Order Exposure in High Frequency Markets*

BIDISHA CHAKRABARTY, TERRENCE HENDERSHOTT, SAMARPAN NAWN, and ROBERTO PASCUAL†

June, 2020

† bidisha.chakrabarty@slu.edu, hender@haas.berkeley.edu, samarpan.nawn@iimu.ac.in, rpascual@uib.es

We thank David Abad, Michael Brolley, Ashok Banerjee, Carole Comerton-Forde, Darrell Duffie, Bart Frijns, Corey Garriott, Sudip Gupta, Frank Hatheway, Vitto Mollica, Marta Khomyn, Venky Panchapagesan, Satchit Sagade, Ajay Shah, Andriy Shkilko, Kumar Venkataraman, Irena Vodenska, Gunther Wuyts, Pradeep Yadav and seminar participants at the Auckland Finance Meeting (Queenstown), Auckland University of Technology, Australasian Finance and Banking Conference (Sydney), Emerging Markets Finance Conference (IGIDR Mumbai), Finance Forum (Madrid), Financial Management Association European Conference (Kristiansand), Financial Market Liquidity Conference (Budapest), Northern Finance Association Meeting (Charlevoix), NSE-NYU Conference, Saint Louis University, Swiss Society for Financial Markets Research (Zurich), Universidad CEU Cardenal Herrera, Wilfrid Laurier University, University of Paris-Dauphine, Stockholm Business School, and European Financial Management Association Meeting (Ponta Delgada), for useful comments.

* This working paper is part of the NSE-NYU Stern School of Business Initiative for the Study of the Indian Capital Markets. The authors acknowledge the support of the initiative. The views expressed in this Working Paper are those of the authors and do not necessarily represent those of the National Stock Exchange of India (NSE) or New York University (NYU). Pascual acknowledges the financial support of the Spanish DGICYT project ECO2017-86903-P, and the Generalitat Valenciana Grant Prometeo/2017/158. Chakrabarty acknowledges the Edward Jones endowment. Hendershott gratefully acknowledges support from the Norwegian Finance Initiative. Hendershott provides expert witness services to a variety of clients.

Order Exposure in High Frequency Markets

ABSTRACT

We examine how traders' technology and motivation to trade impact transparency. We show that algorithmic traders (ATs) are more likely to hide orders and their hidden orders receive better execution. High frequency traders (HFTs, a subset of ATs) extensively use small hidden orders that are aggressively priced near the best quotes. Theory suggests that traders hide orders to limit their option value, delay information exposure, and limit competition for liquidity provision. Our results suggest that HFTs hide orders to limit competition when the expected profitability of liquidity provision is higher, while non-algorithmic traders hide orders to limit their option value.

Keywords: Hidden orders, high frequency trading, order exposure

JEL Classification: G11; G12; G14; G15, G24

Financial markets are not fully transparent. Traders strategically disclose their trading interest by partially or fully hiding their desired quantity. In equity markets this opacity has increased: the Securities and Exchange Commission (SEC) finds that hidden volume's contribution to trades increased from 15% to over 30% in the US between 2012 and 2018.[1] For limit orders near the best quotes on Nasdaq, 30% (50%) are hidden on low (high) volatility days (see Figure 1).

[Figure 1]

Another significant recent trend is the rise of algorithmic and high frequency traders (ATs and HFTs). ATs encompass both algorithms used to execute customers' orders (agency ATs; AATs) and algorithms used for low-latency speculation, market making, cross-market arbitrage, and other short-term speculative trading (HFTs). Despite the increase in both hidden orders and ATs, relatively little is known about how ATs use hidden orders. Empirically studying hidden order usage by different traders is challenging because publicly available trade and quote data do not identify hidden orders prior to execution, nor identify who placed the order. We use proprietary data from two markets – the National Stock Exchange of India (NSE) and the Nasdaq – that both identify hidden orders and provide information on the order submitter.[2] Using these data we study how the technology and motivations to trade affect the economics of order exposure.

Hidden orders have lower priority than exposed orders at the same price.[3] This lowers their probability of execution and lengthens their expected time to execution (Bessembinder, Panayides, and Venkataraman, 2009, hereafter, BPV). Traders weigh these costs of hiding against the benefits of concealing their trading intentions. Order exposure is costly if other traders trade ahead of the order or withdraw liquidity on the opposite side of the market. Other traders may take advantage

[1] https://www.sec.gov/marketstructure/datavis/ma_exchange_hiddenvolume.html; inclusion of the volume executed in dark pools makes the temporal rise in opacity even larger.

[2] The NSE database contains details on each individual order that is not present in the Nasdaq database. The Nasdaq exchange made available a database identifying HFTs to some researchers. Several studies (for example, Carrion, 2013; Brogaard, Hendershott, and Riordan, 2014) use it. While the Nasdaq database contains all trades, is does not include complete order-level information. Instead it contains summary one-minute snapshots of the order book for a subset of dates in each quarter. Furthermore, it only distinguishes between HFTs and non-HFTs. Where possible, we show that our findings from the NSE data also hold in this Nasdaq HFT data set.

[3] Opacity in financial markets come from orders that are fully or partly hidden (iceberg) in lit exchanges, as well as dark pools that are completely opaque. In lit markets (the focus of this study), fully hidden orders are more prevalent in North American markets (e.g., US, Canada) while iceberg orders are widespread in Europe and the Asia-Pacific (e.g., Spain, France, India, Australia). Hidden orders are subject to the exposure priority rule, whereby hidden volume loses time priority to displayed volume at the same price

of exposed orders if they believe that the limit order submitter has access to private information or if they believe individual limit orders are a small part of a larger order that will continue to execute over time. Order exposure also increases the risk of being adversely selected, since an exposed order provides the counterparty with a free option to trade (Copeland and Galai, 1983). Order exposure is beneficial if it attracts traders to execute more quickly against the visible order.

Traders with different technology and motivations weigh the costs and benefits of order exposure differently.[4] Using data before the growth of AT, BPV find that traders with relatively larger order sizes and less information are more likely to use hidden orders. We find this continues to be true for AATs and non-algorithmic traders (NATs). However, AATs' use of technology to monitor the market enables their hidden orders to execute more often than NATs' orders. This reduces the cost of hiding orders and may explain why AATs are more likely to hide their orders than NATs.

HFTs, on the other hand, use technology similar to AATs but differ in their motivation to trade. Unlike AATs, HFTs generally do not trade stocks based on fundamental valuation. Furthermore, they do not acquire positions or liquidate orders over a series of smaller trades. HFTs follow a variety of strategies including market making, statistical arbitrage, and other short-term speculative trading (Hagströmer and Nordén, 2013; Boehmer, Li and Saar, 2018). Because the decision to hide is only relevant for non-marketable limit orders, HFTs decision to hide should primarily relate to their liquidity-provision/market-making strategies.

Why would an HFT market maker choose to hide her trading interest? Market makers' profits depend on trading volume, so hidden orders' lower execution priority is costly. One reason to accept a lower probability of execution is to preserve some informational advantage. Theoretical models suggest that informed traders may hide orders to obscure their trading intentions (Moinas, 2010), thereby reducing the expropriation of informational rents (Boulatov and George, 2013). Studies show that HFTs' trades (Brogaard, Hendershott, and Riordan, 2014) and orders (Brogaard, Hendershott, and Riordan, 2019) carry information. HFTs' limit orders could be motivated by small informational advantages arising from signals gleaned from order flow (Hirschey, 2018;

[4] See Hendershott and Riordan (2013) for an analysis of AT and non-AT trades and displayed limit orders.

Korajczyk and Murphy, 2019) or public news (Bhattacharya, Chakrabarty, and Wang, 2020). However, these informational advantages are likely to be short-lived, making hidden orders' unattractive due to their lower execution priority. Another possible motive for HFTs hiding orders could be to reduce competition in liquidity provision (Buti and Rindi, 2013, Foley and Putnins, 2016). If other traders competing to provide liquidity are aware of the size of the HFTs' order, they may be more likely to offer a better price.

While HFTs may hide their orders to both conceal information and reduce competition, these explanations have different empirical predictions. If HFTs hide orders to conceal information, their hidden orders should contain more information and impact the efficient price more than their displayed orders. If, on the other hand, HFTs hide their orders to reduce competition for liquidity provision, then their hidden orders should be placed when expected profits to liquidity provision are high, and will more likely be used to undercut recently placed orders.

We test the above predictions on hidden order use by different types of traders using data from two sources. Our primary data come from the NSE, the fifth largest market in the world by number of trades.[5] The NSE data furnish rich details on trader accounts, using which we classify each order as coming from one of three mutually exclusive trader types: proprietary algorithmic traders (HFTs), other ("agency") algorithmic traders (AATs), and non-algorithmic traders (NATs). The NSE allows iceberg orders and we can identify both the displayed and the hidden portions of each order.[6] Our second data source is the Nasdaq exchange in the U.S. which allows fully hidden orders. The Nasdaq data provide one-minute snapshots of the order book that identify all standing limit orders, whether they are hidden or displayed, and whether they were placed by HFTs or non-HFTs. We use the term hidden limit orders (HLOs) for both fully hidden and iceberg orders, noting that the NSE HLOs are iceberg orders while the Nasdaq HLOs are fully hidden.

We group the above discussions for hidden order use into three broad questions: "whether", "how", and "why" different traders use hidden orders. For whether hidden order usage varies

[5] In our sample period, HFTs contribute 33% of the total daily volume on the NSE. See https://www.nseindia.com/research/content/1314_BS6.pdf

[6] The NSE allows traders to hide a portion of their total order size by choosing from the Display Size menu and displaying only a fraction of the total order size. The minimum display size is 10% of the original volume. This is similar to the iceberg orders allowed on Euronext Paris and studied in BPV. We provide more details on order types in Section I. A.

across traders with different technologies (ATs versus NATs) and motivation (HFTs versus AATs) for trading, we find that every trader group uses HLOs, although the usage intensity varies significantly. In the NSE sample 8.92% of all limit orders submitted by HFTs are HLOs. Corresponding numbers for AATs (NATs) are 26.28% (9.14%). In the Nasdaq market HFTs hide 21.8% of all limit orders while non-HFTs hide 15.39%. Thus, in both markets, ATs (AAT and HFT) as a group make more use of hidden orders than non-algorithmic ones. While the hidden order use of NATs is consistent with the findings of BVP, it is somewhat surprising that ATs (both AATs and HFTs) pay the cost in terms of loss of priority because they have the technology to monitor markets in real time and can cancel displayed orders faster, should the need arise.

The contrast between ATs and NATs persists when we analyze order placement in different layers of the order book. In the NSE, HFTs place 52.21% of their hidden orders within the best quotes and another 19.39% at the best quotes. In fact, 97.74% of HFT's HLOs in large stocks are placed within five ticks from the best quotes. In contrast, NATs place 38.47% of their HLOs away from the five best ticks. AATs' hidden order placement falls between the HFTs' and NATs'. Similarly, in the Nasdaq order book HFTs place 30.26% of their limit orders at the best quotes and only 10.61% far from the best quotes, while for the non-HFTs the pattern is the opposite with 22.45% at the best quotes and 30.34% far from the best quotes. Thus, in both markets ATs place more aggressive hidden orders than non-ATs.

BPV find that traders hide larger orders. In examining the share sizes of hidden orders, we find that while the BPV result holds for NATs and AATs in our data, HFTs' hidden orders are much smaller in size. In the NSE sample, 68.3% of HFTs' HLOs are below 50 shares in size while for NATs this number is 28.12%. These patterns are also present in the Nasdaq data, where 82.83% of HFT orders in the 0 to 100 share size category is hidden while the corresponding number for non-HFTs is 42.01%

Descriptive statistics from both markets indicate that HFTs' hidden order usage contrasts sharply with that of the NATs, with AATs' falling somewhere in between. But how efficiently does each group use these orders? We model this part of the investigation on BPV who find that HLOs have a lower probability of completion and take longer to execute compared to similar displayed limit orders (DLOs), although DLOs incur higher implementation cost. We find that

HFTs' HLOs have the *highest* likelihood of execution, both for buy and sell orders. In fact, HLO execution probabilities for HFTs is over twice that of the NATs. We also model the time to completion (i.e., full execution) of HLOs vis-à-vis other orders using survival analysis, as in Lo, MacKinlay, and Zhang (2002). This test is particularly relevant in our context, since iceberg orders may mechanically induce a protracted time to completion. We find that HLOs placed by HFTs execute faster compared to those placed by other traders.

The benefits of concealing orders must be weighed against the costs they incur. To estimate costs, we use the implementation shortfall metric (IS; Perold, 1988). This metric has two components: effective cost (price impact), and opportunity cost of non-execution (which measures forgone profits). We find that HFTs face higher effective cost for hidden orders, which is expected since HFTs use more aggressive HLOs (the majority placed within or at the best quotes). However, their opportunity cost of non-execution is significantly lower, indicating less adverse price movements after their hidden order submissions. When combined, the lower opportunity costs either compensate for, or exceed, the higher execution costs and, overall, HFTs' HLOs have much lower implementation shortfall. These findings suggest that HFTs use HLOs more efficiently than non-HFTs, with the IS of AATs falling in between these two trader groups.

Our final set of tests address the why question. What motivates HFTs' use of hidden orders? The share-size results presented earlier indicate that NATs' hidden orders are larger and the majority of them are placed away from the best quotes, similar to the results in BPV. But HFTs' hidden order use is different. Thus, we next test the two competing motives for hidden orders discussed earlier – limiting information revelation and competing for liquidity provision.

To test the information-revelation motive we examine whether HFTs' HLOs carry more information and impact the efficient price more than their own DLOs, using impulse response functions estimated from an extended structural VAR model in the spirit of Hasbrouck (1991a). We find that HFTs' HLOs have the lowest informational impact of trades when compared to their own DLOs, as well as when compared to the other two trader types' orders. As a more direct test, we also measure the information share (Hasbrouck, 1995) of each trader type–order type combination and find that HFTs' HLOs have the lowest information share of all trader types.

Overall, our findings do not support the hypothesis that HFTs use HLOs to trade on time sensitive information.

To test the liquidity provision motive (Foley and Putnins, 2016), we conduct several tests. First we model HFTs' decision to hide an order using logistic regressions. These tests confirm that HFTs are more likely to place HLOs when expected realized spreads (anticipated profits to liquidity provision), conditional on order aggressiveness, are higher. Additionally, we find that the size of HFTs' hidden orders is larger when expected realized spreads are high. We follow up with direct tests of whether HFTs use aggressive hidden orders to undercut standing orders to compete for liquidity provision. We define an undercutting order as a limit order that (i) is placed immediately after another submission on the same side, (ii) is placed within 10 milliseconds of the previous order, and (iii) improves upon the previous price. We find that HFTs are more likely to use HLOs than DLOs to undercut existing orders at the best quotes. We also find that actual realized spreads for HLO executions are higher for HFTs than for other traders. However, actual effective and realized spreads for HFTs' DLOs are higher, consistent with the fact that HFTs place their hidden orders more aggressively than their DLOs, which narrows the spread and thereby reduce the profits to liquidity supply.

