

*“Modern” Market Makers**

Katya Malinova[†]

Andreas Park[‡]

March 21, 2016

Abstract

Using proprietary, trader-level data, we study the order submission and cancellation behavior of high-frequency market makers. Studying a multi-market setting enables us to provide novel evidence for the existence of the so-called quote-fade phenomenon (quotes disappear market-wide immediately after orders) and latency arbitrage by high frequency market makers, and we identify the intra-day determinants of the phenomena. Using an event that eliminated latency between two of the three main markets, we find that reductions in latency exacerbate quote-fade and latency arbitrage. As market makers accumulate inventories, they post on average more conservative prices, and at the same time, they post more orders that are aggressively priced, presumably to trade out of these inventories. As trading in the market becomes one-directional, market makers post fewer orders against the market. High frequency market makers thus temporarily improve posted bid- and ask prices, even though they generally do not lean against the order flow.

*We thank IIROC and in particular Victoria Pinnington and Helen Hogarth. Large parts of this work were done while we visited Copenhagen Business School, and we thank CBS for its hospitality. Support from the Center for Financial Frictions (FRIC), grant no. DNRFF102, and from the Global Risk Institute for the project “Risk Management and Market Liquidity” is gratefully acknowledged. This report is based on a research report that we prepared for the Investment Industry Regulatory Organization of Canada titled “Liquidity Provision and Market Making by HFTs”.

[†]University of Toronto, Department of Economics, katya.malinova@utoronto.ca.

[‡]University of Toronto Mississauga, Department of Management and Rotman School of Management, andreas.park@rotman.utoronto.ca.

In today’s equity markets most trades involve autonomously operating computerized traders on at least one side of the trade.¹ These traders are often referred to using the umbrella-expression of “high-frequency traders”, and much work has been dedicated in the literature to understanding the impact of these new, autonomous traders on markets.² The impact of trading by these algorithmic traders is, by now, well-understood in single-market environments,³ and more recently several studies described their trading behavior in multi-market settings.⁴ The *order* submission behaviour, is, however, less-well understood, particularly in multi-market settings.⁵ In this paper, we fill the gap and study order submissions across different equity trading venues.

For this study, we were granted access to a proprietary data set that contains (masked) trader-level information on all trades, quotes, and orders for all 11 Canadian equity trading venues. Our particular focus is on traders that engage in market-making behavior so as to understand the differences between “modern” and, traditional, widely studied markets, such as the old-style NYSE with specialists. Apart from the computerization and automation of trading, the most critical difference between modern and traditional market environments is the ability to trade the same security simultaneously on multiple, competing, and electronically linked venues.

We classify traders as voluntary, *de facto* market makers if they persistently post similar volumes of non-marketable limit orders on both sides of the market across many securities; in the appendix we argue that most of these market makers are also “fast”. Our analysis then proceeds in four steps. First, we study how these fast market makers react to trades. Second, we analyze how their reaction to trades changes after a major market structure change eliminated the latency between two of the main markets. Third, we study the intra-day dynamics of these traders’ order submissions in response to market-wide demand- and supply pressures. Finally, we analyze whether and how these traders manage their inventories and how their inventory management affects markets.

¹See Jones (2013) and references therein.

²In trying to circumvent the overused and indiscriminate term “high frequency trader”, we use “Modern” in the title of the paper. The choice was inspired by name of the HFT community’s political lobby group, the so-called “Modern Market Initiative”.

³See, e.g., Menkveld (2013), Hagströmer and Norden (2013), or Brogaard, Hendershott, and Riordan (2014a).

⁴See, e.g., Brogaard, Hendershott, and Riordan (2014b), Korajczyk and Murphy (2014) and Boehmer, Li, and Saar (2015).

⁵See Malinova, Park, and Riordan (2013) or Subrahmanyam and Zheng (2015) for single markets.

Our work is in the tradition of Biais, Hillion, and Spatt (1995): exploiting very detailed data, we describe in detail how market makers post orders across multiple markets.

We classify any trader as a (voluntary) market maker if this trader regularly posts similar passive volume on both sides of the market on many days and across many securities. As a first step in our analysis, we study the behaviour of market makers subsequent to trades. Theoretical models of market making predict⁶ that market makers adjust their quotes subsequent to trades, as (a) these trades may reveal information about fundamentals and (b) taking an inventory exposes the market maker to risk. van Kervel (2015) models market-making across multiple venues. In his model, market makers trade-off a higher execution probability with increased adverse selection. We empirically study such a multi-market setting and describe how market makers adjust quotes across markets in the first few milliseconds following a trade. Similarly, we study whether market making traders submit aggressive orders, e.g., to “take out” other traders’ stale quotes.

We perform this part of the analysis for a smaller subset of securities, specifically, highly liquid, non-crosslisted securities that are in the TSX60 index. We focus on two measures: cancellations in the opposite direction of trades (e.g., cancellations of sell-orders following a buy)⁷ and aggressive order submissions in the same direction as the trade (e.g., marketable buy orders following a buy). Aggregating across all trades, we observe a declining number of cancellations and aggressive submissions per millisecond by market makers (HFTMMs) following the trade. To the naked eye, which can at best observe market movements after 250ms, the quotes would appear to have been cancelled concurrently with the trade. The total number of occurrences is not large though: only about 3.4% of trades are followed by aggressive HFTMM orders, and only 17.6% of trades are followed by HFTMM cancellations.

A formal regression analysis confirms the observations from the aggregate data. In this part of the analysis, we additionally address two questions. First, we ask which factors contribute to the probability of observing HFTMM aggressive orders (“latency-arbitrage”) or cancellations (“quote-fade”). Second, we study how a major technological

⁶See, for instance, Kyle (1985), Glosten and Milgrom (1985), Glosten (1994), Biais (1993).

⁷Theoretically, when moving a quote, a market maker would cancel orders on both sides of the market resubmit orders on both sides at different prices. This part of our analysis is both computationally intensive and difficult to present concisely when discussing both same-side and opposite-side cancellations. To simplify the exposition, we focus only on opposite-side cancellations.

change, which all but eliminated the latency between two markets, affected the occurrences of HFTMM aggressive orders and cancellations. The contributing factors for quote-fade and latency arbitrage are similar: both phenomena are more likely to occur when the trade absorbs the entire local depth, when the trade is larger, and, when restricting attention to trades by “directional” traders, when the trader has already traded a lot. When the liquidity was supplied by an HFTMM, the probability of observing an aggressive order is lower, but the chance for a cancellation is higher. Smaller bid-ask spreads and trades that go against the recent return trend increase the probability for both aggressive orders and cancellations. After the market structure change which eliminated latency between two of the three major markets, we observe a significant increase in the occurrences of trades followed by aggressive orders and cancellation for the smaller of the two “merged” marketplaces.

In the third part of our analysis, we study the intra-day dynamics of order submission behavior of market makers. For this part of the analysis we split the day into volume intervals. Each interval contains the orders, cancellations and trades that occur while 1% of average daily volume is being traded (excluding block trades). We first study how market makers change their order submission behaviour as the market demand (or supply) becomes unbalanced in the sense that there are, for instance, many more buyer- than seller-initiated trades.⁸ We are particularly interested in the impact of trade imbalance on the (im-)balance of market maker order submissions, on the average prices at which market makers post their buy and sell orders, and on the location of market makers’ new orders in the order book. We find that as the trade imbalance increases, market makers post fewer orders against the direction of the market and they post less aggressive prices. Moreover, they post more conservatively in the sense that they post their orders away from best prices — both in and against the direction of the trade imbalance.

In the forth part of our analysis, we study the dynamic relationship of market makers’ order submissions and their inventories. We first observe that market makers’ inventories and trade imbalances are negatively related, implying that indeed market makers “lean against the wind” and take the other side of the trade imbalance. We then observe that as market makers accumulate inventories, they reduce order submissions in the direction

⁸The imbalance loosely relates to Easley, López de Prado, and O’Hara (2012)’s measure of order flow toxicity.

of their inventories (e.g., if they are long, they post fewer buy orders). They also post on average lower-priced buy orders and higher-priced sell orders as they accumulate long and short positions, respectively.

As a last step we study how the changes in the market makers' posting behavior affect liquidity and volatility. Here we observe that for the stocks with the most competitive liquidity provision, changes in the market makers' posting behavior in response to changes in their inventories are associated with improvements in liquidity. This finding contrasts findings in the literature, e.g., in Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010), where NYSE specialists' inventory management lead to worsened liquidity. We also find, however, that the inventory management is associated with increased volatility.

Our paper contributes to the rapidly growing literature on high frequency trading. Subrahmanyam and Zheng (2015) study HFTs' order placements on NASDAQ; they find that HFTs use order cancellations in anticipation of short-term price movements. Wah (2016) studies multi-market quote-changes and trades using the SEC's MIDAS data. She quantifies the aggregate cost of latency arbitrage, which she defines as the occurrence of crossed markets (one venue's bid price exceeding another venue's ask price). Brogaard, Hendershott, and Riordan (2014b) study the multi-venue trading behavior of HFTs and their contribution to the price discovery process. Boehmer, Li, and Saar (2015) identify the correlation among high frequency traders' trading strategies across multiple markets, and they show that there is a negative relation between HFT activity and short-term stock volatility. Key features of our approach are that we study the millisecond-level interactions and reactions of high frequency market makers, and their order submission behavior (as opposed to their trading behavior) across multiple markets.

Our work also relates to the interaction of high frequency traders and large institutional orders. Using the same dataset as this paper, Korajczyk and Murphy (2014) find that HFT liquidity provision is significantly reduced for large trades and that such trades face higher bid-ask spreads. van Kervel and Menkveld (2015) study the interaction of HFTs and institutional orders on a single venue, using data on the so-called child orders from four institutional traders. Both Korajczyk and Murphy (2014) and van Kervel and Menkveld (2015) find that early in a string of institutional trades, HFTs provide liquidity, but for later portions of an institutional order, they trade in the same direction as the

institution. Focussing on the millisecond-level HFTMM reaction to trades, we find that the more dollar-volume a single-directional trader has already submitted, the stronger the market maker reaction; our results are thus consistent with these two papers.

