* measure the trading volume recorded at the three MTFs from the time of sampling until time *s*, on the same and opposite sides as the order in question, respectively. The subsample with order lifetimes shorter than 2 seconds includes 14,108 limit orders and the subsample of orders that are sampled at the time of submission includes 6,138 limit orders. All limit orders are sampled from the OMXS 30 stocks in May 2014 on Nasdaq. For details on the sampling procedure, see Section II.C. Each subsample order is tracked over up to 100 events, resulting in 136,103 observations for the subsample of orders with lifetimes shorter than 2 seconds and 245,873 observations for the subsample of orders that are sampled at the time of submission. For each estimate, the *t*-values are reported in parentheses and ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively. Standard errors are clustered on order *i*. Here, HRC is the change in the cancellation hazard rate following a one standard deviation change in the explanatory variable in question, holding all other explanatory variables constant, and STC is the change in the restricted mean survival time following a one standard deviation increase in the explanatory variable in question.

Variable		Fleeting Orders			Orders at Submission		
		Coefficient	HRC	STC	Coefficient	HRC	STC
Level covariates: $Level_{ij}$							
Bid-Ask Spread	BAS_{ij}	-0.022*** (-5.73)	-8.4%	0.7	-0.017** (-2.92)	-7.9%	2.6
Back of Queue Quantity	$\overleftarrow{Q}_{ij}$	-1.311*** (-18.18)	-18.3%	1.8	-0.995*** (-3.83)	-6.1%	2.0
Front of Queue Quantity	$\vec{Q}_{ij}$	0.153*** (3.29)	3.7%	-0.3	0.023 (0.40)	0.6%	-0.2
Order Quantity		-0.053 (-0.58)	-0.6%	0.0	0.164 (1.46)	2.3%	-0.7
Lagged Absolute Return		-0.001 (-0.16)	-0.2%	0.0	0.015 (1.21)	1.2%	-0.4
Lagged Volume		-0.059*** (-7.53)	-7.8%	0.7	-0.017 (-1.38)	-1.8%	0.6
Internal Match Rate		-0.058 (-0.56)	-0.5%	0.0	-1.000*** (-6.03)	-8.0%	2.6
Time-varying covariates: $\Delta\Pi_{ijt}$							
Bid-Ask Spread	ΔBAS_{ijt}	-0.060*** (-8.89)	-14.0%	1.3	-0.086*** (-10.97)	-28.3%	10.6
Opposite-Side Quantity	ΔQ_{ijt}^{Opp}	0.249*** (10.05)	14.7%	-1.0	0.144*** (6.87)	13.0%	-3.7
Back of Queue Quantity	$\Delta\overleftarrow{Q}_{ijt}$	-1.240*** (-12.83)	-43.3%	6.2	-0.464*** (-7.81)	-37.1%	14.7
Front of Queue Quantity	$\Delta\vec{Q}_{ijt}$	-0.340 (1.52)	-26.6%	2.9	-0.050 (1.00)	-7.4%	2.4
Time-varying covariates: $Control_{ijt}$							
Firm Inventory		-0.024 (-1.75)	-2.2%	0.2	-0.025 (-1.21)	-2.5%	0.8
Same-Side MTF Trades		-0.033*** (-5.20)	-3.4%	0.3	-0.010 (-1.05)	-0.8%	0.3
Opposite-Side MTF Trades		-0.005 (-0.41)	-0.7%	0.1	0.018 (0.94)	2.4%	-0.7