This study sits at the interaction of two important issues facing investors and regulators – market opacity and high-speed trading. Research shows that when markets allow traders the facility to hide orders, they substitute non-displayed for displayed orders and change their trading aggressiveness (Bloomfield, O'Hara, and Saar, 2015). Meanwhile, improved (pre-trade) transparency can increase liquidity and the informational efficiency of prices (Boehmer, Saar, and Yu, 2005). Because transparency is a cornerstone of the SEC's investor protection function, current trends have regulators worried that opacity may be attractive to "bad-actors."[7] The growth of high speed trading, and in particular the increasing prevalence of HFTs, has also been accompanied by a frenzy of media commentaries on its inherent unfairness. Although studies find that HFT has both positive (Brogaard, Hendershott, and Riordan, 2014) and negative (Budish,

[7] See, for example, Chairman Clayton's speech: https://www.sec.gov/news/speech/speech-clayton-2017-11-08.

Cramton, and Shim, 2015) effects, there has been no evidence to date linking the motivations versus technology of trading to market opacity.

To our knowledge this is the first study to document that ATs in general, and HFTs in particular, make extensive use of HLOs in lit markets. The hidden orders of HFTs are different in characteristics (e.g., size, aggressiveness), information content, and usage (e.g., liquidity supply, undercutting) than the HLOs of non-HFTs. Our results hold for both consolidated (NSE) and fragmented (US) markets, and for iceberg as well as fully hidden orders, limiting the role of market structure and the difference between partial versus full non-exposure as explanatory factors.

We structure the rest of the paper as follows. Section I describes the institutional details of the NSE market, identification of trader account types, hidden orders on Nasdaq and a description of the two samples. Section II addresses the "whether" question and shows hidden order use by different types of traders. Section III examines the "how" question and documents the efficiency with which each trader type uses hidden versus displayed orders. Section IV tests the "why" question to confirm whether hidden orders are placed for information and/or liquidity provision reasons. Section V catalogs robustness tests and Section VI concludes. Additional information and robustness tables are presented in the accompanying Internet Appendix.[8]

I. Data

A. NSE sample and trader identification

With over 80% of the total traded volume, the NSE is the dominant market for its 1300+ listed stocks. It is a completely order driven market. Like the Euronext platform, the NSE allows traders to place hidden orders by choosing the "iceberg" option with a mandatory minimum exposure of 10% (of the original volume). Once the first tranche is executed, the next tranche is automatically displayed. All tranches are of the same size (10% or greater of the original order). The market operates on price-exposure-time priority whereby non-displayed volume loses time priority to any displayed volume at the same price. Thus, the iceberg order provision of the NSE is identical to that used on the Euronext and analyzed in BPV. Since there are no dark pools in this

[8] https://doc-share-c.s3.amazonaws.com/OrderExposure_InternetAppendix.pdf

market, traders who want to hide orders have to use iceberg orders in the lit market. For more details on the institutional features of the NSE, we refer readers to Kahraman and Tookes (2017).

We obtain order and trade data directly from the NSE. For each trading day we access a message file and a trade file. The message files contains every message for each stock including the ticker symbol, price, quantity and timestamp in jiffies (one jiffy is 33.3564 picoseconds or $(1/2^{16})^{th}$ of a second). Similar to the ITCH data of the Nasdaq platform, for every order the message file includes order entry, modification, execution and cancellation events. In addition, it provides several additional flags. For this study we use three flags: *Client*, *Order Entry Mode*, and *Modifier condition*. The trade file contains analogous information for each trade. By allowing temporal tracking of each order and matching orders to trades, these data allow us to build the complete limit order book (LOB) at any instant of time.

Client classifies trader accounts into *Custodian*, *Proprietary* and *Others*. *Custodian* represents traders who are members of the exchange but do not conduct their own clearing or settlement. This group comprises primarily of foreign institutional investors, mutual funds, and financial institutions. The *Proprietary* flag applies to members of the exchange who trade for their own proprietary accounts, and *Other* applies to all other customers of the exchange who employ their own clearing member.

Order Entry Mode flag applies to each *Client* flag and shows one of the two possible order entry systems used to interact with the NSE's limit order market: *Algorithmic* if order entry and management is done using an algorithm and *Non-Algorithmic* if a client uses a manual system. The product of the three *Client* flags and two *Order Entry Modes* enables us to identify six distinct trade originations. Our particular focus in this study is on the *Proprietary* client using *Algorithmic* order entry mode to trade on their own account. That is the definition of HFTs (SEC, 2010) and we can cleanly identify these orders and trades in our data.[9] We group all other orders and trades with the *Algorithmic* flag into the agency algorithmic trader (AAT) type and all orders and trades flagged with *Non- Algorithmic* order entry mode as non-algorithmic traders (NATs).

[9] See https://www.sec.gov/rules/concept/2010/34-61358.pdf

A key advantage of our identification is that, unlike previous studies, we classify HFT at a finer (message level) granularity. For example, when a trader conducts proprietary trades using algorithms, we classify those trades as HFT, but if this same trader conducts client trades using algorithms, we do not count those as HFT. This overcomes some known limitations of popular HFT identifying databases that group all HFTs as pure-play (e.g., the Nasdaq HFT database used in Brogaard, Hendershott, and Riordan, 2014) or allow for mixed categories that cannot be exactly classified as HFT (e.g., the EUROFIDAI data used in Boussetta, Lescourret, and Moinas, 2017).

Finally, the *Modifier* flag identifies iceberg orders and shows the minimum display volume, allowing us to see both the lit and dark proportions of each iceberg order.

B. Nasdaq sample

Where possible, we also use Nasdaq data with hidden order and trader type identifiers to corroborate the results from the NSE data. The Nasdaq data of this study has been used in other recent papers, for example Brogaard, Hendershott, and Riordan (2014) and Carrion (2013). Unlike these studies that use the trade files, we use the LOB files provided in these data.[10] Nasdaq allows traders to fully hide an order. In our data, all available liquidity on the Nasdaq book is shown in one-minute snapshots from 9:30 a.m. to 4:00 p.m. (inclusive). We have data for the first full week of the first month of each quarter during our sample period (2008 and 2009) the crisis week of Sept 15 – 19, 2008, and the week of Feb 22 – 26, 2010. Data records have a buy/sell indicator to denote the trade initiating side, price, a flag to indicate HFT and non-HFT, a flag to denote if liquidity is displayed or hidden, the ticker symbol, aggregate available shares at that price level/HFT status/display status combination, book snapshot time and date. For each snapshot, order by order records representing the ten best price levels (displayed and hidden) on each side of the market are also shown.

C. Sampling and descriptive statistics from each dataset

HFTs have a greater propensity to trade large stocks (Brogaard, Hendershott, and Riordan, 2014). To provide initial descriptive statistics that ensure adequate representation of both HFTs

[10] Yao and Ye (2018) also use the HFT limit order data.

and other trader types, from the NSE data we select a (market cap) stratified sample of 100 stocks as follows. We begin with the 1254 listed stocks in the NSE in September 2013, filter out 286 stocks that are not in continuous trading session during October to December 2013 (61 trading days). We also exclude firms that (i) have a closing price of Rs. 1 or lower, (ii) have fewer than 100 trades per day on average, (iii) trade less than 1000 shares a day, (iv) have a traded value per day of less than Rs. 100000 over the sample period, (v) have market-cap values in the Bloomberg and CMIE Prowess databases that diverge by over 10%, (vi) are involved in NSE or MSCI index changes. These filters reduce our universe of stocks to 695, which we sort by market capitalization and group into deciles. From each decile we select 10 stocks to generate the sample of 100, with 30 large-cap stocks, 40 mid-cap stocks and 30 small-cap stocks. Company information is from CMIE Prowess (analogous to Compustat), a database of Indian firms which covers approximately 80% of the NSE stocks (Kahraman and Tookes, 2017). Panel A of Table I shows the descriptive statistics of this sample.

[Table I]

The average firm in our sample has over 448 billion rupees market capitalization (about 6.3 billion USD per the exchange rate on 02/01/2020). Large-cap firms have a market capitalization of about 1465 billion rupees (20 billion USD), which is smaller than the large cap firms in the Nasdaq HFT dataset (Brogaard, Hendershott, and Riordan, 2014). Volume and number of trades are higher, and relative spread (ratio of the quoted spread to the quote midpoint) is much smaller for the large firms than mid-sized and the small firms, as expected. While both the accumulated displayed and hidden depths in the LOB are higher for large firms than mid- and small-sized firms, the differences are larger for displayed than for hidden depth.

To benchmark our direct identification of HFTs against much of the literature that uses proxies for HFT activity, in Panel B of Table I we report message traffic and cancellation statistics by trader types and across the three market cap groups. Comparing across each row, we see that HFTs account for much greater message traffic (defined as the sum of submissions, cancellations, and revisions) than either the AATs or the NATs in the large cap stocks, but not in the mid-sized or the small stocks. However, when we scale message traffic by the number of trades executed, HFTs show a bigger presence even in the mid- and small-cap firms. This preponderance of HFTs to generate large message traffic volume echoes similar findings from the US equity markets (Hendershott, Jones, and Menkveld, 2011).

The Nasdaq dataset consists of trades and quotes for a sample of 120 stocks, stratified by market capitalization and evenly split between Nasdaq and NYSE listing. Detailed description of this sample selection is in Carrion (2013).

[Table II]

Table II shows some descriptive statistics of this sample, including HFTs' presence at the top of the order book and their hidden order usage. For the full sample, over 71% of the time there is hidden volume at the best quotes; this number goes up to almost 80% for large stocks and is about 67% for small cap stocks. HFTs are at the best quotes 67.41% of time in the full sample and 93.23% (46.11%) in the large (small) cap stocks.

II. The "whether" question: Hidden order use by trader types

A. Relative importance of hidden orders

BPV find that (non-algorithmic) traders use hidden orders. Do ATs and in particular HFTs also hide orders? In this section we address this question. We present all results for the large cap firms, and report analogous measures for mid and small cap companies in the Internet Appendix. We begin by providing an in-depth look at the relative importance of hidden order use by the different trader types – HFTs, AATs and NATs on the NSE, and HFTs versus non-HFTs on the Nasdaq. To do so, we examine the placement of displayed versus hidden orders in the LOB and compute the accumulated displayed and non-displayed depth, both in the number of orders and in share volume. Table III reports the results.

[Table III]

In Panel A we show results from the NSE sample for the proportion of HLOs relative to all limit orders submitted, both for the number of orders (*Ord*.) and the volume (*Vol*.) of shares. Comparing across the first row, 8.92% (8.62%) of all orders (volume) submitted by HFTs are HLOs. AATs use hidden orders more; they place 26.28% (34.26%) of all orders (share volume) as HLOs. NATs' relative use of hidden orders is comparable to that of the HFTs in terms of orders, but not in terms of volume. In Panel B we show comparable results for the Nasdaq sample. Here we find that the relative use of hidden to displayed orders is significantly higher for HFTs (21.80%) than for non-HFTs (15.39%). In terms of volume, our findings suggest that non-HFTs use larger HLOs than HFTs, both, for NSE and Nasdaq.

B. Disaggregated look at the layers of the order book

Position in the limit order queue is valuable for all traders (Hoffmann, 2014), but it may be especially important for HFTs (Yao and Ye, 2018). Therefore, we examine where in the LOB HFTs place their hidden and displayed orders. We build the NSE order book at every order submission time and identify the position of order placement at four levels relative to the top of the book – price improving or better than the standing best bid and ask quotes ("*Better*"), the best bid and ask ("*At*"), up to the first five ticks from the best bid and ask ("*Near*") and the rest of the book ("*Far*"). For the Nasdaq order book, since the database provides one-minute snapshots of the ten best levels of the LOB, we group the LOB levels into three segments – at the best quotes ("*At*"), from the best quotes up to five ticks away ("*Near*"), and the rest ("*Far*"). Table IV presents hidden and displayed order placement by the different trader types at these LOB depths for both markets.

[Table IV]

Panel A shows results for the NSE sample. Comparing across the first row, we find that while 52.21% (48.98%) of HFTs hidden orders (volume) are placed at "*Better*" than the best quotes, only 16.39% (10.43%) of NATs' hidden orders are at this highest level of the LOB. In fact, while HFTs place very little (2.25% of orders and 3.14% of volume) away from the top of the book, the opposite is true for NATs with 38.47% of orders and 36.09% of volume at the "*Far*" level. AATs' hidden order placement generally falls in between the HFTs' and NATs.' In Panel B, for the Nasdaq sample, we find similar patterns: HFTs place hidden orders more aggressively at or near the best quotes compared to non-HFTs. While 10.61% (14.64%) of HFTs' hidden orders (volume) are placed away from the five best ticks, non-HFTs place 30.34% (40.13%) of their hidden orders (volume) away from the five best ticks. In both panels, non-parametric tests show that the HLO placement at various levels of the LOB is significantly different across trader types.

C. Order size of hidden versus displayed orders

Theory posits that the decision to hide orders may be contingent on order attributes (Buti and Rindi, 2013). For example, when traders wish to trade large positions, they may hide their trading interest in order to limit the option value of their limit orders. Empirically BPV find this to be true. Thus, we next examine the sizes of hidden vis-à-vis displayed orders placed by the different trader types. We define size categories in total shares for both displayed and hidden orders to compare the order size distributions of HLOs and DLOs submitted by the different traders. Table V shows the hidden and displayed order sizes placed by HFTs, AATs and NATs for the NSE (Panel A) and HFTs versus non-HFTs for the Nasdaq (Panel B).

[Table V]

In Panel A we find that while HFTs' hidden orders are predominantly (68.4%) in the smallest share size category of 0 - 50 shares, the pattern reverses for NATs with only 28.12% hidden in this size category. Similarly, for displayed orders also these traders show opposite usage patterns: while HFTs place only 4.98% of their displayed orders in the smallest size category, NATs display 64.71% of their limit orders in this share size category. Comparing numbers along columns, we find that HFTs use larger share sizes when they expose their trading interest. For example, the largest percentages of displayed orders used by HFTs are in the 100 – 1000 share categories. The Nasdaq sample (Panel B) supports the same conclusion. HFTs' hidden orders (82.83%) in the 0 - 50 share size category is almost twice the comparable number (42.01%) for non-HFTs. Two-sample Kolmogorov-Smirnov (Massey, 1951) tests support the conclusion that the distributions of displayed versus hidden orders differ significantly across order size categories between HFTs versus other traders in both the NSE and Nasdaq.

Thus, in the use of order sizes as well, we find a contrast between the HFTs and the NATs. While the NATs' order size choice for hiding their trading interest is consistent with previous literature (BPV for example), HFTs behave in quite the opposite way. In fact, these small-sized hidden orders placed by HFTs bear out O'Hara's (2015) prescient summing up of the relationship between HFTs, small trades, and the ability to conceal trading interest. As she puts it "small trade sizes reflect the influence of HFTs because [these] silicon traders can spot (and exploit) human traders by their tendency to trade in round numbers, … all trading is converging to ever smaller sizes and is being hidden whenever possible."

A visual synthesis of the results in Tables IV and V clearly illustrate the contrast in order exposure between the different trader types.