In summary, our analysis indicates that after, say, trading with a buyer, market-makers cancel their sell orders quickly and may also submit aggressive buy orders. This latter behavior can be interpreted as market makers either trading in anticipation of future orders or taking advantage of and eliminating mis-priced, stale quotes. The more buys the buying trader had already submitted, the stronger this effect. Furthermore, market makers submit, for instance, fewer bids as the number of sellers in the market grows, and they post more conservatively on average as they accumulate inventories. At the same time, despite the more conservative average prices, they appear to be posting very aggressively priced orders as their inventories grow, presumably to trade out of these inventories, and they then temporarily improve posted liquidity. Overall, however, based on our analysis we conclude that there is little to no firm evidence that modern market makers truly lean against the order flow.

I. The Institutional Setting

A. Core rules governing trading in Canada

The Toronto Stock Exchange (TSX) is the primary listing venue for large companies in Canada, small and mid-cap companies are typically listed on the TSX Venture exchange. As in other major markets around the world, trading in TSX-listed stocks is fragmented across multiple exchanges and Alternative Trading Systems (ATS), and many TSX-listed companies are also listed on U.S. exchanges. Securities trading and the activities of market participants in Canada are regulated by the Investment Industry Regulatory Organization of Canada (IIROC), the members of the Canadian Securities Administrators, and are governed by the Universal Market Integrity Rules (UMIR).

Most of the core elements of the UMIR are similar to those governing trading in the U.S. equities markets. Brokers and marketplaces are required to respect the order protection rule, which mandates that orders must be routed to the marketplace with the best-priced orders available on lit markets. Brokers are also subject to obligations regarding best execution for client orders.

A significant volume of trades is pre-arranged off-exchange, before entering orders on a public marketplace. These trades must still be executed on a public marketplace, respecting all the applicable rules. There are usually very few such deals each day, but they are large and, on average, account for roughly 10-12% of trading value. We omit such trades from most of our analysis.

B. Marketplaces and their trading rules

The data in our sample contains observations for ten marketplaces. These marketplaces are separately, but anonymously identified in our data. For our study, we focus on the trades on the six “lit” marketplaces, and we label them alphabetically. During our sample period, three of the lit marketplaces together account for about 90% of the dollar volume traded, respectively.

Towards the end of our sample period (end-May 2013), a new marketplace started operating. The special feature of this new venue was the so-called “taker-maker” or “inverted” pricing schedule under which the liquidity taking side of a transaction would be paid a fee rebate and the liquidity making side would pay a fee. Trading in May in this market was sparse and to avoid any confounding effects, we omit this venue from the analysis.

C. A Major Technological Change

As of April 29, 2013, markets A, B and D moved to the same trading platform. This switch involved a number of changes (e.g., regarding the available order types), but the most important changes are the system integration and the physical move of market A’s servers (and, presumably, the colocated entities) to the data centre where markets B and D were located. Before this switch, there were three main locations, and the main markets A, B, and C were physically separated. After the change, markets A and B were at the same location and on the same system. As part of the move we would suspect that A’s systems were upgraded (and thus made faster) and thus, despite the longer physical distance from A to C, it is possible that the total latency between the venues did not change much. The main change is thus the much-reduced latency between A and B. We will facilitate this change to gain better insights into the importance of latency.

II. Data and Sample

Data. The data for this study is provided by the Investment Industry Regulatory Organization of Canada (IIROC). The dataset contains detailed records on all trades, orders, order cancellations, order amendments, and updates to marketplaces best bid and offer quotes from IIROC’s real-time surveillance system, for all trading on all regulated Canadian marketplaces. Each order-related record includes, in particular:

- The marketplace where the order was sent (masked).
- Size, price, and the direction (buy or sell) of an order.
- Broker ID (masked), user ID (masked), and account type (e.g., specialist, client, options-trader, or inventory).
- Other characteristics, including the duration of an order (for instance, good-till-cancel or immediate-or-cancel), whether an order was transparent or non-transparent, whether the order was a seek-dark-liquidity order, and a unique identifier for each order.

For trades, the data additionally specifies the aggressive and passive (liquidity-providing) side of a trade. The data also identifies the aforementioned intentional broker-crosses, which we omit from the analysis. The information for marketplaces, brokers and users is masked in the sense that IIROC provides a scrambled identifier. The masking is applied consistently so that the same marketplace, broker and user are always assigned the same identifier. Marketplaces time-stamps are reported with millisecond precision for our sample period. Brogaard, Hendershott, and Riordan (2014b), Korajczyk and Murphy (2014), Comerton-Forde, Malinova, and Park (2015) and Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014) contain further information of the data.

Many Canadian companies, in particular the large and frequently traded ones, are cross-listed with U.S. markets; for instance, of the 60 constituents of the S&P/TSX60 index, Canada’s large-cap index, more than 2/3 are also listed on U.S. exchanges, and around 50% of volume for these firms trades in the U.S. For our analysis, we instead focus on a sample of frequently traded, non-crosslisted securities, because for such securities, we know all the trades and orders, and we can reasonably assess the traders’ market-wide behavior in reactions to trades and quotes in the security.

To classify traders into different categories, however, we rely on a larger sample of securities because we believe that using a large sample enables us to capture general trading characteristics most accurately.

Classification Sample. We base our classification on the 307 securities that are classified as “highly-liquid” securities by IIROC during the entire sample period. Loosely, a security qualifies as highly-liquid for a given day if over a 60-day period it traded more than 1,000 times per trading day and had an average trading value of at least \$1M. IIROC compiles a list of highly-liquid securities daily; we include a security in our sample if that security is on the list of highly liquid securities at the end of each month in our sample period. We applied no further filters, in particular, there are no corrections for stock splits, corporate actions, halts, etc. For these securities, we consider the period from January 1 to May 31, 2013.⁹

Analysis Sample. For our analysis of order submission and cancellation behavior, we focus on the 17 constituents of the S&P/TSX60 index that are not cross-listed with U.S. markets during the months of March and May 2013. All of these 17 securities are also “highly liquid.” We determine the cross-listing status based on the June 2013 TSX e-review publication.

Outliers. We eliminated two days from our samples: January 21 (Martin Luther King Day) and May 27 (Memorial Day); these days are public holidays in the U.S., and trading activity on Canadian markets on such days is very low.

III. Market Maker Classification

All traders access the marketplaces via brokers. We base our classification on the analysis of order submission and trading behavior by trader IDs, where we define a trader ID as the combination of broker ID plus user ID, plus the account type (client, specialist, inventory, option market maker, and non-client). The user ID is the most granular identification that is available to regulators in Canada; IIROC researchers describe the

⁹We end the sample at the end of May for two reasons. First, at the beginning of June, a large number of high-activity trader IDs disappeared. At the same time, several new high activity IDs appeared (for the same brokers), and the time horizon of the data is too short to reliably classify many of the new IDs. Second, IIROC’s public market share statistics illustrate that a new marketplace rapidly gained market share. The entry of this marketplace has been associated with changes in behavior that we might not be able to fully capture as our data ends too early (end June 2013).

usage of user IDs in detail in their research reports (IIROC (2012), Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014), and Devani, Anderson, and Zhang (2015)); we provide further details in the appendix. We classify traders based on order submission behavior in the classification sample; in this sample, we found a total of 4,892 distinct trader IDs. Our general approach is to provide a classification based on a trader’s behaviour across a large number of securities.

A defining feature of a market maker is that the trader posts on both sides of the market and stands ready to trade. We thus expect that a market-making trader would submit passive buy and sell orders on both sides of the market so as to earn the bid-ask spread (and maker-taker fees) on as many trades as possible. We then compute for each trader, day, and security, the market maker index defined as

$$\text{market maker index} = \left| \frac{\text{passive buy order volume} - \text{passive sell order volume}}{\text{passive buy order volume} + \text{passive sell order volume}} \right|. \quad (1)$$

By construction, this index is between 0 and 1, where an index close to 0 indicates that the trader submits about as much as sell order volume. A trader’s market making index is the median index over all days and securities.¹⁰

Upon visually examining the classification data, it is apparent that there is a structural break for an imbalance score of 0.2 for our sample. We thus classify a trader ID as market-making if the trader ID has an imbalance score below 0.2.¹¹ We identify a total of 138 IDs as market makers. Of these, 94 IDs trade on the passive side in our Analysis Sample, and they provide liquidity for 45% of the transactions and 30% of dollar value.

In the appendix, we discuss the classification further, namely, we describe that about half of the market-maker IDs in our analysis sample are very fast traders that account for the vast majority (87%) of the liquidity provided by this group and for around 95% of order submissions. By focussing on this group, we very likely capture the “modern” (high-frequency) market makers, and we henceforth abbreviate these traders as HFT-MMs. Notably, all TSX listed stocks have so-called equity specialists (identified by the

¹⁰For traders with low median scores, the averages are similar to the medians. Many traders frequently post perfectly unbalanced scores. By using the median, we can ensure that this high frequency of unbalanced submissions is properly reflected in the score.

¹¹One challenge in the data is that a single entity, such as an HFT firm, may use multiple IDs and they may use, for instance, one ID to post buy orders and another for sell orders. In the appendix, we describe how we address this issue by identifying IDs in clusters.

ST account moniker in our data). None of these traders qualifies as a market maker in our analysis.

Inventories of Market Makers. A common perception is that high-frequency trading firms aim to hold no or only very small overnight inventories. We observe that most trader IDs that we classify as market makers hold substantial median end-of-day inventories, even in non-interlisted securities and even though their posting behavior is balanced. Furthermore, several of the fastest trader IDs (see the appendix for details on the classification of “fast”) trade more than 85% passive, have order-to-trade ratios in the 99th percentile, and yet hold median inventories of 70% or more of their daily trading volume.¹² As Stephen Cavoli from Virtu explained during a recent industry conference,¹³ Virtu *hedges* with related securities when they accumulate an inventory so that they would end the day “flat” in terms of *risk* — but not necessarily in terms of their position.

The concern for our analysis of market making is that without knowing such hedging strategies, we cannot assess how such behavior affects markets. Cavoli’s statement shows, however, that not all market making HFTs manage their inventory on a stock-by-stock basis. Thus even though we believe to have market making behavior correctly identified by persistent two-sided order submissions, not all traders in our sample appear to be managing their inventory on a stock-by-stock basis.