[Figures 2 and 3]

Figure 2 plots the estimated cross-sectional daily average probabilities of hidden order submission by HFTs, AATs, and NATs, conditional on the order size and aggressiveness. It is clear that while HFTs (Fig. 2a.) have a higher probability of placing small sized hidden orders at all distances from the best quotes, they have the highest likelihood of placing such orders at the best quotes, followed by near the best quotes. Their use of hidden orders of larger size is significantly less. The pattern is the reverse for both AATs (Fig. 2b.) and NATs (Fig. 2c.), who use larger hidden order sizes placed further away from the best quotes. Figure 3 shows that the same behavior is evident in the Nasdaq, where HFTs place more hidden orders in the smaller trade size categories while the non-HFTs place more hidden volume using larger trade sizes.

III. The "how" question: Determinants and efficiency of order exposure

Given the preceding evidence of extensive hidden order use by ATs in general and HFTs in particular, and the contrast in how they use these orders, we next examine whether there are differences in how efficiently each trader type uses hidden orders.[11]

A. Likelihood of order execution

HLOs lose time priority and, as limit orders, they have no guarantee of execution. Thus, it is of interest to ask how the different trader types manage the *a priori* higher non-execution risk of hidden orders. To estimate this likelihood by specifically contrasting HFTs with the other trader groups, we use the following pooled Ordered Logit model, in the spirit of Ranaldo (2004) and Pascual and Veredas (2010):

[11] Results presented here and in all following sections require continuous LOB variables for model estimation. Therefore, these results are based only on the NSE sample since the Nasdaq data presents only snapshots of the LOB.

$$
\begin{aligned}
EXEC_{ij} = {} & \alpha_0 + \alpha_1 HFT_{ij} + \alpha_2 AAT_{ij} + \alpha_3 HLO_{ij} + \alpha_4 HLOHFT_{ij} + \alpha_5 HLOAAT_{ij} + \alpha_6 Aggr_{ij} + \\
& + \alpha_7 OrdSize_{ij} + \alpha_8 Rsprd_{ij} + \alpha_9 DepthSame_{ij} + \alpha_{10} DepthOpp_{it} + \\
& + \alpha_{11} LOBImb_{ij} + \alpha_{12} FirstHalfHour_{ij} + \alpha_{13} LastHalfHour_{ij} + \alpha_{14} OI_{it} + \\
& + \alpha_{15} TrdFreq_{ij} + \alpha_{16} Mom_{ij} + \alpha_{17} Volat_{ij} + \varepsilon_{ij}
\end{aligned}
\qquad [1]
$$

where $i = \{1, \ldots, 30\}$ represents the sample stocks and $j = \{1, \ldots, n_i\}$ represents the non-marketable limit orders of stock i. The dependent variable (EXEC) is an ordinal variable that takes three possible values: 1 if the order is cancelled before execution, 2 if the order is partially executed and then cancelled, and 3 if the order is fully executed. Our focus is on the likelihood of execution of non-marketable limit orders. So, we exclude market and marketable limit orders. We drop fleeting orders since they are not typically intended for execution but used for other purposes like fishing for hidden orders that better the opposing quote (Hasbrouck and Saar, 2009), and orders placed more than 20 ticks away from the prevailing best quotes.[12] Revisions of non-executed orders are treated as the same order while revisions of partially-executed orders are treated as new submissions. As explanatory variables, we include a dummy for hidden orders (*HLO*), dummies for HFT and AAT orders (*HFT* and *AAT*, NAT is absorbed in the intercepts), and interactions of trader type and order type dummies. As control variables, we include those used in the order exposure decision model in De Winne and d'Hondt (2007), and additional variables from the logistic regression model in BPV. These include order characteristics (aggressiveness, order size), state of the LOB (relative spread, same and opposite side depth, order book imbalance), and trading conditions (order imbalance, trading frequency, volatility, momentum, and time-of-the-day effects). All variable definitions are in the Appendix. The model is estimated with stock and day fixed effects and White-robust standard errors. Table VI reports the results.

[Table VI]

In Panel A, we report the estimated coefficients of the model. The first column shows results for all orders and the second and third columns are results disaggregated by buy and sell limit orders respectively. Coefficients are similar across all three columns and show consistent results. Specifically, we find that ATs' DLOs are less likely to be fully executed than NATs' DLOs

[12] As in Hasbrouck and Saar (2009), we define fleeting order as any non-marketable limit order that is cancelled unexecuted within 2 seconds from submission.

(negative *HFT* and *AAT* coefficients), likely reflecting ATs' higher cancelation-to-trade ratios. While NAT's HLOs are less likely to be fully executed than their own DLOs (negative *HLO* coefficient), we find the reverse pattern for ATs, most notably for HFTs (*HLO* + *HLOHFT* > 0).[13] Our findings therefore indicate that technology facilitates more effective use of HLOs, so that these orders, which lose time priority per the exchange trading rules, still have a higher execution probability. However, HFTs' HLOs are more likely to be fully executed than AATs' HLOs, indicating that the motivation behind placing these orders also matters.

In Panel B, we provide the average predicted probabilities of full execution (EXEC = 3) of a HLO placed by each type of trader, derived from the fitted model in Panel A. To estimate these probabilities, we fix the control variables at their cross-sectional average and then evaluate the model by setting all other trader types to zero (in effect estimating each case as if every HLO is placed by the same trader type). The resulting estimated probability of full execution for a HFT's HLO is about 87%, which is larger than for AATs (76%) and NATs (41%). Note that this estimated probability is comparable to the unconditional cross-sectional average frequency of full execution of HFTs' HLOs in our sample, which is 79.1% when the orders are placed at or within the best quotes, and 83.42% when placed within the second and fifth best quotes of the LOB. Overall, our results indicate that HFTs' HLOs have a higher likelihood of execution compared to the hidden orders of other traders.

B. Time to completion

To complement the previous analysis on execution probability, we examine the time to full execution of hidden orders placed by HFTs vis-à-vis other traders using survival analysis. Survival analysis can accommodate an important feature of limit order execution times: censored observations. An order cancelled unexecuted 30 minutes after submission apparently provides little information about execution time, but the fact it survived for 30 minutes is useful information. Such information contained in non-executed orders is used in survival analysis. We model time to completion using the econometric model specified in equation (2), which follows Lo, Mackinlay,

[13] In the Internet Appendix (Table IA-IX), we show that *HLOAAT* and *HLOHFT* have marginal positive effects on the probability of full execution and marginal negative effects on the probability of non-execution.

and Zhang (2002). The model describes an accelerated failure time specification of limit order execution under the generalized gamma distribution. The model is as below:

$$\begin{aligned} TIME_j = {} & \alpha_0 + \alpha_1 HFT_j + \alpha_2 AAT_j + \alpha_3 HLO_j + \alpha_4 HLOHFT_j + \alpha_5 HLOAAT_j + \\ & + \alpha_6 Aggr_j + \alpha_7 LastBuy_j + \alpha_8 DepthSame_j + \alpha_9 DepthSame^2{}_j + \\ & + \alpha_{10} DepthOpp_j + \alpha_{11} OrdSize_j + \alpha_{12} TrdFreq_j + \alpha_{13} RelTradeFreq_j + \varepsilon_j \end{aligned} \qquad [2]$$

where $TIME_j$ is the time to full execution of the j^{th} order, or the time survived in the book for a cancelled or expired order, with a positive censorship dummy. Control variables are the same used by Lo et al. (2002), which include order characteristics (aggressiveness, size), depth in the order book (same and opposite side), and market conditions (trading frequency, direction of the last trade). We add order type (*HLO*) and trader type (*HFT* and *AAT*) dummies, and interactions between them (*HLOHFT* and *HLOAAT*). Variable definitions are in the Appendix. As in the previous test, we exclude market and marketable limit orders, filter out fleeting orders, and drop orders placed more than 20 ticks away from the prevailing best quotes. Revisions of non-executed orders are treated as the same order while revisions of partially-executed orders are treated as new submissions. We estimate the model on a stock-by-stock basis. As before, we model the determinants of the execution of all orders, as well as buy and sell limit orders separately. Significance is evaluated using aggregated *t*-statistics computed as in Chordia, Roll, and Subrahmanyam (2005). Results are reported in Table VII.

[Table VII]

In Panel A, we provide the cross-sectional average estimated coefficients. To facilitate the interpretation, in Panel B we provide the predicted average time to completion for HLOs placed by different type of traders. BPV find that HLOs take longer to fully execute than similar DLOs. We find that is not true for HFTs, their hidden orders take shorter time to complete. Notably, the positive and significant coefficient on *HLO* indicates that NATs' limit orders behave similarly to the findings in BPV. However, for all orders, as well as for buy and sell limit orders separately, the negative and significant coefficient on *HLOHFT* offsets the positive coefficient for *HLO* (the coefficient on *HLOAAT* is also negative but about half the magnitude compared to *HLOHFT*). As a result, the estimated average time to full execution for HFTs' HLOs is significantly shorter than for AATs and NATs, as shown in Panel B. Differences in the estimated expected time to execution of HLOs between these groups are statistically significant.

C. Implementation shortfall of hidden order execution

The results in Tables VI and VII indicate that HFTs manage their hidden order execution more efficiently than the other traders such that they have a higher probability of, and lower time to, execution. But at what cost? We next estimate the costs HFTs face for their hidden order execution. To compute execution costs, it is important to note that iceberg orders, the type that are allowed in the NSE, are single (or parent) orders that are broken up into a sequence of smaller (child) orders. As the child orders are executed, they are recorded in the data as multiple smaller transactions in a sequence of orders. However, as Perold (1988) pointed out, the cost incurred by the trader is not a function of a single transaction but rather the entire sequence of child orders. To accommodate such order splitting, Perold (1988) introduced the "implementation shortfall" metric to measure transaction cost for the parent order. Implementation shortfall compares the value of a paper portfolio with no transaction costs to the real portfolio obtained by actual trading. It has been used in empirical work by Keim and Madhavan (1997), BPV, and Engle, Ferstenberg, and Russell (2012), among others.

For each limit buy order j, either DLO or HLO, we compute the implementation shortfall (IS) as the sum of the effective cost of execution (EFC) and the opportunity cost of non-execution (OPC), where

$$IS_j = EFC_j + OPC_j = ks\left(\overline{p} - q_0\right) + \left(1-k\right)s\left(q_C - q_0\right). \qquad [3]$$

The EFC_j component is the difference between the average execution price ($\overline{p}$) and the mid-quote at the time of order submission (q_0), multiplied by the amount of shares executed (ks), where s is the order size (in shares) and k is the fill rate of the order. The OPC_j is the difference between the closing price on the day of order submission (q_c) and q_0, multiplied by the unexecuted part of the order $(1-k)s$. Metrics for sell orders are analogously computed but conveniently signed.

Results are based on non-marketable limit orders. As in the previous analyses, we exclude fleeting orders and orders placed more than 20 ticks away from the prevailing best quotes. Revisions of standing limit orders are common in our data. We treat revisions of non-executed orders as the same order. In such cases, the IS is computed using *s* as the order size after the last revision. Revisions of partially-executed orders are treated as new submissions. After estimating the IS, EFC, and OPC for each order, we regress them on trader type dummies, controlling for order attributes, market conditions (trading frequency and volatility) in the 5 minutes prior to order submission, and order book characteristics. We use OLS to estimate the pooled regression model in equation (4) with stock and day fixed effects and White-robust standard errors.

$$\begin{aligned} IS_{ij} = {} & \alpha_0 + \alpha_1 HFT_{ij} + \alpha_2 AAT_{ij} + \alpha_3 HLO_{ij} + \alpha_4 HLOHFT_{ij} + \alpha_5 HLOAAT_{ij} + \\ & + \alpha_6 Aggr_{ij} + \alpha_7 OrderSize_{ij} + \alpha_8 Buy_{ij} + \alpha_9 TrdFreq_{ij} + \alpha_{10} Volat_{ij} + \\ & + \alpha_{11} FirstHalfHour_{ij} + \alpha_{12} LastHalfHour_{ij} + \varepsilon_{ij} \end{aligned} \qquad [4]$$

Here *i* represents the i^{th} stock and *j* represents the j^{th} order for stock *i*. All variable definitions are in the Appendix. For the execution cost component we provide results for all limit orders, and also for partially executed limit orders (fill rate > 0%); for the opportunity costs component, we provide results for all orders and also for non-executed or partially executed orders (fill rate < 100%) separately. Note that a fully executed order has zero opportunity cost and a completely non-executed order has zero execution cost. Results are reported in Table VIII.

[Table VIII]

In Panel A we show the regression coefficients for IS_j, and its components EFC_j and OPC_j. To facilitate interpretation, in Panel B we provide the estimated average IS, EFC, and OPC for the HLOs of each trader type. The relatively large negative and significant coefficient (-20.58) on the *HLOHFT* dummy indicates that, overall, the implementation shortfall for hidden orders placed by HFTs is the lowest among the three trader groups (2.43 in Panel B). While AATs also have a lower IS (13.6) than NATs (18.68), their hidden orders are not as effectively implemented as those of HFTs. Note also that the *HLO* dummy has a positive and significant coefficient, indicating that for NATs, hidden orders have a higher IS than DLOs. The same is true for AATs. Therefore, our results corroborate the trade-off between time to execution and IS reported in BPV for hidden orders in the Paris Bourse back in 2003.

To probe how HFTs manage to achieve such reduced IS for their hidden orders, we examine the estimates for its two components – EFC and OPC. Comparing the EFC across all three trader types, we find that HFTs' HLOs have the highest cost of execution. The *HLOHFT* coefficient is positive and significant (2.07) while it is negative for both AATs (-1.16) and NATs (-6.50). If we consider only those orders with fill rates greater than zero (that is orders that were at least partly executed), the magnitude of EFC for HFTs increases significantly. This indicates that HFTs face a higher effective cost when their hidden orders are executed, as shown in Panel B. This is consistent with our previous results which show that HFTs' HLOs are more aggressively placed in the order book than other traders' HLOs. The next two columns focus on the opportunity costs of non-execution. Here we find that HFTs' hidden orders have significantly lower opportunity cost of non-execution compared to the other traders. The *HLOHFT* dummy in these two regressions are much lower than for AATs and NATs, which results in substantially lower average OPC, as shown in Panel B. Thus, although HFTs face higher execution costs for hidden orders, their non-execution costs are lower, and much larger in magnitude, so that when aggregated the latter (OPC) effect dominates the former (EFC), leading to an overall lower implementation shortfall.

IV. The "why" question: Competing hypotheses

The widespread prevalence of HLOs placed by HFTs and their efficient execution lead us to ask why HFTs use these hidden orders? In other words, for what purposes are these orders used? Earlier we alluded to two hypotheses – trading on information and competing for liquidity provision – as possible explanations. Next, we test if either (or both) are confirmed in the data.