For this reason, we perform our analysis of inventory management based on two groups of market makers. The first is the one that we have employed thus far. The second group is the sub-group of trader IDs that maintains an average intra-day inventory of under $\pm 20\%$ for our analysis sample. This group consists of 20 trader IDs.

¹²This observation alone highlights the importance of understanding the usage of trader and user IDs in different jurisdictions and in different datasets. In Canadian markets, a single DMA client may use multiple trader IDs (IIROC (2012) and IIROC (2014)), and it is thus possible that an HFT firm is assigned multiple user IDs. Furthermore, a single user ID may be used for trading activity of multiple entities, for instance, for all the brokerage’s retail order flow (which is balanced, on average). As a consequence, low end-of-day inventories are neither a necessary nor a sufficient attribute of an HFT trader ID in our dataset.

¹³The 16th Annual TD Securities Portfolio Management and Market Structure Conference, on November 5, 2015.

IV. Market Maker Reactions to Trades

A. Background

Practitioners often describe the phenomenon of the so-called quote-fade, a situation when following a trade, available liquidity at the best prices disappears market-wide. Models of asymmetric information, such as Glosten and Milgrom (1985) and Glosten (1994), describe that a market maker adjusts quotes after observing trades, usually by adjusting posted prices upward after buys and downward after sells. The reason for the adjustment is that trades, on average, convey some information about the fundamental value. For instance, a sell (at least on average) reveals that the seller believes the stock to be overvalued and thus, upon observing a sell, the market maker adjusts the price downward. An alternative view is that as a market maker takes an inventory, e.g., when buying from a seller, he assumes a liability. If the stock were to fall in price, he would lose. Assuming that the market maker is risk averse, he would accept additions to his inventory only at lower prices. In a single-market environment, the market maker can post a schedule of buy and sell limit orders to the order book that accounts for future trades and he may adjust them from time to time as he learns from the arrival of new orders. Indeed, both information and risk-aversion would thus imply that market makers adjust quotes on the opposite side of the market relative to the trade, i.e., they should adjust the bid following a sell and the ask following a buy.¹⁴

In an environment with multiple markets, the situation is more complex. If markets would be fully integrated, i.e. if there was no latency between markets, then, loosely, market makers would have to split the quantity that they post on a single venue between multiple venues. The reason is that they should expect that a trader with a market order would try to access multiple markets at the same time. However, if there is latency between the market, then the market makers can try to quote on multiple markets and adjust quotes as he observes trades on different markets. In this paper, we will assess whether this latter situation arises.

By canceling quotes, a market maker aims to ensure that he is not “picked off” in

¹⁴One can further make a case that the market maker should also adjust the quote on the same side of the market, where the latter argument usually relies on a competition-based argument. We focus on the opposite side here to simplify the exposition.

the future. In addition, it is possible that there are other traders' orders in the book at now "wrong" prices, e.g., buy orders that are at bid prices that became "too high" after a sell. These stale quotes can present a profit opportunity if the fundamental value has moved sufficiently. Thus assuming that non-HFTs are slow to react to the information contained in prices and that such traders' orders are present in the order book, we should observe that fast traders, including HFT market makers, take advantage of stale quotes and trade against them. An alternative view is that market makers have predictive power over future order flow and thus trade in anticipation of such flow. Namely, large buy-side orders ("parent" orders) are usually not traded in one large chunk, but they get split into smaller ("child") orders that are traded over time. If, by observing the order flow, market makers can detect the presence of a large parent order (for instance, because the buy-side child orders are traded with a poorly designed algorithm), then market makers may try to trade ahead of the rest of the large order and, as a consequence, would trade against existing orders in the book. Both views presented here would result in a situation where, subsequent to the trade, HFT market makers submit aggressive trades in the direction of the trade.

The analysis in this section involves determining, for each trade, the immediate reaction of market participants in terms of order cancellations and aggressive order submissions, at the trader level and across multiple markets. This type of analysis is computationally complex and data-intensive and, to the best of our knowledge, has not been performed at this level of granularity for multiple markets. Key to the analysis is that we have the best available information regarding the order of events at the highest available level of granularity, and thus we can observe who did what at which time. However, there are intrinsic limitations to this approach: as is well known quantum theory, events can happen in sequence but parties involved with the events may not know of the other, even when they can be ordered by time-stamps. Specifically, to draw a causal link from an event occurring on market A to an event on market B, it is insufficient to observe that the event on market B occurred after the event on market A for it is not clear that the trader acting on B knew of the event on A. Even if we assume that the market B trader monitors market A permanently, information still has to flow from market A to the trader so he can act on market B. What we aim to do is to highlight statistical regularities (or lack thereof).

B. Variable Construction

Using the Analysis Sample, we proceed as follows. First, we match all orders with the prevailing local bid and offer prices and determine whether or not the order was aggressive in the sense that it was either marketable (based on whether or not it “crosses the spread”) or immediate-or-cancel (IOC) (which, arguably, is an order submitted with aggressive intent). Second, we determine the visible order cancellations for the lit marketplaces; we do not count cancellations of IOC orders because the cancellation of such an order is automatic when the order doesn’t trade, and we also exclude dark orders because they do not contribute to the visible quote.¹⁵ We further determine the type of trader that submitted or cancelled the order, and the marketplace where the order was submitted.

Many marketable orders trigger multiple transactions. We aggregate trades that originated from the same trader on the same marketplace within 5 milliseconds. We note that the time stamp that we observe is not the time stamp of when the order was received or processed, but when the event is reported, and the report can be delayed when multiple events occur on the server at the same time.¹⁶ The 5ms time horizon ensures that we can collect related trades. We aggregate the total value for related trades, and base our matching with cancellations and aggressive orders on the time of the first transaction that is part of the trade.

At times, traders also trade at multiple marketplaces at the same time or in quick succession. At the millisecond level, such trades require the use of smart order routers (SORs). Multi-venue trades can arise when the broker instructs the SOR to access multiple markets or when the exchange sends an unfilled portion of an order to a different market to abide by the order protection rule. We classify SOR trades as a trade for which all of the following three conditions are met: (1) the trade originated from the same

¹⁵Determining cancelled volume is a non-trivial task because several marketplaces only report a cancellation and not the cancelled volume. To remedy this issue, we thus match each cancelled order with its original order. Moreover, some marketplaces allow amendments of existing orders. We count a price-amendment as a cancellation, we count a volume reduction as a cancellation of the amount by which the order was reduced, and we count a volume increase as a new submission for the additional volume.

¹⁶Specifically, we observed instances when the exchange-issued sequence numbers are not in single increments, even though the time stamps coincide — this indicates that other events occurred on the marketplace, thus slowing down the server’s reporting.

trader, (2) the trade involves transactions on at least two different marketplaces within 5 milliseconds and (3) the time from the first to the last transaction does not exceed 9 milliseconds. For each trade, we further determine the type of trader that submitted the marketable order, the type that provided liquidity, and we determine whether the order exhausted the posted depth at the local marketplace.

For every trade, we then compute the volume, value and number of orders cancelled by market makers and by all traders during the 1, 2, 3, 4, 5, 10, 20, and 50 milliseconds subsequent to the trade on marketplaces other than the one where the trade occurred. We omit events in the same millisecond as the trade because we cannot determine the correct order of events across multiple venues.

According to Hibernia Networks, a low-latency data line provider, the geographical latency between major Canadian data centers was around 400 microseconds; within Toronto, the latency is around 41 microseconds. Both of these latencies are, unfortunately, smaller than the smallest time increment in our data. That being said, the marketplaces' hardware (e.g., routers and servers) and the market participants' own systems add further latency.

Our focus is on aggressive order submissions in the same direction of the trade (buy-orders if the trade was buyer-initiated) and on cancellations in the opposite direction of the trade (e.g., cancellations of passive buy-orders following a marketable sell order).

C. Aggregate Observations

For comparative purposes, we split each trading day from 9:30 a.m. to 4:00 p.m. into 5 millisecond intervals and computed the number of such intervals with 0, 1, 2, ..., 9 and 10 or more cancellations and marketable order submissions. There are 4.68 million such intervals in a day, and most of them show no observations; for comparison: the average number of cancellations per day per security is about 50,000. In our data, we observe that a total 33.8 million of the 5ms intervals have one or more cancellations; 25.4 million of these have one or more HFTMM cancellation.

We now compare these aggregate numbers to those that we will analyze. For trades, we only consider the subset of trades that are not initiated by an HFTMM. Of the 2.1M trades in our sample, 0.55M are initiated by HFTMMs. We observe that 9% of the 33.8M (3.1M) cancellation intervals are within 5ms of a trade where the cancellation

occurs on a venue other than the one where the trade occurred; 3.04M intervals have HFTMM cancellations.¹⁷

Next, overall we observe 6.38 million intervals with aggressive orders, 1.7M stemming from HFTMMs.¹⁸ For HFTMMs, 6% of their aggressive order time-intervals follow a trade, whereas for all traders combined, only 2% of these intervals follow a trade within 5ms on a venue other than that of the trade (excluding only-HFTMMs intervals, only 1% of intervals follow a trade directly). For cancellations, 0.46M with-cancellation intervals follow a trade that occurred on different venue within 5ms, and the vast majority, million, of these, 0.43M involve HFTMMs. However, 92% of trades are not followed by an aggressive order, and 58% of trades are not followed by a cancellation on a different venue.

In what follows, we focus on the trade data, where we study those trades that are not initiated by an HFTMM; there are 1.55 million such trades. Of these, about 47K/250K are followed with an other-market HFTMM aggressive order/cancellation respectively within 1ms, and 180K/716K are followed by an HFTMM aggressive order/cancellation within 5ms.

D. Determinants of Trade-Reactions in a Trade-by-Trade Regression Analysis

In this subsection, we determine which factors contribute to the occurrence of aggressive orders and cancellations by HFTMMs subsequent to a trade on a different venue. As a first step, we note that not all situations lend themselves to examine “quote fade” and “latency arbitrage”. Namely, aggressive traders can only hope to find a counterparty on a different venue if the national best depth in the direction of the trade (the ask for buys, the bid for sales) exceeds the local depth. We thus restrict attention to such situations; of the 1.55M trades in our sample, 1.05M occur at a time when the national depth at the best price exceeded to local depth at the same price (for buy trades, at the ask, for sell trades at the bid).