A. Information content of HLOs: Impulse response functions (IRFs)

Evidence in BPV from non-high-frequency markets suggest that HLOs are generally uninformed. Is this true of HFTs' hidden orders? To measure information content, we calculate the permanent price impact of each trader type's orders using an extended version of the Vector Autoregressive (VAR) model in Hasbrouck (1991a). The model is defined in event time (t), where an event may be a limit order submission, cancellation of a standing limit order, or a trade (market or marketable limit order submission). Revisions that improve (degrade) prices or increase (decrease) quoted depth are treated as limit order submissions (cancellations). For each trader type (HFT, AAT, and NAT) we consider both displayed and hidden orders. Recall that earlier tests showed that the likelihood of hiding orders increases with order aggressiveness, and that HFTs use more aggressive HLOs than other traders. Thus, we split each trader's displayed and hidden orders into aggressive and non-aggressive. We classify as aggressive (non-aggressive) any limit order placed at or within (beyond) the prevailing best quotes.[14] As a result of these partitions, the VAR model has 18 equations: one for the quote midpoint return and 17 (6 types of events x 3 types of traders-1) for order-flow related variables. The optimal number of lags is determined using the Schwarz' Bayesian Information Criterion for each stock-day. We exclude stock-days with less than 20 orders for each event, where an event is as defined above. The trade variable takes the value +1 (-1) for buyer- (seller-) initiated trades. Displayed, hidden, or cancellation variables on the ask (bid) side of the LOB take the value -1 (+1). We reset the trading process at the end of each day, resetting all lagged values to zero. The model is estimated in event (not transaction) time, so contemporaneous correlation is negligible. Nonetheless, we compute the IRFs such that any correlation is accounted for (methodological details are in the Internet Appendix). Estimates are cross-sectional averages with standard errors clustered by stock and day (Thompson, 2011). In Table IX we report the estimated impulse responses. We boldface those average estimated impacts for AATs and NATs that are significantly different from corresponding impacts estimated for HFTs.

[Table IX]

As expected, trades have the largest IRFs for all trader types. Among the three trader types, HFT-initiated trades have the largest estimated impact (1.16), which is significantly different from both NATs and AATs. This is consistent with recent findings for Canadian stocks (Broggaard, Hendershott, and Riordan, 2019). Of greater interest to us, however, is whether the permanent price impact of hidden orders placed by HFTs is larger than that of their own DLOs, which would suggest that HFTs choose to hide their orders when they want to trade on valuable signals. We find that HFTs' aggressive HLOs have an insignificant long-term cross-sectional average price impact (0.18, *p*-value = 0.11), suggesting they convey little information. Meanwhile, their aggressive DLOs have a significant and positive cross-sectional average price impact (0.25, *p*-value < 0.01). Thus, this test suggests that HFTs do not use HLOs to trade on information. In

[14] We drop the non-aggressive HFT HLOs because these orders lack enough observations to be modelled in a separate equation.

contrast, cross-sectional average permanent price impacts for both AATs' (0.35) and NATs' (0.49) HLOs are positive and significant, so we cannot reject the null that their HLOs are informationally motivated. These results also highlight that while HFTs and AATs both share the same technology for trading, their choice of order exposure is contingent on their differential motivation to trade.

B. Information content of HLOs: Information shares

As a second test of the information hypothesis, we estimate the information shares of HFTs' hidden orders vis-à-vis their displayed orders, as well as the orders of the other traders. To do this, we use the approach proposed in Hasbrouck (1995). Although much of the literature, including Hasbrouck (1995), uses this set-up to estimate information shares in a multi-market setting, it has also been used to assess information shares across different trader types. In our application, we assume that HFTs', AATs' and NATs' quotes share a common long-term (efficient) price process, and the information share attributable to each of these traders is the relative contribution of their innovations to the formation of the common efficient price. Methodological details are in the Internet Appendix. Table X presents the results.

[Table X]

Information share is greater for displayed orders than for hidden orders for all trader types. Within hidden orders of the different traders, HFTs have the lowest average information share at 6.13%, compared to 7.62% for AATs and 11.87% NATs. Tests of statistical significance show that these differences – both between HFTs and AATs and between HFTs and NATs – are significant. Overall, the results in Tables IX and X highlight that HFTs' hidden orders convey less information than their displayed orders, or the hidden (and displayed) orders of the other two trader types. Therefore, we cannot accept the hypothesis that HFTs use HLOs to trade on information.

C. Hidden orders for market making: Expected profit

Having found little support for the information hypothesis, we next test the alternative: HFTs hide their orders to limit competition in liquidity provision. To examine the validity of this hypothesis, we first model HFTs' decision to hide orders as a function of the expected market making profits. Following De Winne and d'Hondt (2007) and BPV, we use logistic models of order characteristics and market conditions to evaluate a trader's decision of whether to hide or display a limit order. Additionally, we use Tobit models of order characteristics and market conditions to study how, contingent on choosing to hide orders, a trader chooses the amount of shares to hide. We approximate the expected profits of market making by the average realized spread over recent trades (*ERzdSpr*) (Huang and Stoll, 1996).[15] For a limit order to sell (buy), *ERzdSpr* is the average of the realized spread for the five most recent buyer-initiated (seller-initiated) trades executed at least one second earlier. We impose the one second time constraint because realized spreads are computed using quote midpoints one second after the time of the trade. We exclude all market and marketable limit orders, fleeting orders, and orders placed beyond 20 ticks from the prevailing best quotes.

The model, shown in equation (5) below, is estimated on a stock-by-stock basis. In the logistic model, the dependent variable equals 1 when a trader submits a HLO and 0 otherwise. In the Tobit model, the dependent variable is the amount of shares hidden, normalized by the stock's average daily trading volume. Control variables in both models are the same and are described in the Appendix. In Table XI, we report cross-sectional averaged coefficients. Statistical significance is evaluated using the aggregated *t*-statistics as in Chordia, et al. (2005). We report results for the logistic model in Panel A, and the Tobit model in Panel B. For comparison purposes, we estimate both models for HFTs, AATs, and NATs separately.

$$\begin{aligned} Y_j = {} & \alpha_0 + \alpha_1 Aggr_j + \alpha_2 OrdSize_j + \alpha_3 ERzdSpr_j + \alpha_4 DepthSame_j + \\ & + \alpha_5 DepthOpp_j + \alpha_6 Volat_j + \alpha_7 WaitTime_j + \alpha_8 TrdFreq_j + \\ & + \alpha_9 HidVolSame_j + \alpha_{10} LOBImb_j + \alpha_{11} LastTrdSize_j + \\ & + \alpha_{12} MVolat_j + \alpha_{13} LastHalfHour_i + \varepsilon_i \end{aligned} \qquad [5]$$

[Table XI]

The coefficient of interest here is on *ERzdSpr*. For HFTs this coefficient is positive and significant, indicating that when the expected realized spread (profit to market making) increases, HFTs are more likely to hide orders (Panel A) and likely to hide a greater amount of shares per order (Panel B). By contrast, this is not true for NATs and is weaker for AATs. Thus, the

[15] Realized spread is calculated as two times the difference between the quote midpoint h = {1,5,30} seconds after a trade and the trade price. For seller-initiated trades, we multiply the above magnitudes by -1.

competition for liquidity provision hypothesis finds confirmation in our data. Note that the positive and significant coefficient on price aggressiveness (in both panels) for AATs and HFTs indicate that these algorithmic traders show an interest in assuming positions aggressively and place orders closer to the prevailing best quotes, but hide them so as not to expose their trading interest. The same is not true of NATs.

D. Hidden orders for market making: Undercutting

If HFTs use hidden orders to aggressively compete to provide liquidity, then we should see greater use of these orders to undercut existing orders at the top of the book. To test if this is true, we identify all limit order submissions (including revisions) that offer price improvement for a standing limit order and/or increase the size of standing limit orders, for each trader/order type. We define an undercutting limit order as a limit order that (a) is placed immediately after another submission on the same side of the market, (b) comes in under 10 milliseconds of the previous order, and (c) improves the price of the previous one. In Panel A of Table XII we present unconditional cross-sectional average statistics on the use of undercutting orders restricted to the five best quotes; conclusions remain unchanged if we consider only the best quotes.

[Table XII]

Of the three trader types, HFTs use the highest proportion of hidden orders to undercut. For example, they use 5.02% of hidden orders to undercut at a lower level of stock activity (at least 20 orders of each type – hidden or displayed – by each trader type per stock-day). Not surprisingly, they also use displayed orders (3.01%) to trade ahead of standing quotes. Expectedly, NATs show the least amount of such undercutting activity, both for hidden and displayed orders. Results are similar for higher (50 orders) levels of stock activity.

This evidence, while illustrative, does not take account of market conditions. From BPV and our earlier results, we know that order exposure is affected by both stock and market attributes. Thus, we next estimate the logit regression in equation (6) to examine whether the observed higher likelihood of undercutting by HFTs' hidden orders remains after controlling for market conditions and the state of the order book.

$$
\begin{aligned}
Und_j = {} & \alpha_0 + \alpha_1 HFT_j + \alpha_2 AAT_j + \alpha_3 HLOHFT_j + \alpha_4 HLOAAT_j + \alpha_5 HLONAT_j + \\
& + \alpha_6 DispSizeUnd_j + \alpha_7 AggrUnd_j + \alpha_8 HidVolSame_j + \alpha_9 Rspr_j + \\
& \alpha_{10} DepthSame_j + \alpha_{11} DepthOpp_j + \alpha_{12} Volat_j + \varepsilon_j
\end{aligned} \qquad [6]
$$

where the dependent variable (*Und*$_j$) is a dummy that takes the value of one if an order is an undercutting order as defined earlier, zero otherwise. As in previous tests, we include dummies to control for trader types (*HFT* and *AAT*, NAT is captured in the intercept) and the interaction of trader types with HLO. For controls, first we include the displayed size of the undercut order (*DispSizeUnd*$_j$). We expect that when the undercut order has a larger displayed size, HFTs are more likely to undercut it. Second, we consider the aggressiveness of the undercut order (*AggrUnd*$_j$). Aggressiveness is defined as the number of ticks away from the best quote on the same side. The further the undercut order is from the best quotes, in other words less aggressive, the less likely it is to be undercut. Thus, we expect a negative relationship between the aggressiveness of the undercut order and its chance of being undercut. Finally, we include a variable that gauges the possibility of hidden order detection. This variable, *HidVolSame*, is a dummy that takes the value of one if the presence of hidden volume on same side has been revealed, zero otherwise. Hidden volume is revealed at the time an undercutting order is placed if the quantity traded at the prevailing best quote is greater than the displayed depth, which is only possible if there was additional (hidden) volume at the best quotes (see Pardo and Pascual, 2012). Additional controls are as defined in the Appendix.

Panel B of Table XII presents the estimation output. The model is estimated on a stock-by-stock basis. We report cross-sectional average coefficient estimates and evaluate statistical significance using aggregated *t*-statistics using the approach in Chordia et al. (2005). Aligning with our expectation, displayed size of the undercut order is positively related to the likelihood of an order used for undercutting, confirming that larger orders are more likely to be undercut. Likewise, when an order is closer to the top of the book (more aggressive), it is more likely to be undercut (negative and significant coefficient on *AggrUnd*). *HidVolSame* is positive indicating that when traders can infer the presence of hidden volume at the best quotes, they are more likely to place orders to trade ahead of these hidden orders. In fact, the odds ratio shows that this likelihood is

1.56 times (or 56% more) compared to the use of displayed orders by NATs (we use the DLOs of NATs as the reference group for the odds ratio calculation in this Panel).

The main variable of interest is *HLOHFT*. The coefficient on this variable is 0.41 and it is significant at the 1% level. Compare this to the negative coefficients on *HLOAAT* and *HLONAT*. Clearly, HFTs use hidden orders for undercutting, while the two other trader groups are less likely to do the same. In fact, the odds ratio for *HLOHFT* is greater than 1 (1.51) while for both other groups it is lower than 1, indicating that while HFTs use hidden orders to undercut the standing quotes at or near the top of the order book, the two other groups are less likely to use hidden orders for the same purpose.

E. Effect of undercutting on profits to market making

We round out our investigation by showing how undercutting for liquidity competition affects HFTs' profits to market making. We use effective and realized spreads to evaluate the profitability of hidden and displayed orders submitted by the different trader types. For a buyer-initiated trade, the effective spread is twice the difference between the trade price and the quote midpoint prevailing before the trade. We estimate the pooled regression model (7), defined on a trade-by-trade basis, where we regress the realized spread of each trade on *passive*[16] trader type (HFT and AAT) indicators, passive order type (HLO) indicators, trade characteristics (size, duration), LOB characteristics (depth imbalance, depth on the same side), market conditions over the last minute (volume, volatility, order imbalance), and time of day dummies. The model is estimated by OLS with stock and day fixed effects and White-robust standard errors.

$$\begin{aligned} RzdSpr_{it} = {} & \alpha_0 + \alpha_1 HFT_{it} + \alpha_2 AAT_{it} + \alpha_3 HLO_{it} + \alpha_4 HLOHFT_{it} + \alpha_5 HLOAAT_{it} + \\ & + \alpha_6 LOFirst_{it} + \alpha_7 RelSize_{it} + \alpha_8 TDur_{it} + \alpha_9 TDurSame_{it} + \alpha_{10} HFTtaker_{it} + \\ & + \alpha_{11} AATtaker_{it} + \alpha_{12} LOBImb_{it} + \alpha_{13} DepthSame_{it} + \alpha_{14} Vol_{it} + \\ & + \alpha_{15} Volat_{it} + \alpha_{16} OI_{it} + \alpha_{17} HVolOpp_{it} + \alpha_{18} FirstHalfHour_{it} + \alpha_{19} LastHalfHour_{it} + \varepsilon_j \end{aligned} \qquad [7]$$

We provide the estimated coefficients in Table XIII together with Wald tests on specific null hypotheses. Consistent with HFTs placing HLOs more aggressively than DLOs, we find that both effective spreads and realized spreads for all *h* are larger for HFTs' DLOs than for HFTs'

[16] This is the counterparty to the trade-initiator.

HLOs (*HLO* + *HLOHFT* < *0*). We have previously shown that HFTs hide limit orders in anticipation of higher realized spreads (Table XI), and they often do so by undercutting standing orders near the top of the book (Table XII). Results in Table XIII indicate that this aggressive competition to exploit profit opportunities in market making ends up cutting into HFTs' profits. However, and despite the fact that HFTs are more aggressive than other traders in placing their HLOs (Table IV), they manage to obtain significantly higher realized spreads for these orders than AATs (*HFT* + *HLOHFT* > *AAT* + *HLOAAT*) and NATs (*HFT* + *HLOHFT* > 0), as the Wald tests show. Therefore, HFTs extract higher rents than other traders when using HLOs for market making purposes.

V. Robustness checks

We conduct numerous robustness checks to confirm the validity of our findings. For brevity, here we simply catalog the tests and provide the tabulated results with detailed table headers in the Internet Appendix. In Table AI-I (AI-II) we examine hidden order usage in mid- and small-cap stocks in the NSE (Nasdaq) and find that HFTs place substantial proportions of HLOs in these firm size categories as well, in both markets. Thus, their hidden order activity is not limited to large cap stocks only. In Table AI-IV, we expand the (NSE) sample to three months and find the same relative percentage use of hidden orders by HFTs, AATs and NATs (compare with Panel A of Table III for large cap stocks, and Table AI-I for mid and small cap stocks). Thus, our results are temporally stable.

In Table AI-V we show hidden order placement in the order book for mid and small cap firms. Results corroborate those presented in the paper, in that HFTs place the majority of their HLOs within or at the best quotes, while NATs place over 40% of their HLOs at the "far" depths of the order book and AATs falling somewhere in between. Table AI-VI shows that these results hold when we implement filters that have been used in previous studies, in particular when we include only those orders that were placed within the first 30 minutes of the trading session (Panel A), after the first 30 minutes of the trading session (Panel B), filter out fleeting orders as in Hasbrouck and Saar (2009) (Panel C), drop orders placed more than 20 ticks (Panel D) or 10 ticks (Panel E) away from the prevailing best quotes. Table AI-VII extends the sample to three months and shows that the results on HFTs' (and other traders') HLO placement in the order book hold for the longer time period as well.

We showed in Table V the size distribution of HLOs and DLOs placed by HFTs, AATs, and NATs for our main sample. In Table AI-VIII we show that the same trends also hold for mid and small cap stocks. HFTs use smaller sized HLOs in both mid and small cap stocks compared to AATs and NATs. In fact, this result is more pronounced in these smaller firm sizes than in our main (large cap) sample. Table AI-IX extends the sample period and finds similar results. Table AI-X presents the estimated marginal effects and predicted probabilities of hidden order execution associated with the results in Table VI and shows that HFTs have a lower probability of non-execution and partial execution, and higher probability of full execution of hidden orders. Table AI-XI models the survival analysis, now by disaggregating trader types. Results corroborate the findings in Table VII in that HFTs' HLO sit the least amount of time in the order book, followed closely by AATs. NATs' HLOs, on the other hand, survive significantly longer. Tables AI-XII and AI-XIII verify the implementation shortfall results after excluding the filters used in the main analysis and using stock-by-stock (instead of pooled) regressions, respectively.

To test whether HFTs' HLOs are motivated by information reasons, we conduct an additional test. For each stock-day, we decompose the efficient price variance due to the order flow (OF) into its components using the Hasbrouck (1991b) approach. This method provides an estimate of the efficient variance that is split into an OF-related and an OF-unrelated component. We estimate the share of each trader type in the OF-related component; our interest is in examining the information content of HFTs' hidden orders. In Table AI-XIV we show that HFTs' HLOs contribute the smallest (0.46%) to OF-related price variation among the three trader types and shows, yet again, that HFTs' hidden orders convey less information into price, when compared with either their displayed orders, or with the displayed and hidden orders of AATs and NATs. In Table AI-XV we show that HFTs place more hidden orders and choose to hide larger amounts in anticipation of higher expected profits to market making (*ERzdSpr*) even when we remove the filters from our main analysis (Table XI). Finally, in Table AI-XVI we corroborate the result that HFTs use hidden orders to undercut by restricting the analysis to the undercutting of orders placed either at or within the prevailing best quotes. Figure IA-I shows the probability of hiding an order, conditional on order size and volatility, separated out by HFTs versus non-HFTs in the Nasdaq sample and shows that the order size results (in Figure 3) obtain even when we examine only high volatility days (the five days with the highest daily realized volatility).

VI. Discussion and Conclusion

The impact of transparency is a much-debated topic in the literature with mixed conclusions.[17] Consistent with transparency having drawbacks, all major exchanges allow traders to hide their trading interest by placing hidden orders. This paper expands and updates our understanding of endogenous transparency by studying how different traders with differing technology and motivations to trade use hidden orders.

Past research concluded that patient liquidity providers use the option to hide when they want to transact large quantities and avoiding picking off risks (BPV, De Winne and d'Hondt 2007, Buti and Rindi, 2013). However, these empirical and theoretical findings are from markets and models without ATs, who now make up the majority of volume in many markets (the U.S., Japan, and Europe, for example) and an increasing fraction in others (India, and China, for example). How order exposure is impacted by the availability of faster trading technology as well as the differing motivations of different trader types has not been examined.

In this paper we provide, to our knowledge, the first comprehensive account of hidden order use by ATs in general and HFTs in particular. This study is made possible by our access to data from the NSE that identify the types of traders and the order handling system they use. Thus, we can identify proprietary traders who use algorithmic order entry and management systems (HFTs), algorithmic traders who trade for clients (AATs), and non-algorithmic traders (NATs) and

[17] Bloomfield and O'Hara (1999) use an experimental financial market to show that increased post-trade transparency increases bid-ask spreads. Madhavan, Porter, and Weaver (2005) use a change in dissemination of data on the limit order book on the Toronto Stock Exchange and find that increased pre-trade transparency also increases bid-ask spreads. On the other hand, Flood, Huisman, Koedijk, and Mahieu (1999) find that increased pre-trade transparency leads to lower spreads and increased trading volume in an experimental multiple-dealer market. Boehmer, Saar, and Yu (2005) find that the improvement of pre-trade transparency on the NYSE via the introduction of Open Book generally increases in stock market liquidity.

examine their hidden order use. Where possible, we also use data from Nasdaq with HFT and non-HFT identifiers to validate our findings from the NSE.

We find that ATs make extensive use of hidden orders. HFTs do not use hidden orders to transact large volumes or avoid picking-off risk; instead they choose to hide small-sized orders placed aggressively nearer the top of the book. This pattern is different from the NATs, who hide large orders placed away from the best quotes, a finding that was documented in earlier studies. HFTs' hidden orders are more efficiently executed with a higher likelihood of execution and shorter time to completion compared to the hidden orders of other traders as well as their own displayed orders. Although HFTs' aggressive HLOs have a higher effective cost of execution, they have very low rates of non-execution, which reduces their overall implementation shortfall.

Given the contrast in use of hidden orders between HFTs and NATs, we ask why HFTs hide orders in the first place. HFTs and AATs have the technology to monitor markets in real-time and to cancel standing quotes faster than other traders. So why give up time priority and place hidden orders? We offer two hypotheses – to trade on time-sensitive information, and to compete for liquidity provision, in both cases without revealing their presence.

To test the first hypothesis, we assess the information content of hidden orders using two different tests –the permanent price impact as represented by impulse response functions, and the Hasbrouck (1995) information share measure. Both tests indicate that hidden orders placed by HFTs have lower information content than their displayed orders, as well as the hidden orders of the other two trader groups. Thus, HFTs' hidden order use does not support the information motive. To test the market making hypothesis, we examine and find that HFTs place hidden orders when expected realized spreads are higher. To that end, they use hidden orders to undercut standing orders at or near the best quotes. Such undercutting reduces the realized profits to market making.

The fact that different types of traders use hidden orders so differently raises the question of whether hidden orders favors certain traders over others. Given that algorithmic trading and high-frequency trading are so prevalent, this question is of interest to both market participants and regulators. Consistent with the literature from before algorithmic trading, we find that NATs use hidden orders to transact relatively larger orders with less information. Technology allows ATs to make greater use of hidden orders, suggesting that they benefit from the ability to hide. The impact of HFTs' use of hidden orders is more difficult to understand. HFTs appear to hide orders to limit competition in liquidity provision. If markets precluded HFTs from hiding orders, this could potentially increase competition. However, it could also lead HFTs to place fewer limit orders, reducing competition in liquidity provision. Overall, our findings underscore the need for further study to better understand how hidden order usage impacts different types of traders, as well as the market overall.

References

Bhattacharya, Neil, Bidisha Chakrabarty, and Xu Frank Wang. 2020. High-frequency traders and price informativeness during earnings announcements. *Review of Accounting Studies*, forthcoming.

Bessembinder, Hendrik, Marios Panayides, and Kumar Venkataraman. 2009. Hidden liquidity: an analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics* 94(3), 361-383.

Bloomfield, Robert, and Maureen O'Hara. 1999. Market transparency: Who wins and who loses? *The Review of Financial Studies* 12(1): 5-35.

Bloomfield, Robert, Maureen O'Hara, and Gideon Saar. 2015. Hidden liquidity: Some new light on dark trading. *Journal of Finance* 70(5), 2227-2274.

Boehmer, Ekkehart, Dan Li, Gideon Saar. 2018. The competitive landscape of high-frequency trading firms. *The Review of Financial Studies* 31(6): 2227–2276.

Boehmer, Ekkehart, Gideon Saar, Lei Yu. 2005. Lifting the veil: An analysis of pre-trade transparency at the NYSE. *Journal of Finance* 60(2): 783–813.

Boulatov, Alex. and Thomas. J. George. 2013. Hidden and displayed liquidity in securities markets with informed liquidity providers. *Review of Financial Studies* 26(8): 2096-2137.

Boussetta, Selma, Laurence D. Lescourret, and Sophie Moinas. 2017. The role of pre-opening mechanisms in fragmented markets. *SSRN Working Paper.*

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. 2014. High frequency trading and price discovery. *Review of Financial Studies* 27(8): 2267-2306.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. 2019. Price discovery without trading: Evidence from limit orders. *Journal of Finance*, 74(4): 1621-1658.

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

Buti, Sabrina, and Barbara Rindi. 2013. Undisclosed orders and optimal submission strategies in a limit order market. *Journal of Financial Economics* 109(3): 797-812.

Carrion, Allen. 2013. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets* 16(4): 680-711.

Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam. 2005. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics* 76(2): 271-292.

Copeland, Thomas E., and Dan Galai. 1983. Information effects on the bid-ask spread. *The Journal of Finance* 38(5): 1457-1469.

De Winne, Rudy, and Catherine D'Hondt. 2007. Hide-and-seek in the market: placing and Detecting hidden orders. *Review of Finance* 11(4): 663-692.

Engle, Robert, Robert Ferstenberg, and Jeffrey Russell. 2012. Measuring and modeling execution cost and risk. *The Journal of Portfolio Management* 38(2): 14-28.

Flood, Mark, Ronald Huisman, Kees Koedijk, and Ronald Mahieu. 1999. Quote disclosure and price discovery in multiple-dealer financial markets. *Review of Financial Studies* 12 (1): 37-59.

Foley, Sean, and Talis J. Putnins. 2016. Should we be afraid of the dark? Dark trading and market quality. *Journal of Financial Economics* 122(3): 456-481.

Hagströmer, Björn, and Lars Nordén. 2013. The diversity of high-frequency traders. *Journal of Financial Markets* 16(4): 741-770.

Hasbrouck, Joel. 1991a. Measuring the information content of stock trades, *Journal of Finance* 46(1): 179–207.

Hasbrouck, Joel. 1991b. The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4(3): 571–595.

Hasbrouck, Joel. 1995. One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50(4): 1175–1199.

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

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. 2011. Does algorithmic trading improve liquidity? *The Journal of Finance* 66(1): 1-33.

Hendershott, Terrence, and Ryan Riordan. 2013. Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis* 48(4): 1001-1024.

Hirschey, Nicholas. 2018. Do high-frequency traders anticipate buying and selling pressure? *SSRN Working Paper*.

Hoffmann, Peter. 2014. A dynamic limit order market with fast and slow traders. *Journal of Financial Economics* 113(1): 156-169.

Huang, Roger D., and Hans R. Stoll, 1996. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics* 41(3): 313-357.

Kahraman, Bige, and Heather Tookes. 2017. Trader leverage and liquidity. *The Journal of Finance* 72(4): 1567-1610.

Keim, Donald B., and Ananth Madhavan. 1997. Transactions costs and investment style: an inter-exchange analysis of institutional equity trades. *Journal of Financial Economics* 46(3): 265-292.

Korajczyk, Robert A., and Dermot Murphy. 2019. High frequency market making to large institutional trades. *Review of Financial Studies* 32(3): 1034-1067.

Lo, Andrew W., A. Craig MacKinlay, and June Zhang. 2002. Econometric models of limit-order executions. *Journal of Financial Economics* 65(1): 31-71.

Madhavan, Ananth, David Porter, and Daniel Weaver. 2005. Should securities markets be transparent? *Journal of Financial Markets* 8(3): 265-287.

Massey Jr. Frank J. 1951. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American Statistical Association* 46(253): 68-78.

Moinas, Sophie. 2010. Hidden limit orders and liquidity in order driven markets. Unpublished working paper. TSE Working Paper Series 10-147.

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

Pardo, Angel, and Roberto Pascual. 2012. On the hidden side of liquidity. *The European Journal of Finance* 18(10): 949-967.

Pascual, Roberto, and David Veredas. 2010. Does the open limit order book matter in explaining informational volatility? *Journal of Financial Econometrics* 8(1): 57–87.

Perold, Andre F. 1988. The implementation shortfall: Paper versus reality. *The Journal of Portfolio Management* 14(3): 4-9.

Ranaldo, Angelo. 2004. Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7(1): 53-74.

Thompson, Samuel B. 2011. Simple formulas for standard errors that cluster by both firm and time. *Journal of Financial Economics* 99(1): 1-10.

Yao, Chen, and Mao Ye. 2018. Why Trading Speed Matters: A Tale of Queue Rationing under Price Controls. *Review of Financial Studies* 31 (6): 2157–2183.

Table I
NSE sample descriptive statistics

This table provides daily cross-sectional average statistics for 100 stocks listed on the National Stock Exchange (NSE) of India. Cross-sectional averages are computed from daily averages per stock. The sample period is December 2013 (21 trading days). The sample comprises market-capitalization-based subsamples of 30 (largest), 40 (medium), and 30 (smallest) stocks. Market capitalization is in billions of Rupees. Volume is in 10,000-share units, number of trades is in 100-trade units, depth is in 1000-share units, and Price is in Rupees. Daily volatility is 100((maximum price/minimum price) –1). Relative bid-ask spread is the ratio of the quoted spread to the quote midpoint, in basis points. Relative effective spread is two times the difference between the average trade price and the quote midpoint, divided by the quote midpoint. Displayed (hidden) depth is the accumulated displayed (non-displayed) depth in the limit order book (LOB). MT (message traffic) is the number of order messages (sum of submissions, cancellations, and revisions) in 1000-message units. MT/Trd is the ratio of MT to trades, and CAN/Trd is the ratio of cancellations to trades. Share in MT is each trader type's share in message traffic, where trader types are high frequency traders (HFT), agency algorithmic traders (AATs), and non-algorithmic traders (NATs). Metrics are generated from 1-minute snapshots of the LOB and averaged across observations. Statistical significance is evaluated using the non-parametric Wilcoxon rank-sum test. In Panel A, "***", "**", "*" on "Mid" ("Small") indicate statistically different from the "Large" ("Mid") subsample at the 1%, 5%, and 10% levels, respectively. In Panel B, significance under the AATs column tests for the difference between HFTs and AATs, and in the NATs column tests the differences between algorithmic traders (HFTs and AATs) and NATs.