Our goal is to assess the probability that a trade is followed by a cancellation or aggressive order submission by an HFTMM in the 1 and 5 milliseconds following the

¹⁷We recognize that the comparison is imperfect because the aggregate number is computed on a fixed grid while our after-trade cancellation metric is relative a preceding trade, and there is the possibility that the same cancellation is more than once.

¹⁸The main reason why the number of aggressive orders vastly exceeds the number of trades is that we count IOC orders, even if they do not result in a trade.

trade. We thus ran a probit regression based on the following equation

$$\begin{aligned}
DV_{t+m} = & \alpha + \beta_1 \text{takebook}_t + \beta_2 \{\text{vol}_t > 200\} + \beta_3 |\text{lmb}_t| + \beta_4 \text{early}_t + \beta_5 \text{late}_t \\
& + \beta_6 \text{momentum}_{t-10,t} + \beta_7 q \times r_{t-10,t}^\pm + \beta_8 \text{qspread}_t + \beta_9 \ln(\text{cumval}_{it}) + \beta_{10} \text{totaltrans}_{it} \\
& + \beta_{10} \% + \beta_{11} \text{totaltrans}_{it} + \beta_{12} \text{time since first}_{it} + \beta_{13} \text{voltime since first}_{it} \\
& + \beta_{14} \text{VXX}_t + \beta_{15} \text{takebook}_t \text{SOR}_t + \beta_{16} \text{takebook}_t \text{not SOR}_t \\
& + \beta_{17} \text{not takebook}_t \text{SOR}_t + \beta_{18} \text{HFTMM passive}_t + \epsilon_t,
\end{aligned} \tag{2}$$

where DV_{t+m} is the dependent variable that measures the volume of cancelled orders in the opposite direction of the time t trade in the $m = 1, 5$ milliseconds following the trade at time t and the remaining variables are as described in detail in what follows. Some of the above explanatory variables are substitutes, and we thus only include a subset of the variables in each regression specification, as indicate in our regression tables. Each trade is an observation. We employed clustered standard errors at the security level. Namely, we believe that the following variables are of interest a priori.

1. Aggressiveness of the trade.

Trades that exhausts the local depth are likely most aggressive, and we employ a dummy for such trades, takebook_t . Furthermore, a trader may also attempt to access liquidity on multiple marketplaces using a smart-order router. We thus add a dummy for such trades, SOR_t ; we use this dummy in specifications where we split trades three ways: those that trade the full local depth and are SOR, those that do the same and are not SOR, not SOR_t , and those that are SOR but do not trade the full local depth.

2. Trade size.

Most trades are for 100 or 200 shares; we thus use a dummy for trades that are for 300 shares or more, $\{\text{vol}_t > 200\}$; these are conceivably large enough so that market makers may suspect impeding trades to occur at other venues.¹⁹

3. The absolute value of the day's cumulative trade imbalance.

Large trade imbalances are associated with market movements, which in turn may be caused by a prevalence of buyers or sellers, $|\text{lmb}_t|$.

¹⁹The 75th percentile of HFTMM aggressive trades is 200 shares; for retail it is 400, for institutions 300.

4. Time dummies for the first and last half-hour of trading.

The time close to the open and close often involve most of the activity and behavior during this time may be different; $\text{early}_t, \text{late}_t$.

5. Features of the submitter of the marketable order.

Of special interest are traders that build positions. For each trader, stock and day, we thus determine if this trader performs all of his trades in a single direction. Of these, we focus on those that trade at least 10 times. We perform two regressions: one for all trades, and one where we focus on those that follow under this classification (roughly half of the sample).

Furthermore, for all traders and all trades we compute the total number and the fraction of the day’s total number of trades that the trader has completed at that point, totaltrans_{it} and $\%\text{totaltrans}_{it}$; the logarithm of the total value traded for this trader at the time of the trade $\ln(\text{cumval}_{it})$, and the passage time (measured in calendar time, $\text{time since first}_{it}$, and volume time, $\text{voltime since first}_{it}$) from the trader’s first trade to the current trade (the first trade has value 1). This variable will help us determine whether “later” in their order they are more or less likely to experience front-running.

6. Short term momentum.

We consider two variables. First, we determine whether the trade is in the direction of recent price movements, $\text{momentum}_{t-10,t}$. For this situation, we determine the cumulative midpoint return over the last 10 trades. If the cumulative return was positive and the trader buys, or if the return was negative and the trader sells, then there is a momentum trade. Second, we also compute the cumulative midpoint return over the last 10 trades and multiply this return with the trade direction, $q \times r_{t-10,t}^{\pm}$.

7. The level of the spread.

When the spread is wider (in cents), latency arbitrage may be more difficult because traders may be more cautious in submitting orders to multiple venues, qspread_t .

8. The type of liquidity provider.

We use a dummy for whether or not one of the liquidity providers was an HFTMM, HFTMM passive_t . The idea here is that it is possible that the trader gets notified about the fill before the general market and may thus race to the next venue to cancel the order quickly.

9. Marketwide Intra-day volatility.

We proxy marketwide intra-day volatility with the natural logarithm of the quoted midpoint of the exchange traded note VXX, which tracks the U.S. volatility index VIXI VXX_t . In untabulated regressions, we also use the change in the VXX since the beginning of the day; the results are robust.

As explanatory variables we use dummies for whether or not there was a cancellation/aggressive trade in the 1 and 5 milliseconds after the trade.²⁰ In our description we thus discuss not only the regression outcomes themselves, but also the difference between the effects for the 1ms and 5ms settings.

Summary Observations. Table II provides some summary statistics for the after-trade aggressive order cancellations and aggressive order submissions; Figure 1 illustrates these numbers graphically. We represent the percentages of HFTMM of the total order submissions/cancellations and the total number of observations per millisecond. These numbers are based on the total number of aggressive trade and cancellation observation for the entire sample for those trades where the local market was not the only market at the NBBO. These figures represent order submissions on venues other than the one where the trade occurred.

We observe that the HFTMMs make up between 24% and 40% of all the aggressive orders submitted after trades. Although this fraction may appear small, we note that aggressive orders include those from “spray”-SORs that try to access multiple venues simultaneously. Moreover, there are other HFTs that may have more aggressive strategies in the first place.

Next, about 93% of order cancellation against the direction of the incoming trade stem from HFTMMs. This number drops to 90% after 50ms.

However, the numbers are not large. Only about 4.2% of trades are followed by other-venue aggressive trades by HFTs. We note, however, that the number of per-

²⁰In untabulated regressions, we also performed the analysis for the 2,3,4 etc. milliseconds. However, the estimates for these time intervals were either very similar or they showed a trend from 1 to 5ms.

millisecond aggressive trades drops continuously from the time of the trade, suggesting that the occurrence of the aggressive order submission is not random. Focussing on situations when aggressive trades after the trade do occur (by any trader, including those from spray SORs), there are much more HFTMM aggressive orders and other-side cancellations.

Results for Aggressive HFTMM Orders. Table III contains our regression results; we provide only the marginal effects as estimated by STATA; the coefficient estimates are available from the authors. We observe that absorbing the book increases the likelihood of observing an aggressive order by 4%. Similarly for trades of 300 shares or more. For several specifications, there are indications of weaker effects in the morning and stronger effects towards the close. The cumulative trade imbalance at the time of the trade appears to play no role. The total value that the trader has traded plays no role for the total sample, but for the restricted sample that only considers trades by directional traders, there is a positive effect, indicating that such traders can be detected as they accumulate a position. However, the number of transactions or the passage time (either in hours or in the day’s volume-time) appears to have no effect. For transactions, in fact, the relation is negative. Momentum is associated with a negative effect when accounting for the size of the return, implying that aggressive orders by HFTMMs are observed when the trade was against the flow. Put differently, HFTMM appear to follow short-term contrarians. For the general population, having an HFTMM on the passive side marginally reduces the probability of observing an aggressive trade; for the directional trader subsample, there is no effect. There is thus no indication that HFTMMs can use the (possibly faster-received) information in trade reports to pick off stale orders. Finally, SOR trades are generally associated with increased likelihood of aggressive orders.

The difference between the 1 and 5 millisecond estimations is that the coefficient estimates are generally larger, roughly by a factor of 3 to 4. One explanation for the larger magnitude is that there are three times as many trades with aggressive order submissions within 5ms than 1ms.

Results for HFTMM Order Cancellations. Table IV contains our regression results. The variables that increase the probability of a aggressive order also increase the probability of HFTMM order cancellations, including relating to the attributes of the trade-

initiating trader. The magnitude of the effect is generally larger (by factor 2), mirroring the larger rate with which cancellations actually occur. Notable additional piece of insight relate to the market-wider order imbalance. Here we observe that HFTMM cancel fewer orders as the imbalance grows (we will address this issue in the following sections when we discuss their posting behavior — as we argue there, when imbalances increase, HFTMMs quote less and thus have less to cancel). Furthermore, we observe that there is a strong relation between the cancellation rate and HFTMM being on the passive side of the trade. This latter point indicates that HFTMM don't always intend to actually trade the consolidated size that they post on multiple venues.

V. Causality of Latency

In this section, we study how the technological change of moving marketplace A to the same platform as marketplace B affected the frequency of after-trade aggressive trades and cancellations.

For this part of our study, we aggregate the occurrences of after-trade aggressive orders and cancellations by stock and day and perform a panel analysis. Specifically, we estimate the following regression equation

$$DV_{it} = \sum_{j \in \{A, \dots, F\}} \alpha_j \times \text{quantum}_t \times \text{market}_j + \beta_1 \times \%HFTMM_{ti} + \beta_2 \times VXX_t + \delta_i + \epsilon_t, \quad (3)$$

where the dependent variable DV_{it} is the aggregate after-trade submission of aggressive orders and cancellations; quantum_t is a dummy that is 0 before May 1, and 1 thereafter, market_j is a dummy for market $j \in \{A, B, C, D, E, F\}$; $\%HFTMM_{ti}$ is the percentage of trades per stock per day that involve an HFTMM on the passive side, and VXX_t is the trade-weighted average of the volatility ETF VXX. The specification also contains stock and market fixed effects, and we cluster standard errors by date and security.

As a first step, we estimated (3) using some liquidity related variables as the dependent variables, namely: $\%HFT_{ti}$, i.e. the percentage of trades per stock per day that involve an HFTMM on the passive side; the by-marketplace quoted spread (in bps of the midquote); the by-marketplace time-weighted dollar-depth; and the fraction of time that the respective markets are at the NBBO. For the aggressive order submis-

sions and order cancellations, we use several different measures: the total number of orders/cancellations, the total number of occurrences trades with HFTMM and non-HFTMM orders/cancellations, and the fraction of HFTMM orders/cancellations.²¹

Results. The most interesting venues for our analysis are marketplaces A, B, and C, as these markets combined account for 85% of all trades and 92% of dollar-volume traded for the trades that we consider in our analysis. We note that after markets A and B joined on the same platform, market A’s market share of dollar volume and number of trades increased by about 2% and 4.5% respectively, whereas market C saw a decline of 1.5% and 3.5% respectively.

We begin by discussing our results for the liquidity-related analysis. Table V contains our findings. The first three row entries contain the markets that are now on the same trading platform, the next three rows are for the marketplaces that remained at their locations.

We observe that the participation rate of HFTMMs on the passive side declined significantly for markets A and C and weakly increased for market B. The declining HFTMM participation is then reflected in a decline in the posted depth for these markets (though depth also (weakly) decline on market B). In untabulated regressions, we analyzed whether this drop in by-market depth led to a reduction in consolidated depth across markets but we found no significant change. We further observe no change in by-market posted bid-ask spreads (and neither for the NBBO), but we note that the by-market spread is a very noisy measure: the time-weighted spread for most markets is rather large, even though the NBBO and the spread for the main market, B, is small (around 2 cents on average). Finally, we observe that market A is at the NBBO (weakly) less often.²² Overall, our results suggest that there are noteworthy changes to the posting behavior of HFTMMs.

We next turn to the occurrences of aggressive order submissions; Table VI contains our findings. For the 1ms after-trade window, we observe a significant decline in HFTMM activities for market C and an increase (albeit not statistically significant) for market A. We also note that there is an increase in non-HFTMM activities on markets A (and

²¹We also used the occurrences and the total number of orders and cancellations *per trade*; the results are similar and thus omitted.

²²The average (not included in tables) time per day for being at both the best bid and ask declines from 39% to 34%; for market C, we observe the same magnitude of a decline, from 37.5% to 33%. Market D declines from 13% to 6%. For comparison: market B is at the best 81% to 82.5% of the time.

F). For the 5ms window, the situation is similar with regard to the total number of occurrences: there is an increase trades with HFTMM aggressive orders for market A and a decrease for market C.

For the relative fraction of HFTMMs for aggressive orders, the result is stark: there is a relative increase for trades from all the venues that are in the joint location and there is a relative decrease for all venues that are in other locations. The results for markets A and C are particularly noteworthy. To be able to take advantage of stale quotes, there have to be quotes on the other market and it must be possible to reach this market quickly. Market A is now much closer to the main market B, which is also the most liquid. Thus observing trades on market A now likely gives HFTMMs easier access to market B right after trades on A.

Results for cancellations are in Table VII. We observe increases for market A, and decreases for markets B and C for the 1ms and 5ms windows (depending on the measure with stronger and weaker statistical significance). It is important to relate this finding to our analysis of HFTMM participation. Our results there showed that HFTMMs participate less in markets A and C, suggesting that they also post less on these markets. This decline explains the decline in after-trade cancellations for trades on market B: HFTMMs have fewer orders on the other markets. Likewise, the increase in occurrences of after-trade cancellations for market A can be explained by the lower latency to market B: after observing trades on A, HFTMMs can now rush to B and cancel their orders there and avoid being picked off.

In untabulated regressions we re-ran the probit analysis to estimate whether there are differences in the contributing factors for aggressive orders and cancellations. We found little change.

Overall our results suggest that market A, which joined market B, saw an increase in quote fade and latency arbitrage whereas market C, which is now physically removed from one more market, A, saw a decrease in such activities.

VI. Market Maker Reaction to Market-wide Developments

We will now analyze how market makers behave as the market moves. For this part of the analysis, we segment the trading day into volume intervals. Using volume intervals allows us to study behavior in “event”-time as opposed to calendar time. Arguably,

market maker behavior is most interesting to study not so much depending on the time of the day but, rather, when markets move. Moreover, calendar time poses econometric and interpretational difficulties in the sense that it is difficult to understand times during which little or nothing happens; such situations arise often if one chooses a fine time grid (e.g., minutes). When choosing a wider grid, one may miss or misinterpret situations when, for instance, market makers aggressively manage their inventories over very short stretches of time.

Specifically, each volume interval is based on 1% of the average daily volume per stock, where the average is computed over our sample horizon. Thus on an average day, we would observe 100 volume intervals. For each volume interval we then compute market's trade imbalance, and the market's cumulative trade imbalance. The market's trade imbalance is the difference of buyer- and seller-initiated volume relative to total volume; the cumulative trade imbalance is the computed akin to the measure for contemporaneous volume, except that all volume figures are summed since the first volume interval of the day.

For market makers, we are interested in third measures. First, the by-trader non-marketable order volume imbalance, which is defined as the difference between newly submitted buy and sell order volume relative to total order volume in that volume interval by that trader. This measure captures whether, in response to market developments, market makers change the relation of buy and sell orders. Note that we count an amended order both as a new order submissions and as a canceled order. Second, we compute the average price that market makers submit relative the last prevailing bid and ask prices in the preceding volume interval. And third, analogously to the preceding section, we compute the frequency of order submissions relative to the best posted prices.

Another measure that one may be interested in is imbalance of cancelled order volume imbalance, defined analogous to the order volume imbalance, except that it uses canceled volume. Arguably, a market maker who leans against the buying pressure would not cancel sell-orders. However, cancellations are more difficult to interpret because by looking at the aggregate imbalance only, we cannot determine where in the book orders were cancelled. In principle, a market maker could cancel the sell orders that are far from the best prices and post them closer. We thus do not include cancellations in this part of our analysis.

In the analysis in this section, we are interested in the reaction of market makers to aggregate order flow. In our regression analysis, we thus use a simple specification where we regress a dependent variable that captures market maker behavior on lagged trade imbalances (either single-period or cumulative), where we distinguish positive (buyer-dominated) and negative (seller-dominated) imbalances.

A. Imbalance of Buy vs. Sell Orders.

We analyze how the lagged market-wide order imbalance affects the submission of new buy and sell orders by market makers in the sense of possibly tilting the market makers towards submitting more buy or sell orders. Generally speaking, if there is buying pressure, a market maker who leans against the price would submit more sell orders whereas a market maker going with the flow would submit more sell orders. Specifically, we estimate the following equation

$$DV_{it} = \alpha \times \text{trade imbalance}_{i,t-1} + \delta_i + \epsilon_t, \quad (4)$$

where DV_{it} is the imbalance of buy- and sell-order volume discussed above for firm i at time t , $\text{trade imbalance}_{i,t-1}$ is the difference between buying and selling trading volume relative to all trading volume in volume either only for interval $t - 1$ or cumulative since the beginning of the trading day; and δ_i is a firm fixed effect. We compute the imbalance measure by aggregative order volume across all market makers.

Our estimation results are in Table IX. We consider both the group of all market makers as well as the group of inventory-managing market makers. We note that the imbalance of market makers' order submissions and the trade imbalance are positively related, and that the relationship is stronger for the cumulative imbalance. This result implies that market makers submit their orders in the direction of the market.

B. Prices of Buy and Sell Orders.

Generally speaking, liquidity provision is the willingness to assume a risk at a price. The prices that market makers post relative to existing prices will thus signify how much the market has to pay liquidity providers to trade. In the most extreme view, market makers that truly provide liquidity would “lean” against the order flow and post at

prices that are not worse relative to the existing ones (e.g., they submit sell orders no higher than the last posted ask).

Formally, we now analyze how the lagged signed market-wide order imbalance affects the prices of new buy and sell orders by market makers relative to the last posted NBBO bid and ask prices in the last volume interval. Formally, we use $\text{vwap}_{t,b}^o, \text{vwap}_{t,s}^o$ for the volume weighted average prices of all buy and sell orders respectively in volume interval t and $\text{bid}_{t-1}, \text{ask}_{t-1}$ are the last prevailing NBBO bid and ask prices in volume interval $t - 1$. We then define

$$\Delta^{\text{bid}} \text{vwap}_{t,b}^o = \frac{\text{vwap}_{t,b}^o}{\text{bid}_{t-1}} - 1, \quad \Delta^{\text{ask}} \text{vwap}_{t,s}^o = \frac{\text{vwap}_{t,s}^o}{\text{ask}_{t-1}} - 1,$$

so that these two numbers will capture the submission prices of buy order relative to the last bid and sell orders relative to the last prevailing ask. We then estimate the following regression equation

$$DV_{it} = \alpha \times \text{trade imbalance}_{i,t-1} \times \text{buy}_{it} + \beta \times \text{trade imbalance}_{i,t-1} \times \text{sell}_{it} + \delta_i + \text{Int}_t + \epsilon_t, \quad (5)$$

where DV are the dependent variables that we discussed above that capture market maker behavior, $\text{trade imbalance}_{t-1}$ is the difference between buying and selling trading volume relative to all trading volume in volume either only for interval $t-1$ or cumulative since the beginning of the trading day; buy_t and sell_t are dummies that are 1 if the trade imbalance is positive (more buys than sales) and negative, respectively; and Int_t is a time-trend for the day to capture if/when imbalances are larger towards the end of the trading day. A negative number for $\Delta^{\text{bid}} \text{vwap}_{t,b}^o$ and a positive number for $\Delta^{\text{ask}} \text{vwap}_{t,s}^o$ signify that traders submit orders that improve the current bid and ask prices.

Our estimation results are in Table X. We present only the estimated coefficients for the inventory-managing market makers; the estimated coefficients using all market makers are very similar. Most of the statistically significant findings relate to the cumulative imbalance whereas the per-interval imbalance has little explanatory power. Moreover, most results relate to buying order volume (which form bid prices) whereas we find few or no significant effects for selling volume. The results for buying volume indicate that market makers do not lean against the order flow. Namely, the more positive trade imbalances become (more buyers than seller), the smaller the buying prices that market

makers submit, i.e., the worse are the prices. For more the negative imbalances become, the higher the prices at which they are willing to buy. Results for selling prices are statistically insignificant.