Panel A: Sample statistics

		Market-capitalization-based subsamples		
	All	Large	Mid	Small
Market capitalization (billions)	457.31	1493.84	21 ***	2.52 ***
Volume ('0000)	88.92	226.38	44.93 ***	10.11 *
Number of trades ('00)	103.74	297.69	29.14 ***	9.26 ***
Volatility	40.81	29.41	42.48 ***	49.97 **
Relative bid-ask spread (bps)	39.33	6.15	38.89 ***	73.11 ***
Displayed depth ('000)	112.85	217.12	90.81 ***	37.98
Hidden depth ('000)	28.68	57.26	21.29 ***	9.96
Price (Rupees)	317.56	619.02	263.32 ***	88.4 ***

Panel B: Message traffic per trader type and subsample

		Trader types		
Subsample	Variable	HFTs	AATs	NATs
Large	MT	1098.98	127.74 *	37.09 ***
	MT/Trd	233.31	23.47 ***	2.31 ***
	CAN/Trd	9.55	1.90 ***	0.25 ***
	Share in MT	56.41	26.27 ***	17.32 ***
Mid	MT	7.34	12.21 *	5.61
	MT/Trd	228.03	121.99	3.02 ***
	CAN/Trd	24.91	2.27	0.55 **
	Share in MT	16.71	42.67 ***	40.62 ***
Small	MT	1.17	4.77 ***	1.86 ***
	MT/Trd	30.88	112.62	3.49 ***
	CAN/Trd	3.84	1.48	0.68 **
	Share in MT	6.74	54.85 ***	38.42 ***

Table II
Nasdaq sample LOB descriptive statistics

We provide limit-order-book-based cross-sectional average descriptive statistics on the 120 stocks included in the Nasdaq's HFT database. The analysis is based on order by order data collected from one-minute snapshots of the 10 best ask and bid LOB levels. We consider the whole sample period covered by the database, except the second week of September 2008 (Lehman Brother's failure). The database includes a HFT "flag" that identifies the HFT orders in the LOB. It also includes a "HLO" flag that allows us to distinguish HLOs from DLOs. We provide separated statistics for the displayed LOB, that is, ignoring hidden volume, and for the whole LOB (displayed plus hidden). The "Displayed relative spread", for example, differs from the "Posted relative spread" in that the former ignores HLOs placed within the displayed best quotes. "Hidden volume at the best quotes" is the percentage of time hidden volume is present at the market quotes (either displayed or hidden). Finally, we provide statistics on the HFTs' contribution to liquidity supply. "HFTs at the best quotes" is the percentage of time HFTs post either the best ask or bid quote (either hidden or displayed). Similarly, "HFTs contribution to the best quote depth" is the percentage of total depth (hidden plus displayed) at the best quotes (hidden or displayed) provided by the HFTs' standing limit orders. We split the sample of 120 Nasdaq-listed firms into 3 subsamples of 40 stocks based on market capitalization. Tests are based on the Wilcoxon non-parametric rank-sum test. ***, **, * means statistically different than the corresponding statistic for large caps at the 1%, 5%, or 10% of statistical significance.

		Market-capitalization-based subsamples		
Cross-sectional daily average statistic	All	Large	Mid	Small
Displayed relative spread (bsp)	31.99	7.40	22.57 ***	65.99 ***
Posted relative spread (bsp)	26.27	6.31	19.07 ***	53.41 ***
Displayed depth at the best quotes ($US)	62621.35	153819.71	28198.41 ***	5845.91 ***
Total depth at the best quotes ($US)	96453.74	224053.26	49520.39 ***	15787.58 ***
Displayed depth 5 best quotes ($US)	445666.30	1136466.80	157180.80 ***	43351.34 ***
Total depth 5 best quotes ($US)	555882.43	1362283.70	224485.33 ***	80878.30 ***
Hidden volume at the best quotes (% time)	71.72	78.95	69.30 ***	66.91 ***
Hidden volume within the best displayed quotes (% time)	72.08	48.99	77.75 ***	89.50 ***
HFTs at the best quotes (% time)	67.41	93.23	62.89 ***	46.11 ***
HFTs contribution to the best quotes depth (%)	30.90	50.41	25.36 ***	16.92 ***
HFTs contribution to the 5 best quotes depth (%)	26.80	36.84	23.58 ***	19.97 ***

Table III
Use of Hidden Limit Orders

This table provides cross-sectional average daily statistics on the use of hidden limit orders (HLOs) and displayed limit orders (DLOs) per trader type for a random sample of 30 NSE-listed large caps (Panel A), and the 40 largest stocks in the Nasdaq HFT database (Panel B). In Panel A, we distinguish between high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs). The sample period is December 2013 (21 trading days), and the analysis is performed with all the limit order submissions within that period. In Panel B, we distinguish between HFTs and non-HFTs. The sample period covers selected weeks between 2008 and 2010. The analysis is based on order by order data collected from one-minute snapshots of the 10 best ask and bid LOB levels. In each Panel, we provide the proportion of HLOs, both in the number of orders (Ord.), and the accumulated volume (Vol.), relative to all limit orders submitted by each trader type. We also show each trader type's share of both HLOs and DLOs. In Panel A, significant differences in medians between HFTs and other trader types are shown beside AAT, and NAT numbers. In Panel B, differences between HFTs and non-HFTs are explicitly reported together with the significant test. We use the non-parametric Wilcoxon rank-sum test. ***, **, * indicate statistically different at the 1%, 5%, and 10% levels respectively.

Panel A: NSE

	HFTs		AATs		NATs	
Variable	Ord.	Vol.	Ord.	Vol.	Ord.	Vol.
Relative use of HLOs (%)	8.92	8.62	26.28 ***	34.26 ***	9.14 ***	30.88 ***
Share of DLOs (%)	33.79	55.18	22.50 ***	8.28 ***	43.71 ***	36.55 ***
Share of HLOs (%)	7.06	2.98	54.03 ***	31.00 ***	38.91 ***	66.02 ***

Panel B: Nasdaq

	HFTs		non-HFTs		Diff.	
Variable	Ord.	Vol.	Ord.	Vol.	Ord.	Vol.
Relative use of HLOs (%)	21.80	15.25	15.39	21.74	6.40 ***	-6.49 **
Share of DLOs (%)	42.00	27.14	58.00	72.86	-16.01 ***	-45.72 ***
Share of HLOs (%)	44.00	16.83	56.00	83.17	-12.01 ***	-66.34 ***

Table IV
Hidden Limit Order placement in the order book

We examine the placement of hidden limit orders (HLOs) throughout the LOB by different types of traders. All the statistics are reported in percentage terms over either the number of orders or the share volume. In Panel A, we provide results for the NSE of India. We build snapshots of the limit order book (LOB) at the time of each new order submission and group the LOB levels into four segments: (a) better than the standing quotes ("Better"), (b) at the best quotes ("At"), (c) from the best quotes up to 5 ticks away ("Near"), and (d) the rest ("Far"). We report findings for the 30 largest stocks in a market-capitalization-representative sample of 100 stocks. Regarding the trader types, we distinguish between high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs). In Panel B, we provide results for the Nasdaq. In this case, the database provides order-by-order information on one-minute snapshots of the ten best levels of the LOB. We group the LOB levels into three segments: (a) at the best quotes ("At"), (b) from the best quotes up to 5 ticks away ("Near"), and (c) the rest ("Far"). We report findings for the 40 largest stocks (out of 120) in the database. As for the trader types, we distinguish between HFTs and non-HFTs. Each statistic reported is the cross-sectional average of the time series mean of the daily proportion of orders at the different LOB level groups. We average ask and bid quotes. ***, **, * indicate statistically different than corresponding HFTs' statistic at the 1%, 5% and 10% level, respectively.

Panel A: NSE

wrt. the	Orders			Volume (sh.)		
best quotes	HFTs	AATs	NAT	HFTs	AATs	NAT
Better	52.21	14.53 ***	16.39 ***	48.98	9.55 ***	10.43 ***
At	19.39	45.81 ***	18.56	23.53	42.19 ***	26.21 **
Near	26.14	29.60 ***	26.58	24.36	30.46 ***	27.27 ***
Far	2.25	10.07 ***	38.47 ***	3.14	17.80 ***	36.09 ***

Panel B: Nasdaq

	HFTs	non-HFTs	HFTs	non-HFTs
At	30.26	22.45 ***	32.92	23.70 ***
Near	59.13	47.21 ***	52.44	36.17 ***
Far	10.61	30.34 ***	14.64	40.13 ***

Table V
Order size

We provide cross-sectional average daily statistics on the empirical distribution of the size of hidden limit orders (HLOs) and displayed limit orders (DLOs). In Panel A, we report findings for the NSE of India. The sample consists of the 30 largest stocks in a representative (by market capitalizations) sample of 100 NSE-listed stocks on December 2013. We distinguish between three type of traders: high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs). The analysis is based on order-by-order data that we group according to the full (displayed plus non-displayed) order size. In Panel B, we provide our findings for the 40 largest stocks in the Nasdaq's HFT database. In this case, the analysis is based on order-by-order data for regular one-minute snapshots of the ten best levels of the LOB. We can distinguish HFTs' orders from non-HFTs' orders. Order size categories, defined in total (both displayed and hidden) shares, are market-specific. The order size cutoffs for each market are selected based on the empirical distribution of the order size of all limit orders submissions. We provide the percentage of HLOs and DLOs in each order-size category per trader type. We use the two-sample Kolmogorov-Smirnov (Massey, 1951) test to compare the order size distributions of HLOs and DLOs submitted by the different trader types. ***, **, * indicate statistically different than corresponding HFTs' statistic at the 1%, 5% and 10% level, respectively.

Panel A: NSE

	HFTs		AATs		NATs	
Order size distrib. (%)	DLOs	HLOs	DLOs	HLOs	DLOs	HLOs
(0,50]	4.98	68.34	62.50	57.25	64.71	28.12
(50,75]	0.95	9.93	11.03	8.42	1.54	2.79
(75,100]	1.39	0.88	4.34	6.05	11.25	11.07
(100,200]	25.97	4.54	10.44	11.59	6.66	12.23
(200,500]	43.35	12.96	8.70	10.49	9.23	22.82
(500,1000]	18.36	2.92	1.66	3.63	3.55	10.87
(1000,2500]	3.20	0.39	0.80	1.88	1.62	5.99
>2500	1.79	0.04	0.54	0.70	1.44	6.12
HFTs vs. AATs/NATs (p-value)				0.00		0.00
DLOs vs. HLOs (p-value)		0.00		0.00		0.00
Average size (sh.)	1179.41	255.74	318.40 ***	601.19	320.77 **	1192.80 ***

Panel B: Nasdaq

	HFTs		non-HFTs	
	DLOs	HLOs	DLOs	HLOs
(0, 100]	47.69	82.83	37.14	42.01
(100, 500]	17.30	12.20	14.67	15.01
(500, 1000]	6.19	1.53	6.41	5.95
(1000, 5000]	13.04	1.99	16.69	11.95
(5000, 10000]	3.30	0.39	3.87	2.93
>10000	12.47	1.07	21.22	22.15
HFTs vs. Others (p-value)				0.00
DLOs vs. ULOs (p-value)		0.00		0.00
Average size (sh.)	1607.28	473.51	3060.69 ***	3288.32 ***

Table VI
Likelihood of order execution

We study the likelihood of execution of hidden limit orders (HLOs) in the NSE conditional on trader type. We distinguish between high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs). To model order execution likelihood, we use a Pooled Ordered Logit Model with stock and day fixed effects, and White-robust standard errors. The dependent variable (EXEC) is an ordinal variable that takes three possible values: EXEC = 1 indicates that the limit order is cancelled before execution; EXEC = 2 indicates that the limit order is partially executed and then cancelled; EXEC = 3 indicates that the limit order is fully executed. In Panel A, we report the estimated coefficients for all orders, and buy and sell orders separately. In Panel B, we report the predicted probabilities of full execution for each type of trader derived from the estimated coefficients of the model. The estimation sample consists of the 30 largest stocks (in which HFTs are reasonably active) from our main sample of 100 stocks listed on the NSE. We drop fleeting orders and orders placed more than 20 ticks away from the prevailing best quotes. The sample period is December 2013. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively. In the Appendix, we provide detailed definitions for the explanatory variables.

Panel A: Estimated coefficients

Variable	All orders		Buy orders		Sell Orders	
HFT	-1.06	***	-1.03	***	-1.09	***
AAT	-0.68	***	-0.72	***	-0.64	***
HLO	-0.50	***	-0.54	***	-0.47	***
HLOHFT	0.98	***	0.96	***	0.98	***
HLOAAT	0.66	***	0.69	***	0.63	***
Aggr	158.48	***	207.06	***	133.28	***
OrdSize	-0.05	***	-0.06	***	-0.04	***
Rsprd	2.09	***	2.37	***	1.98	***
DepthSame	-1.74	***	-1.08	**	-4.13	***
DepthOpp	2.37	***	5.24	***	0.84	***
LOBImb	-0.19	***	-0.18	***	-0.19	***
FirstHalfHour	0.11	***	0.12	***	0.11	***
LastHalfHour	0.13	***	0.12	***	0.13	***
OI	-0.09	***	-0.08	***	-0.07	***
TrdFreq	0.22	***	0.20	***	0.21	***
Mom	-1.28	***	5.44	***	-6.65	***
Volat	7.97	***	6.29	***	11.14	***
Obs.	9616454		4746610		4869844	

Panel B: Predicted probabilities of full execution for HLOs

Dummy	Prob.		Prob.		Prob.	
NATs	0.41	***	0.39	***	0.43	***
AATs	0.76	***	0.76	***	0.76	***
HFTs	0.87	***	0.86	***	0.88	***

Table VII
Time to completion: Survival analysis

We study the determinants of the time to full execution of non-marketable limit orders at the NSE. We exclude market and marketable limit orders, fleeting orders, and orders placed more than 20 ticks away from the prevailing best quotes. Revisions of non-executed orders are treated as the same order; revisions of partially-executed orders are treated as new submissions. In Panel A we report the estimated parameters of an econometric model of time-to-completion using survival analysis, as in Lo, et al. (2002). The model describes an accelerated failure time specification of limit order execution times under the generalized gamma distribution. The model is estimated on a stock-by-stock basis, and we report aggregated coefficients and significance levels based on Chordia, Roll, and Subrahmanyam (2005). The estimation sample for this table consists of the 30 largest stocks (in which HFTs are reasonably active) from our main sample of 100 stocks listed on the NSE; the sample period is December 2013. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively. In the Appendix, we provide detailed definitions for the explanatory variables.

Panel A: Survival analysis estimates

Variable	All		Buy orders		Sell orders	
HFT	0.82	***	0.82	***	0.85	***
AAT	0.14		0.18		0.12	
HLO	1.98	***	2.02	***	1.88	***
HLOHFT	-2.71	***	-2.97	***	-2.55	***
HLOAAT	-1.55	***	-1.59	***	-1.40	***
Aggr	-21.45	***	-21.57	***	-21.02	***
Last Buy	-0.04		0.06		-0.13	
DepthSame	352.96	***	420.27	**	415.97	***
DepthSame^2	-305.23		-579.15		-401.99	*
DepthOpp	-236.04	***	-264.55	***	-293.93	***
OrdSize	10.34	**	10.69	**	12.17	***
TrdFreq	-11.80	**	-12.84	*	-11.83	*
RelTrdFreq	-1.60	***	-1.61	***	-1.60	***
Intercept	14.67	***	14.56	***	14.71	***
Obs.	9617067		4801022		4931201	

Panel B: Predicted average time to completion of HLOs

NATs	16.65		16.58		16.58	
AATs	15.24	***	15.16	***	15.30	***
HFTs	14.77	***	14.42	***	14.88	***

Table VIII
Implementation shortfall analysis

We evaluate the effective costs of execution (EFC) and the opportunity costs of non-execution (OPC) of hidden limit orders (HLOs) and displayed limit orders (DLOs) using the implementation shortfall (IS) approach of Perold (1988). For a buy order, EFC is the difference between the average execution price and the mid-quote at the time of order submission, multiplied by the amount of shares traded; OPC is the difference between the closing price on the day the order is cancelled or expires and the quote midpoint at the time the order is submitted, multiplied by the unexecuted part of the order (in shares). Metrics for sell orders are analogously computed but conveniently signed. We estimate pooled regression models with stock and day fixed effects and White-robust standard errors for the whole IS, but also for each component separately. In Panel A we report the estimated coefficients. In Panel B, we provide average estimated IS, EFC and OPC for each trader type (HFTs, AATs and NATs) based on the parameter estimates. Note that a fully executed order has zero opportunity cost, and a fully cancelled order has zero execution cost. So, for the EFC component we provide results conditional on partial execution (fill rate > 0); for the OPC component, we provide results conditional on non-full execution (fill rate < 100%). The estimation sample consists of the 30 largest stocks from our main sample of 100 stocks listed on the NSE; the sample period is December 2013. We consider only non-marketable limit orders and drop fleeting orders and orders placed more than 20 ticks away from the best quotes. Revisions of non-executed orders are treated as the same order. Revisions of partially-executed orders are treated as new submissions. For variables definitions, see the Appendix.

Panel A: Estimated coefficients

	IS	EFC (of execution)		OPC (of non-execution)	
Coef*100	All fill rates	All fill rates	Fill rate > 0	All fill rates	Fill rate < 100%
HFT	4.33 ***	3.43 ***	2.04 ***	0.08	-19.77 ***
AAT	2.38 ***	4.75 ***	2.78 ***	-3.13 ***	-29.74 ***
HLO	9.96 ***	0.44 *	0.45	10.94 ***	-1.83
HLOHFT	-20.58 ***	2.07 ***	4.70 ***	-25.01 ***	-22.59 ***
HLOAAT	-7.45 ***	-1.16 ***	-2.01 ***	-7.66 ***	12.61 **
Aggr	16.07 ***	-0.48	124.37 ***	5.75	84.50 ***
OrdSize	0.07 *	-0.08 ***	-0.09 ***	0.21 ***	0.77 ***
Buy	-8.24 ***	0.64 ***	-0.11	-9.28 ***	-24.32 ***
TrdFreq	-1.25	-5.60 ***	-5.44 ***	4.31 *	21.18 ***
Volat	-0.32	0.05	2.82 ***	-0.54	4.42
FirstHalfHour	0.36	-0.41 ***	0.63 ***	0.85	2.46
LastHalfHour	3.14 ***	1.14 ***	-0.01	1.45 *	-2.20
Intercept	8.71 ***	-6.50 ***	2.00 ***	15.48 ***	62.52 ***
Obs.	9732223	9732223	6601255	9732223	4184203

Panel B: Average estimated costs of a HLOs per trader type (*100)

	IS	EFC All fill rates	EFC Fill rate > 0	OPC All fill rates	OPC Fill rate < 100%
NATs	18.68	-6.05	2.45	26.42	60.69
AATs	13.60	-2.47	3.23	15.62	43.56
HFTs	2.43	-0.56	9.19	1.49	18.33

Table IX
Informativeness

We provide stock-day average impulse response functions (IRF) from an extended structural VAR model in the spirit of Hasbrouck (1991a). The model is defined in event time (*t*), where an event can be a limit order submission, cancellation, or trade. Revisions that improve (degrade) prices or increase (decrease) quoted depth are treated as limit order submissions (cancellations). Aggressive ("a") non-marketable limit orders improve or hit the prevailing best quotes; otherwise, they are non-aggressive ("na"). We differentiate between hidden (HLOs) and displayed limit orders (DLOs) and consider three trader types: high frequency traders (HFTs), agency algorithmic traders (AATs), and non-algorithmic traders (NATs). These partitions produce a VAR model with 18 equations: one for the quote midpoint return (in bps) and 17 for order-flow related variables (6 types of events x 3 types of traders -1) – we drop the HLOna category for HFTs because non-aggressive HFTs' HLOs are not frequent enough. The model is estimated for each stock and day. The optimal number of lags is determined using the Schwarz' Bayesian Information Criterion. "Trade" variables are signed +1 (-1) for buyer- (seller-) initiated trades. "DLO", "HLO" or "Cancellation" variables that happen on the ask (bid) side of the LOB are signed (-1) +1. We assume the trading process restarts each day, resetting all lagged values to zero. We report average estimates across stock-day observations clustered by both stock and day (Thompson, 2011). ***, **, * indicate significance at the 1%, 5% and 10% level, respectively. We boldface those impacts for AATs and NATs that are significantly different from corresponding impact for HFTs. We use order level data for December 2013 on the 30 largest stocks in our representative sample of 100 NSE-listed stocks.

		Trader type		
Message	All traders	HFT	AAT	NAT
Trades		1.16 ***	**0.73** ***	**0.86** ***
		(0.13)	(0.10)	(0.15)
DLOa		0.25 ***	0.24 ***	**0.62** ***
		(0.05)	(0.03)	(0.07)
DLOna		0.01 *	0.00	0.00
		(0.01)	(0.01)	(0.00)
HLOa		0.18	**0.35** ***	**0.49** ***
		(0.11)	(0.04)	(0.05)
HLOna			-0.04	-0.02 **
			(0.02)	(0.01)
Cancellations		0.06 ***	0.05 ***	**0.12** ***
		(0.02)	(0.01)	(0.02)

Table X
Information shares

This table reports the average stock-day information shares for different types of traders and orders in the NSE. Information shares are estimated using the Hasbrouck (1995) approach. We report lower bound (minimum), upper bound (maximum), and average information shares for three types of traders: HFTs, AATs, and NATs and we distinguish between hidden limit orders (HLOs) and fully displayed limit orders (DLOs). At a one-second frequency, we obtain the best quotes for each trader type and order category. The price path of each trader type and order category pair is given by the quote midpoint prevailing at the end of each second. Using the IS approach, we decompose the variation in the unobserved common efficient price into individual components attributable to specific trader type and order category pair. Our purpose is to examine the fraction of price discovery attributable to HLOs and how much of comes from HFTs' and ATs' orders. We use order level data for December 2013 on the 30 largest stocks in our representative sample of 100 NSE-listed stocks. ***, **, * next to a HFTs' or ATs' IS indicates that the IS statistic is significantly different from the corresponding NATs' IS statistic for the same order type.

		Information shares (%)		
Trader type	Order	Min.	Max.	Avg.
HFTs	DLO	15.87	45.83	30.85
	HLO	5.91	6.34 ***	6.13 **
AATs	DLO	8.81 ***	34.44 ***	21.62 ***
	HLO	5.00	10.25 ***	7.62 **
NATs	DLO	16.22	47.62	31.92
	HLO	6.36	17.39	11.87

Table XI
Expected profits of market making and the exposure decision

We study the determinants of the order non-exposure decision of high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs) as a function of expected profits of market making, as measured by the average realized spread of recent trades. For a limit order to sell (buy), *ERzdSpr* is the average of the realized spread for the five most recent buyer-initiated (seller-initiated) trades executed at least 1-second earlier. The 1-second time constraint is imposed because realized spreads are computed using quote midpoints 1-second after trade execution. We use logistic models (Panel A) of order characteristics and market conditions to study the choice between submitting a hidden limit order (HLO) and a fully displayed limit order (DLO). We exclude all market and marketable limit orders, fleeting orders, and orders placed beyond 20 ticks from the prevailing best quotes. The dependent variable equals one (zero) if the trader submits a HLO (DLO). We use Tobit models (Panel B) of order characteristics and market conditions to study the decision of how much of a limit order to hide. The dependent variable here is the amount of shares hidden, normalized by the stock's average daily trading volume. The models are estimated on a stock-by-stock basis, and we report aggregated coefficients and t-statistics using the approach in Chordia, Roll, and Subrahmanyam (2005). The estimation sample for this table consists of the 30 largest stocks (in which HFTs are reasonably active) from our main sample of 100 stocks listed on the NSE. The sample period is December 2013. ***, **, * indicate significance at the 1%, 5%, and 10% level respectively. In the Appendix, we provide detailed definitions for the explanatory variables.

Panel A: Decision to hide - logistic model

Coeff.	HFTs		AATs		NATs	
Aggr	2233.73	***	970.30	***	33.27	
OrdSize	31.78	**	29.40	***	21.94	***
ERzdSpr	309.37	***	73.04	**	-53.81	**
DepthSame	-583.58	***	-241.70	***	-135.48	***
DepthOpp	30.10		45.34	**	-52.03	**
Volat	-1.22		-0.12		-0.69	**
WaitTime	-50.17	*	18.08		18.54	**
TrdFreq	-1.84		-0.61		-0.94	**
HidVolSame	-4.12		0.08		-0.41	*
LOBImb	14.81		0.31		-0.54	
LastTrdSize	-1.31	**	-1.71	**	-0.81	**
MVolat	0.00		0.00		0.00	
LastHalfHour	553.31	***	54.14		-259.31	***
Intercept	-3.74	***	-0.61		-1.76	***

Panel B: Magnitude of hidden volume - Tobit model

Coeff.x100	HFTs		AATs		NATs	
Aggr	23.77	***	12.30	***	-7.89	
OrdSize	0.32	**	0.63	***	0.55	***
ERzdSpr	6.05	***	0.25		-2.21	
DepthSame	-5.33	***	-1.82	***	-3.03	**
DepthOpp	0.09		0.32		-2.49	***
Volat	0.00		0.00		-0.05	***
WaitTime	-1.26		0.06		0.64	***
TrdFreq	-0.05		0.05		-0.05	
HidVolSame	-0.10		0.00		-0.02	
LOBImb	0.18		0.08		0.10	
LastTrdSize	-0.04		-0.01	***	-0.01	
MVolat	0.00		0.00		0.00	
LastHalfHour	4.84		4.79		-10.15	
Intercept	-0.12	***	-0.03	***	-0.26	***

Table XII
Undercutting using Hidden Limit Orders

We present the proportions of hidden (HLOs) and displayed (DLOs) limit orders used for undercutting by the three trader types – HFTs, AATs, and NATs – (in Panel A) and use a logit regression model to study the likelihood of undercutting by these three trader types (in Panel B). We define an undercutting limit order as a limit order that (a) is placed immediately after another submission on the same side of the market, (b) comes in under 10 milliseconds of the previous order, and (c) improves the price of the previous order. We present results using undercutting orders restricted to the five best quotes. We scale the total number of undercutting orders of each type – hidden and displayed –by each trader type – HFT, AAT, and NAT – by all orders submitted of a given type by each trader type. We present those fractions in Panel A. In Panel B we present the coefficients of the logit regression where the dependent variable is a dummy that takes the value of 1 if the order is an undercutting order, 0 otherwise. We also report the odd ratios for the variables of interest. The models are estimated on a stock-by-stock basis and we aggregate coefficients and t-statistics using the approach in Chordia, Roll, and Subrahmanyam (2005). The estimation sample for this table consists of the 30 largest stocks (in which HFTs are reasonably active) from our main sample of 100 stocks listed on the NSE; the sample period is December 2013. ***, **, * indicate significance at the 1%, 5%, and 10% level respectively. Variables are defined in the Appendix.

Panel A: Descriptive statistics on undercutting (% of orders)

First case: At least 20 orders per category and stock-day

Order	TraderType	Bid side	Ask side
HLO	HFT	5.02 ***	5.41 ***
	AAT	3.23 ***	3.41 ***
	NAT	0.82 ***	0.81 ***
DLO	HFT	3.01 ***	3.24 ***
	AAT	4.63 ***	5.02 ***
	NAT	1.14 ***	1.17 ***

Second case: At least 50 orders per category and stock-day

Order	TraderType	Bid side	Ask side
HLO	HFT	5.60 ***	6.07 ***
	AAT	3.40 ***	3.48 ***
	NAT	0.81 ***	0.80 ***
DLO	HFT	2.60 ***	2.73 ***
	AAT	5.17 ***	5.58 ***
	NAT	1.06 ***	1.08 ***

Panel B: Logit model on undercutting

Variable	Coeff.	Odds ratio	CRS t-stat
HFT	0.76 ***	2.14	39.49
AAT	0.99 ***	2.68	40.69
HLOHFT	0.41 ***	1.51	7.67
HLOAAT	-0.19	0.83	-0.06
HLONAT	-0.56 ***	0.57	-3.96
DispSizeUnd	0.00 ***	---	10.03
AggrUnd	-0.07 ***	---	-119.14
HidVolSame	0.45 ***	---	66.72
Rsprd	0.03 ***	---	39.78
DepthSame	0.38 ***	---	10.71
DepthOpp	-0.95 ***	---	-9.27
Volat	0.01 ***	---	22.57
Intercept	-4.07 ***	0.02	-183.04

Table XIII
Realized spreads

We use effective and realized spreads to evaluate the profitability of hidden limit orders (HLOs) and displayed limit orders (DLOs) submitted by different types of traders. We use trade and quote data for the 30 largest stocks from our representative sample of 100 NSE-listed stocks. Effective and realized spreads are obtained for every standing limit order at the prevailing best quotes executed during December 2013. For a buyer initiated trade, the effective spread is two times the difference between the trade price and the quote midpoint prevailing before the trade; the realized spread is two times the difference between the quote midpoint *h* seconds after a trade and the trade price. For a seller-initiated trade, we multiply the above magnitudes by -1. For the realized spread, we report results with $h = \{1, 5, 30\}$. We express each metric relative to the quote midpoint prevailing before the trade and, then, in basis points. We regress the realized spread on passive trader type (HFT and AAT) indicators, passive order type (HLO) indicators, trade characteristics (size, duration), LOB characteristics (depth imbalance, depth on the same side), market conditions over the last minute (volume, volatility, order imbalance), and time of day dummies. See the Appendix for detailed definitions of all the variables. We use pooled regression models with stock and day fixed effects and White-robust standard errors. In addition to the estimated coefficients, we also report Wald tests on specific null hypotheses.