C. Where in the book?

As the final step, we consider where in the book market makers post their orders in reaction to buying or selling pressure. For this analysis, we focus on the group of inventory-managing market makers. Panel A in Figure 2 provides a three-dimensional plot of the order submissions relative to the cumulative trade imbalances; Panel A in Table VIII displays the underlying data. Namely, each column represents for each price the average fraction of their orders that market makers have submitted relative to the best prices, conditional on the level of cumulative imbalances. We note a few general patterns. First, the largest fraction of market-maker orders is submitted at the best price. Second, a noticeable fraction of orders is submitted at prices that are much better and much worse than the best. Third, as the trade imbalance becomes larger, market makers submit fewer orders at best prices and more orders at much improved prices. For trade imbalances below 80%, market makers post (weakly) increasingly far off the best prices.

In a formal regression analysis, we focus on five measures: at the best, ± 1 tick, and ± 2 or more ticks. The submission data has been constructed relative to best prices irrespective of the direction of the order,²³ and as an explanatory variable we will thus use the absolute values of trade imbalances. Moreover, as dependent variables we compute the fraction of the market makers' orders that are submitted, because this measure captures how market makers structure their behavior as the market moves. Table XI displays our results. We observe that market makers react strongly to short term imbalances in the sense that they significantly reduce the fraction of orders submitted at the best or improved prices. However, for cumulative imbalances, the reaction is different: indeed, they submit more at aggressive and at very aggressive but also much worse prices, while at the same time reducing their presence at the best prices. We note that, by construction, none of these orders that we consider are marketable.

²³In a future iteration of this paper, we will further differentiate by order.

VII. Market Maker Inventory Management and Liquidity

In the past section, we studied the relation of order submission behavior following market imbalances. We will now focus on the market makers' own inventories. Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010) show that NYSE specialists manage their inventories and in doing so, affect prices, spreads and investors trading costs. For their analysis, Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010) aggregate inventories by specialist firm and day and show that as NYSE specialists accumulate inventories, they post worse prices and thus reduce liquidity in markets. "Modern" markets, however, differ substantially from the NYSE when monopolist specialists were in charge of organizing trading because in modern markets, voluntary market makers compete for order flow with other market makers and with investors who want to trade with limit orders. Consequently, it is an open, empirical question whether "modern" market makers can affect prices at all and whether and how their inventory management affects prices. Due to data-information constraints, we cannot aggregate inventories by firm. We also do not want to aggregate inventories by day because we would not be able to capture the effects of the "modern" market makers' intra-day inventory management.

An inventory is defined as the difference between buy and sell volume relative to total volume for the day up to and including the volume interval. Our analysis in this section is henceforth similar to that in the preceding section, except that we are now interested in the effect of the one period lagged market maker inventory on market maker order submission behavior. To assess the impact of market maker inventory management, we additionally compute, for each volume interval, the time-weighted quoted bid-ask spread (in cents and in basis points of the prevailing mid-quote), the last quoted bid and ask prices, the volume-weighted average price of trades, the range measure (the largest midquote minus the smallest midquote divided by the time-weighted average midquote), and the time-weighted midquote return.

As a first step, Figure 3 illustrates the distribution of by-trader inventories across all volume intervals. Panels A and B plots the histogram of trader inventories for all traders and all non-market making trader IDs in our sample. Since traders are assumed to start the day with zero inventories, there is necessarily a concentration at zero. However, there are also many $\pm 100\%$ inventories. Panel C For this group, there is a significant

concentration at $\pm 100\%$ intervals and it is not clear to us that these traders actively seek to trade out of this positions or, rather whether these traders actually hold these positions or whether they have been offset within the HFT firm by another trader ID's position. Panel D plots the inventories of the inventory-managing market makers. As can be seen (as as one would expect given their classification), these traders keep very small inventories. However, there is a slight uptick at $\pm 100\%$ intervals. For the remainder of the analysis, we will focus on the group of inventory-managing market makers.

A. *Market Maker Inventory vs. Market-wide Developments.*

As a first step, we seek to understand the relation between market maker inventories and market-wide trade imbalances. If market makers indeed make markets then their inventories should be negatively related to trade imbalances. We determine this relationship by estimating the following equation

$$MM\ inventory_{i,t} = \alpha \times trade\ imbalance_{i,\tau} + \delta_i + \epsilon_t, \quad (6)$$

where $MM\ inventory_{i,t-1}$ is the market makers' time $t - 1$ cumulative inventory (computed across all market makers) for firm i ; all other variables are as in (8); $trade\ imbalance_{i,\tau}$ is the market-wide trade imbalance at times $\tau = t$ (contemporaneous) and $\tau = t - 1$ (one period lagged).

Panel A in Table XII contains the results. For inventory-managing market makers, we observe a negative relation to market-wide trade imbalances both contemporaneously and with a one-period lag, suggesting that these traders operate as market makers.

B. *Inventories vs. Order Imbalance.*

We analyze how the lagged market-wide order imbalance affects the submission of new buy and sell orders by market makers in the sense of possibly tilting the market makers towards submitting more buy or sell orders. Generally speaking, if there is buying pressure, a market maker who leans against the price would submit more sell orders whereas a market maker going with the flow would submit more sell orders.

Specifically, we estimate the following equation

$$DV_{it} = \alpha \times MM \text{ inventory}_{i,t-1} + \delta_i + \epsilon_t, \quad (7)$$

where all variables are as defined for (8) and (6)

Table XIII contains the results. We observe a significant negative relation of market makers order submissions and their lagged inventories, suggesting that, as their inventory increases, they submit fewer buy orders and more sell orders.

C. Inventories vs. Prices

Analogously to the last section, we analyze the relation of submitted priced and market maker inventories. As in the last section, we split by the sign of market maker inventories. We then estimate the following regression equation

$$DV_{it} = \alpha \times MM \text{ inventory}_{i,t-1} \times \text{buy}_{it} + \beta \times MM \text{ inventory}_{i,t-1} \times \text{sell}_{it} + \delta_i + \text{Int}_t + \epsilon_t, \quad (8)$$

where buy_{it} and sell_{it} are dummies for whether or not the market makers have a positive or negative aggregate inventories.

Table XIV contains the results. We observe that as market makers' positive inventories increase, they lower their bid prices, and there is no evidence of them changing their ask prices. Likewise, as their negative inventories increase, they reduce their ask prices and they increase their bid prices; these latter two results are counterintuitive.

D. Inventories vs. Where in the Book?

Building on our analysis of the average prices of orders relative to the NBBO, we now study where in the book market makers post as their inventories grow. Panel B in Figure 2 is constructed analogously to Panel A in the same figure, except that the columns are drawn for market maker inventories; Panel B in Table VIII displays the underlying data. The general patterns in this figure are analogous to Panel A of Figure 2. Additionally, as the absolute value of inventories grows, market makers submit fewer orders at best prices and more orders at much improved prices. For worse prices, there is no general pattern that relates to the inventories.

In a formal regression analysis, we focus on the same five measures as in the preceding section: submissions at the best, ± 1 tick, and ± 2 or more ticks. Table XV displays our results. We observe that market makers react strongly to increases in their inventories in the sense that they significantly reduce the fraction of orders submitted at the best prices or one tick worse or one tick better than the best price. However, they also submit aggressively at prices that are two or more ticks better than the best prices.

E. Inventories and Liquidity

The findings thus far suggest complex reactions to growing inventories. On the one hand, market makers reduce their new buy order submissions as their inventories grow and they are willing to purchase only at lower prices. On the other hand, they also submit much more aggressively priced orders as their inventories grow. Taken together, the effect on liquidity is ambiguous. To determine this effect, we estimate (7), using bid-ask spreads as the dependent variable and using the absolute value of market maker inventories as the dependent variable. However, we find no effect.

To further investigate the mechanism at work, we split the sample into terciles according to the level of competition for liquidity provision. Namely, for each stock and day we compute the inverse of the Hirschman-Herfindahl Index for liquidity provision, computed as the sum of squared market shares, where a market share per trader is defined as the percentage of non-marketable volume of all volume per trader. We then estimate

$$DV_{it} = \sum_{j=1}^3 \alpha_j \times |MM\ inventory_{i,t-1}| \times comp_{ji} + \delta_i + \epsilon_t, \quad (9)$$

where $comp_{ji}$ is a dummy that is 1 if firm i is in the most ($j = 3$), medium ($j = 2$) or least ($j = 1$) competitive group for liquidity provision.

Table XVI contains the results from this analysis. We observe that for the most competitively traded stocks, market maker inventories are associated with tighter bid ask spreads, both when measured in cents and basis points of the midquote. We also observe that larger market maker inventories are associated with higher quote volatility, measured by the range measure. We find no evidence for a relation of inventories and the absolute value of midquote returns.

VIII. Conclusion

The main purpose of this analysis is to study the order submission behavior of modern, electronic market makers. The novelty of our work that sets it apart from the existing literature is that we study trader-level data across multiple markets. A major portion of the paper is devoted to studying latency arbitrage and the quote-fade phenomenon, which can only be observed when looking at multiple markets simultaneously. The work closest to ours is van Kervel (2015) who studies the impact of trades on posted depth across multiple markets, using public data for FTSE100 stocks. Building on his work, in the first part of our paper, we study behavior of high frequency market makers in the first 1-5 milliseconds after a trade (van Kervel aggregates over the first 100 milliseconds), we study the submission of aggressive orders by the market makers, and we perform a detailed analysis of the determinants of the occurrence of such behavior. Finally, we can identify the impact of latency by studying an event that eliminated the latency between two of the three main Canadian marketplaces.

Our study uses Canadian data. Even though the Canadian market shares many attributes and rules with the largest market, the U.S., and has the same high frequency market participants, there are notable differences that are relevant for our analysis. In the U.S., many market participants and, anecdotally, also some marketplaces rely on the slow, consolidated tape, the so-called SIP, for the NBBO — even though faster solutions such as the marketplaces’ direct feeds are available. In Canada, according to TD Securities, most market participants (in particular the brokers) rely on the direct feeds even though a consolidated tape from the so-called Information Processor (IP) exists. Second, geographic latency in Canada is much lower: the geographic latency for the furthest-apart market centers is around 400 microseconds, whereas latency from Chicago to New Jersey is around 40 times larger. Third, Canada has a dominating marketplace, the TSX. Intuitively, one would expect that latency arbitrage happens most often when the depth at multiple markets is thin so that a liquidity seeking trader needs to access multiple markets to fill an order. The TSX, however, usually has a lot of volume at the best prices and it is at the NBBO most of the time. Thus traders can often trade a fair amount of volume just at the single market. To summarize, within Canada, latency arbitrage is more difficult to realize and thus more difficult to find.