	Effective spreads	Realized spreads 1s	Realized spreads 5s	Realized spreads 30s
HFT	1.40 ***	1.86 ***	2.54 ***	3.15 ***
AAT	0.23 ***	0.51 ***	1.00 ***	1.46 ***
HLO	0.17 ***	0.33 ***	0.51 ***	1.18 ***
HLOHFT	-0.70 ***	-1.01 ***	-1.23 ***	-1.34 ***
HLOAAT	-0.50 ***	-0.21 ***	-0.57 ***	-1.42 ***
LOFirst	0.24 ***	-0.16 ***	-0.54 ***	-0.62 ***
RelSize	-0.02 ***	-1.25 ***	-1.31 ***	-1.41 ***
TDur	0.16 ***	0.35 ***	0.31 ***	0.07 ***
TDurSame	-0.19 ***	-0.09 ***	0.17 ***	0.47 ***
HFTtaker	-0.25 ***	-2.13 ***	-2.90 ***	-3.48 ***
AATtaker	-0.06 ***	-0.07 ***	-0.15 ***	-0.36 ***
LOBImb	-0.06 ***	0.40 ***	0.74 ***	0.97 ***
DepthSame	0.11 ***	0.11 ***	-0.05 ***	-0.22 ***
Vol	-0.26 ***	-0.31 ***	-0.37 ***	-0.28 ***
Volat	6.54 ***	3.16 ***	3.20 ***	0.90 ***
OI	0.04 ***	-0.30 ***	-0.88 ***	-1.42 ***
HVolOpp	-0.34 ***	-0.58 ***	-1.00 ***	-1.44 ***
FirstHalfHour	0.13 ***	0.12 ***	0.03 **	-0.14 ***
LastHalfHour	0.08 ***	0.46 ***	0.71 ***	0.72 ***
Intercept	6.10 ***	5.43 ***	6.38 ***	6.05 ***
Obs./1000	11743	11742	11740	11701
Adj.R^2	0.67	0.15	0.08	0.03
Wald tests				
HFTs' DLOs vs HFTs' HLOs:				
H0: HLO + HLOHFT = 0	0.53 ***	0.69 ***	0.72 ***	0.16
HFTs' HLOs vs NATs' HLOs				
H0: HFT + HLOHFT = 0	0.71 ***	0.85 ***	1.31 ***	1.80 ***
HFTs' HLOs vs AATs' HLOs				
H0: HFT + HLOHFT = AAT + HLOAAT	0.98 ***	0.55 ***	0.88 ***	1.77 ***
HFTs' DLOs vs NATs' DLOs				
H0: HFT = 0	1.40 ***	1.86 ***	2.54 ***	3.15 ***
HFTs' DLOs vs AATs' DLOs				
H0: HFT = AAT	1.18 ***	1.36 ***	1.54 ***	1.69 ***

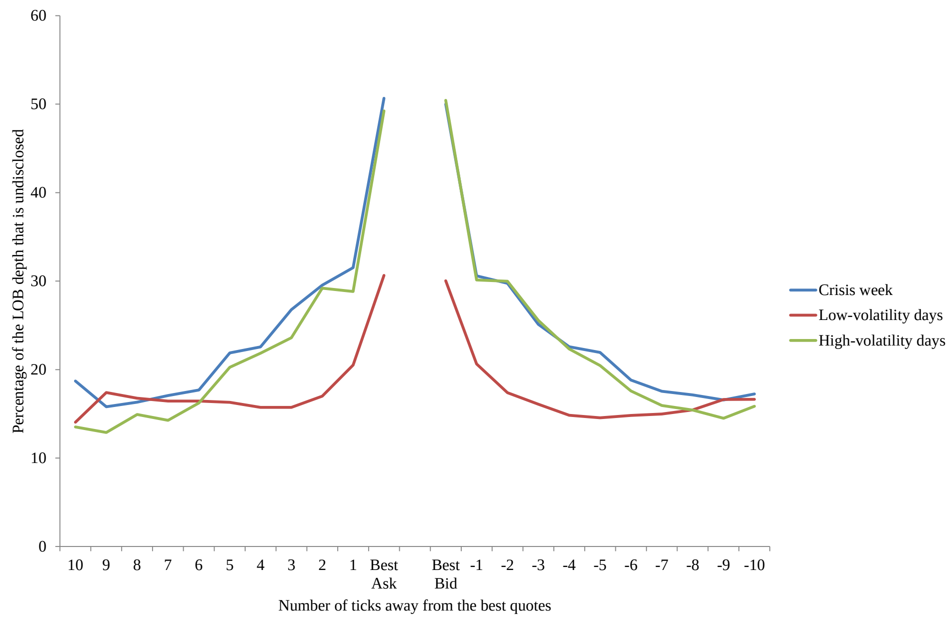

(a) Large Cap stocks

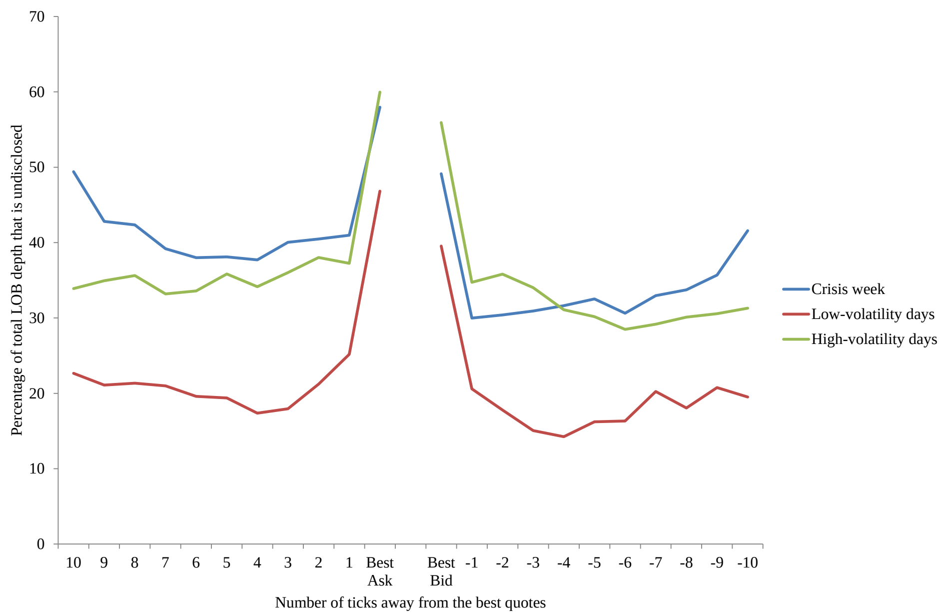

(b) Small Cap stocks

Figure 1. Hidden depth in the Nasdaq Limit order book

We plot the cross-sectional average percentage of the hidden depth in the Nasdaq limit order book for the "crisis week" of September 2008, and five days (per stock) with lowest and highest volatility. We show results here for the large and small cap samples. We use order by order data collected from one-minute snapshots of the ten best ask and bid order book levels.

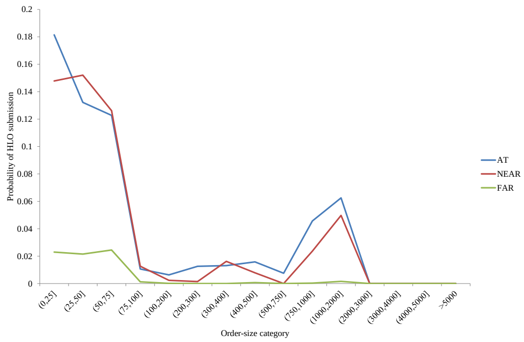

2.a. HFTs’ probability of HLO submission

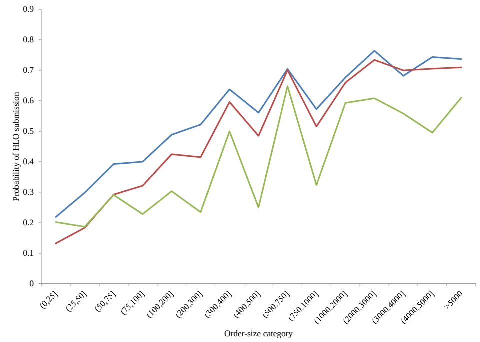

2.b. AATs’ probability of HLO submission

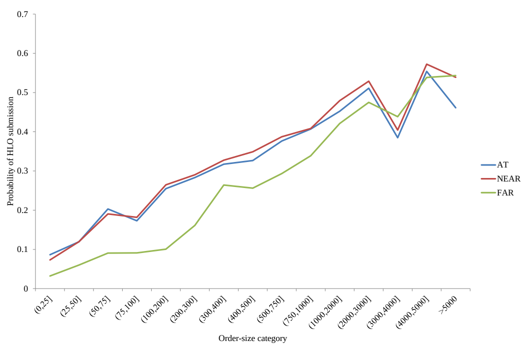

2.c. NATs’ probability of HLO submission

Figure 2. Probability of hidden order submission conditional on order size and aggressiveness (NSE)

We plot estimated cross-sectional daily average probabilities of hidden limit order (HLO) submission in the NSE conditional on order size and order aggressiveness. We distinguish between high frequency traders (HFTs), agency algorithmic traders (AATs) and non-algorithmic traders (NATs). The sample consists of the 30 largest stocks from our size-stratified sample of 100 stocks listed on the NSE between October and December 2013. We combine the limit order book levels into three groups: at the best quotes (“At”); from the best quotes up to 5 ticks away (“Near”), and the rest (“Far”). For each order size, level of aggressiveness, and type of trader, we provide the percentage of HLOs of all non-marketable limit orders submitted. Figures 2.a, 2.b, and 2.c provide the findings for HFTs, AATs, and NATs respectively.

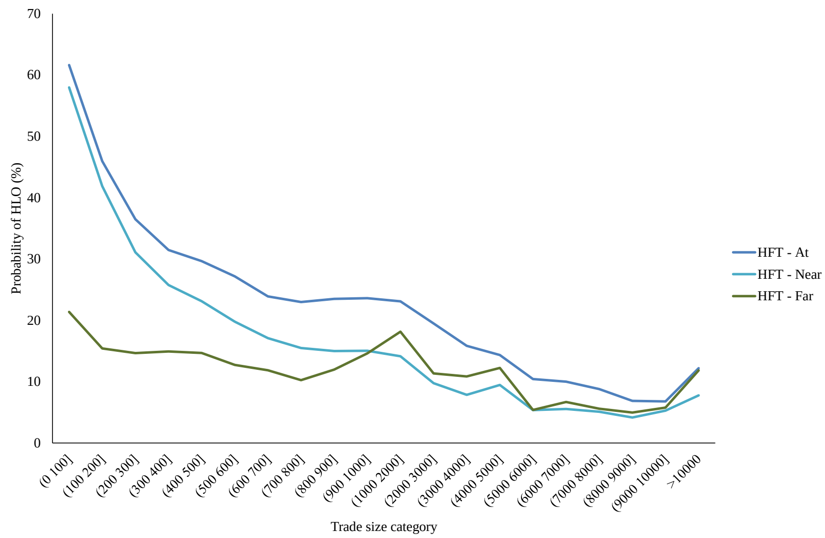

3a. HFTs – Large Caps

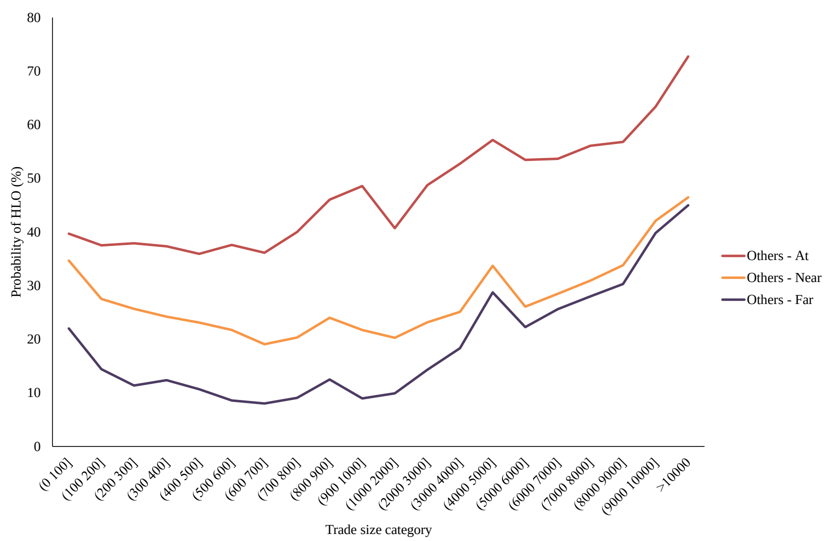

3b. non-HFTs – Large Caps

Figure 3. Probability of hidden order submission conditional on order size and aggressiveness (Nasdaq)

We plot cross-sectional daily average probability of hidden limit order (HLO) submission conditional on order size and aggressiveness. We distinguish between high-frequency traders (HFTs) and non-HFTs ("Others"). The sample period consists of 50 non-consecutive days from 2008 to 2010. We exclude the "crisis week" of September 2008. The sample consists of 120 Nasdaq-listed firms split into three equally-sized subsamples: large, mid-sized (not reported), and small caps. We use order by order data collected from one-minute snapshots of the ten best ask and bid LOB levels. We define three levels of aggressiveness: (a) at the best quotes ("At"); (b) up to five ticks away from the best quotes ("Near"), 6 ticks away or more from the best quote ("Far"). We consider order sizes from 1 to 1,000 shares in increments of one round lot (i.e., (0, 100], (100, 200] … (900 1,000]), from 1,000 to 10,000 shares in increments of 10 round lots (i.e., (1,000 2,000], (2,000 3,000] … (9,000 10,000]), and orders of size greater than 10,000 shares.

Appendix
Variable definitions

DepthSame	Prevailing displayed depth at the best bid (ask) for an incoming buy (sell) order divided by the average daily trading volume
DispSizeUnd	Displayed size of the order being undercut
ERzdSpr	For a submission of a non-marketable limit order to sell (buy), average realized spread over the last 5 buyer-initiated (seller-initiated) trades that are at least one second old
FirstHalfHour	Indicator variable that equals 1 for orders submitted in the first half hour of the trading day, 0 otherwise
HFT	Indicator variable that equals 1 for orders submitted by HFT, 0 otherwise
HFTtaker	Indicator variable that equals 1 when a trade is initiated by a HFT, 0 otherwise
HidVolOpp	Indicator variable that equals 1 if the presence of hidden volume on the opposite side is revealed, 0 otherwise
HidVolSame	Indicator variable that equals 1 if the presence of hidden volume on the same side is detected, 0 otherwise
HLO	Indicator variable that equals 1 for hidden orders, 0 otherwise
LastBuy	Indicator variable that equals 1 if the last trade is buyer initiated, 0 otherwise
LastHalfHour	Indicator variable that equals 1 for orders submitted in the last half hour of the trading day, 0 otherwise
LastTradeSize	Size of the last trade executed in number of shares
LOBImb	Difference between the accumulated displayed depth at the best five bid and ask quotes and divided by the total accumulated depth, suitably signed (positive when same side depth exceeds opposite side depth)
LOFirst	Indicator variable that equals 1 if a standing limit order is executed for the first time
Mom	Continuously compound quote midpoint return in the last 5 minutes
MVolat	Sum of the squared continuously compounded return of the NIFTY-50 over the last 60 minutes
NAT	Indicator variable that equals 1 for orders submitted by NATs and 0 otherwise
OI	Buyer-initiated volume minus seller-initiated volume divided by total volume in last k minutes, where k = 1 in table XIII and k = 5 in Table VI
OrdSize	Total (displayed plus hidden) size of the order divided by the average daily trading volume
RelSize	Full trade size of an incoming market order over the prevailing depth at the best quote of the side it hits
RelTrdFreq	Number of shares traded in the last 30 minutes divided by number of shares traded in the last 60 minutes
Rsprd	Bid-ask spread divided by the quote midpoint
TDur	Time from the previous trade in seconds
TDurSame	Time from the previous trade with the same direction in seconds
TrdFreq	Number of shares traded per second within the last k minutes, where k = 5 in Table VI and Table VIII and k = 60 in Table VII and Table XI
Vol	Number of shares traded in the last minute
Volat	Sum of the squared continuously compound quote-midpoint return over the last k minutes, where k = 1 in Table XII and Table XIII, k = 5 in Table VI and Table VIII and k = 60 in Table XI
WaitTime	Average time between the last three message arrivals on the same side, resetting the clock daily

* The coefficients are sometimes multiplied by 10^n, where n varies as required, for representation purposes (this has no effect on the significances and the results)