Yet even though the institutional and geographical setup is biased against finding

effects, we do find some indication of the quote fade and the latency arbitrage phenomenon in the terms of modern market makers engaging in the activity. One point of our work is to highlight the existence and to outline what to look for in the data. For the Canadian context, in our opinion, the effect itself is too small to be a major concern regarding market integrity. For a market as large and as geographically and institutionally dispersed as the U.S. or Europe, there is likely much more opportunity for latency arbitrage.

REFERENCES

- Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *The Journal of Finance* 48, 157–185.
- Biais, B., P. Hillion, and C. Spatt, 1995, An empirical analysis of the limit order book and the order flow in the paris bourse, *Journal of Finance* 50, 1655–1690.
- Boehmer, Ekkehart, Dan Li, and Gideon Saar, 2015, Correlated high-frequency trading, Working paper Cornell.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014a, High-frequency trading and price discovery, *Review of Financial Studies*.
- , 2014b, Market integration and high frequency intermediation, Discussion paper, IIROC Working Paper.
- Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes, 2010, Time variation in liquidity: The role of market-maker inventories and revenues, *The Journal of Finance* 65, 295–331.
- Comerton-Forde, Carole, Katya Malinova, and Andreas Park, 2015, Regulating dark trading: Order flow segmentation and market quality, Discussion paper, IIROC Working Paper.
- Devani, Baiju, Lisa Anderson, and Yifan Zhang, 2015, Impact of the dark rule amendments, Discussion paper, IIROC Working Paper.

- Devani, Baiju, Ad Tayal, Lisa Anderson, Dawei Zhou, Juan Gomez, and Graham W. Taylor, 2014, Identifying trading groups – methodology and results, Discussion paper, IIROC Working Paper.
- Easley, David, Marcos M. López de Prado, and Maureen O’Hara, 2012, Flow toxicity and liquidity in a high-frequency world, *Review of Financial Studies*.
- Glosten, Lawrence R., 1994, Is the electronic open limit order book inevitable?, *The Journal of Finance* 49, 1127–1161.
- , and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogenously informed traders, *Journal of Financial Economics* 14, 71–100.
- Hagströmer, Björn, and Lars Norden, 2013, The diversity of high-frequency traders, *Journal of Financial Markets* 16, 741 – 770.
- IIROC, 2012, The HOT study, Discussion paper, Investment Industry Regulatory Organization of Canada.
- Jones, Charles, 2013, What do we know about high-frequency trading?, Research Paper No. 13-11 Columbia Business School <http://ssrn.com/abstract=2236201>.
- Korajczyk, Robert, and Dermot Murphy, 2014, High frequency market making to large institutional trades, Discussion paper, IIROC Working Paper.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Malinova, Katya, Andreas Park, and Ryan Riordan, 2013, Do retail traders suffer from high frequency traders?, Working paper University of Toronto.
- Menkveld, Albert J., 2013, High frequency trading and the new market makers, *Journal of Financial Markets* 16, 712 – 740 High-Frequency Trading.
- Subrahmanyam, Avanidhar, and Hui Zheng, 2015, Limit order placement by high-frequency traders, Working paper UCLA and CMCRC Available at SSRN: <http://ssrn.com/abstract=2688418> or <http://dx.doi.org/10.2139/ssrn.2688418>.

van Kervel, Vincent, 2015, Competition for order flow with fast and slow traders, *Review of Financial Studies* 28, 2094–2127.

van Kervel, Vincent, and Albert J. Menkveld, 2015, High-frequency trading around large institutional orders, Working paper VU Amsterdam Available at SSRN: <http://ssrn.com/abstract=2619686> or <http://dx.doi.org/10.2139/ssrn.2619686>.

Wah, Elaine, 2016, How prevalent and profitable are latency arbitrage opportunities on U.S. stock exchanges?, Working paper University of Michigan.

Appendix: Classification of High Frequency Traders.

The user ID is the most granular identification that is available to regulators in Canada; IIROC researchers describe the usage of user IDs in detail in their research reports (IIROC (2012), Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014), and Devani, Anderson, and Zhang (2015)). According to these research reports, marketplaces assign user IDs, and an ID may identify a single trader, a business stream (for example, all orders that originate through a broker’s online discount brokerage system), or a client that accesses trading venues directly (through a direct market access (DMA) relationship). It is our understanding that the brokers separate different types of order flows (e.g., retail vs. institutional) by user ID. For DMA clients, IIROC requires dedicated IDs. However, according to Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014), a DMA client may be assigned more than one user IDs, for instance, to trade through multiple brokers or to trade on different marketplaces, and they may choose to use multiple user IDs for business or administrative purposes.

We classify traders as high-frequency based on their reaction speed to market events, which is, arguably, the definition feature of modern, ultra-fast electronic traders, where we require reaction times that are faster than human reaction times (the average duration of a single blink of a human eye is 100-400 milliseconds, according to the Harvard Database of Useful Biological Numbers). We further require that trader IDs exhibit such fast reaction times across many orders and trades, and in many securities. We use the following three specific criteria to quantify a trader ID’s reaction speed.

The first criterion is the trader ID’s median order-to-cancel time. The order-to-cancel time is the time from the submission to cancellation of the same order; for the purpose

of this classification, we exclude immediate-or-cancel (IOC) orders, because their order-to-cancel time is determined by the processing speed of the marketplace.

The second criterion is the number of trade and order messages that a trader ID submits during a short interval after a daily scheduled public information release. We focus on the first 500 milliseconds after 3:40 p.m., which is when the TSX first publishes the imbalance between the buy and sell orders in its market-on-close facility.²⁴

In aggregate, there is a significant spike in trades immediately after the publication of the market-on-close imbalance, though this spike may not be visible or pronounced on a stock-by-stock basis. Comerton-Forde, Malinova, and Park (2015) includes a plot of the by-minute number of trades (Figure 1 in their paper), aggregated over all securities in their sample (which is similar to ours, albeit for a different time horizon). The dataset that is provided to us by IIROC does not contain information on the market-on-close announcement. Thus, we are not able to determine the time between the publication of the market-on-close imbalance and a trader’s action at the millisecond level. For this reason, we classify trader IDs as HFTs based on their actions during a relatively long interval of 500 milliseconds after the announcement. These first two criteria were also used by Comerton-Forde, Malinova, and Park (2015).

The third criterion is the fraction of orders and (non-IOC) cancellations that a trader submits very quickly after a change in the order book that was not triggered by the trader him-/herself but by another trader.

For each trader ID, stock and day we compute the median order-to-cancel speed, and we compute the total number of orders and aggressive trades during the 500 milliseconds after 3:40 p.m. Furthermore, for each trader ID and security, we computed the number of orders and cancellations that this trader submits and that this trader submits within 1 millisecond following another visible order submission by a trader other than him-/herself in that security. A trader ID is classified as HFT

1. if the median of the trader ID’s median stock-day order-to-cancel speeds is below

²⁴The closing price for TSX-listed securities is determined in a multi-stage process. Before 3:40 p.m., traders may submit market buy and sell orders tagged as market-on-close orders. These orders will trade at the 4:00 p.m. closing price. At 3:40 p.m., the TSX publishes the imbalances of buy and sell orders, and traders then have the opportunity to submit priced limit orders to trade at the market-on-close to off-set the market order imbalance. The market-on-close imbalance is indicative of the closing price and may help predict behavior over the last 20 minutes of trading. The publication at 3:40 p.m. is merely the first publication. Between 3:40 and 4:00 p.m., TMX regularly publishes updates of the prevailing imbalance.

250 milliseconds, or

2. if the trader ID submits more than 1,000 orders or is involved in more than 500 aggressive transactions in the first 500 milliseconds after the market-on-close publication across all securities in our classification sample during our classification period, or
3. if the trader ID submits more than 85% of its orders and cancellations within 1 millisecond of some other trader’s order submission.

We classify a total of 101 trader IDs as HFT; of these, 78 are active in the Analysis Sample of non-crosslisted securities.

In Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014), IIROC researchers discuss their classification of HFTs. In their data, they had direct information about a subset of the existing HFT IDs in the Canadian market (49 IDs) and they used the knowledge of these IDs to apply machine learning techniques to identify further IDs as HFT. Their report does not list the details of the criteria, but Figure 1 in their report shows that “speed” in various forms is a decisive criterion. For their sample period from March to June 2013, which overlaps with ours, they identify 98 IDs as HFT, which is close in number to our 101 HFT IDs.

Cluster Analysis of HFT Groups. The by-trader data displays some pronounced similarities among subsets of traders in the sense that traders have very similar characteristics, for instance, in terms of numbers of trades and orders, or the number of securities traded. As IIROC researchers Devani, Tayal, Anderson, Zhou, Gomez, and Taylor (2014) highlight, HFT firms may use several trader IDs for their strategies. For instance, it is possible that an HFT firm uses one trader ID to submit buy orders and another to submit sell orders. Taking together, these two IDs may have a perfectly balanced end-of-day inventory, whereas individually their inventory is imbalanced.

The usage of multiple IDs is particularly important and presents a challenge for an analysis of market making behavior, in particular with respect to inventory positions and order submission behavior. We thus group together trader IDs using a cluster analysis approach to detect similarities in behaviour. Specifically, we use the following criteria: the average per day number of securities traded, the average per day per stock number of trades and orders, the average daily order imbalance and trade imbalances per stock, the median order-to-cancel time, the total number of orders and trades submitted in the

500ms after 3:40 p.m., and the average percentage of orders and cancels submitted within 1ms of another trader ID’s order submission. For these nine criteria, $crit^k$, $k = 1, \dots, 9$, we then compute the pair-wise absolute-value distance for traders i and j as follows²⁵

$$Dist_{ij} = \sum_{k=1}^9 \frac{|crit_i^k - crit_j^k|}{crit_i^k + crit_j^k}. \quad (10)$$

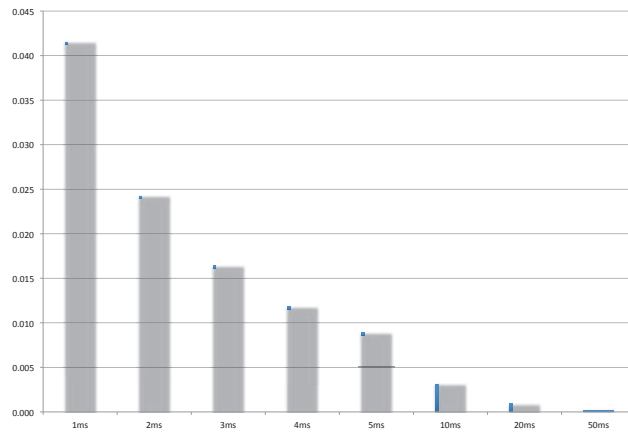
We use visual analysis of the pair-wise distances in an Excel table to identify the clusters. Figure 4 displays the pairwise pair-wise distances based on these criteria, using color-coding to highlight the pairwise distances, where darker colors indicate smaller pair-wise distances. The color-coding in the figure shows that there are groups of securities that have small pair-wise distances. Using a maximum pairwise distance of threshold of .25, we identify four clusters of IDs. Notably, members within a group all have the same underlying broker (but cluster groups have different brokers).

Relation to Market Maker IDs. The identification of market maker IDs in the main text was based on the clustered IDs. As a group, “fast” HFT market makers were on the passive side of 39% of trades and 23% of dollar volume for the Analysis Sample; market maker IDs not identified as fast accounted for 6% of passive transactions and 7% of passive dollar-volume.

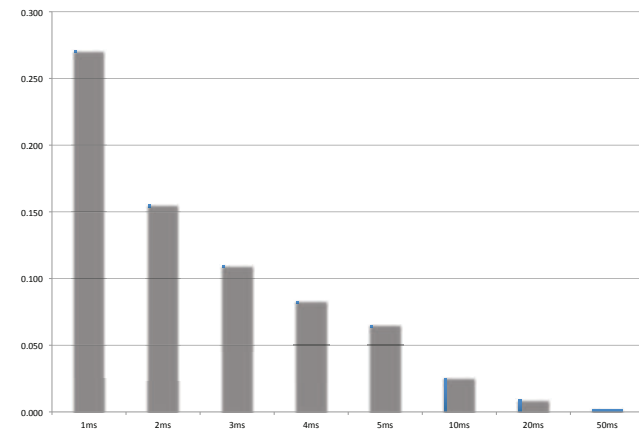
Moreover, a number of non-fast MM IDs were actively providing liquidity on market A before the technological change. At the beginning of May 2015, a number of these IDs (which also had otherwise similar trading characteristics) retired from posting on market A. At the same time, some other IDs showed a significant increase in activity in market A to the point that the number of orders from the disappeared IDs almost coincides with the increase in the number of orders for these HFT IDs. This behavior indicates that these non-HFT IDs may have been part of a market making strategy that used different IDs on different markets which were then were consolidated when markets A’s and B’s system effectively merged. There is thus reason to believe that a fair number of the non-fast market makers are in fact, also HFTs, and that we were just not able to identify them based on our speed criteria.

We thus believe that it is justified to combine HFT and other market making IDs for our analysis of market-making behavior.

²⁵As a convention, when both criteria are 0 for some k for i and j , then we set the distance to 0.



Panel A: Aggressive HFTMM Orders

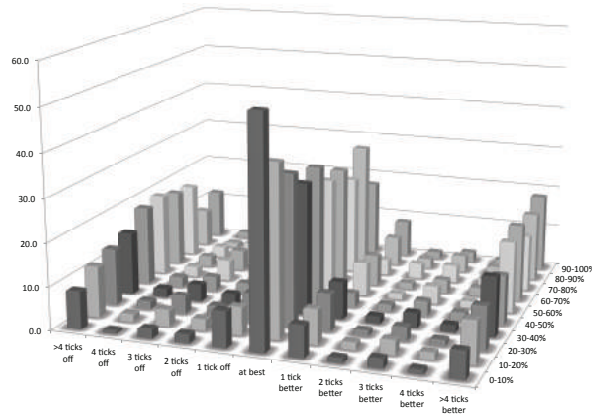


Panel B: HFTMM Cancellations

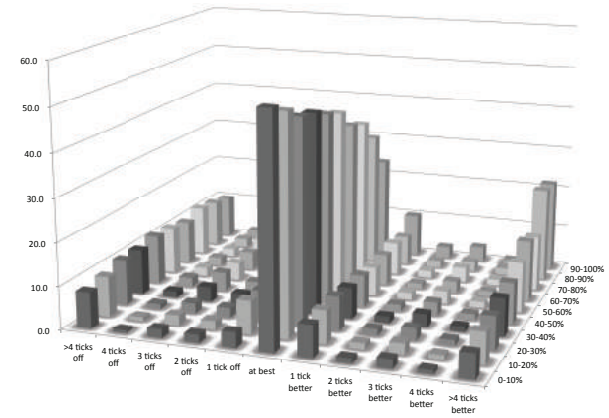
Figure 1

Aggressive order-submission and cancellation rates in response to trades.

The figures plot the same side aggressive order submissions (Panel A) and opposite-side cancellation (Panel B) rates by market-making traders. Table II contains the underlying numbers.



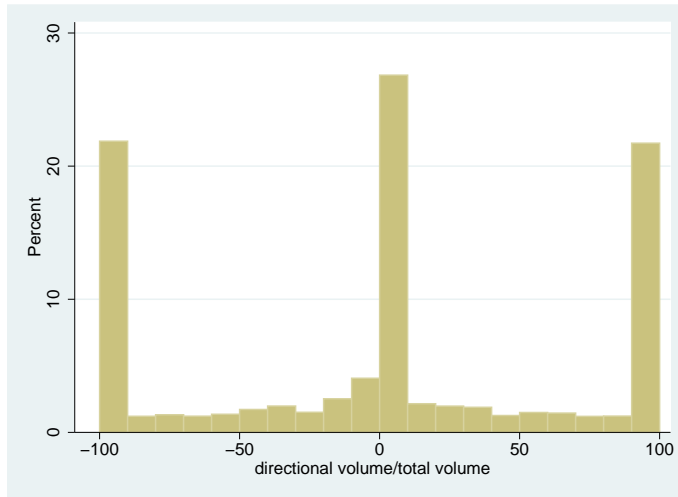
*Panel A: Market Maker Order
Submissions relative to
Trade Imbalances*



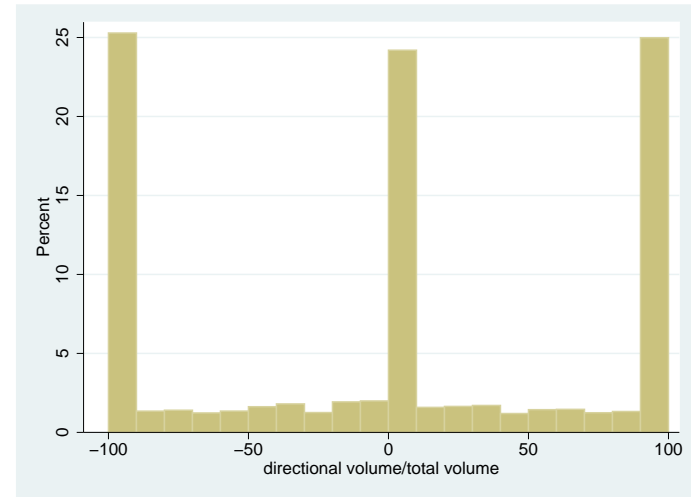
*Panel B: Market Maker Order
Submissions relative to
Own Inventory*

Figure 2
Order Submission Behavior of Market Makers

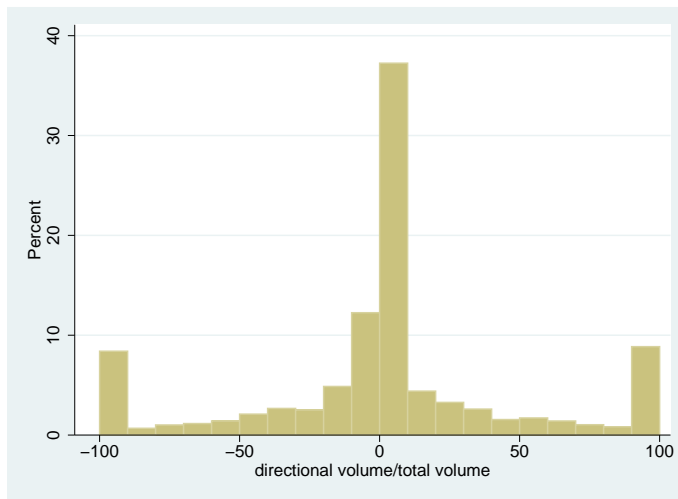
The figure plots the histograms of order submissions by market makers depending on the absolute value of the aggregate marketwise trade imbalance (Panel A) and the absolute value of the market makers inventories (Panel B).



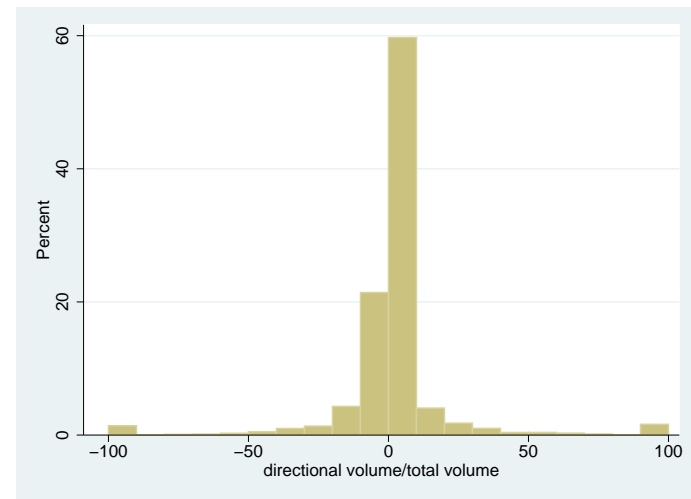
Panel A: All traders



Panel B: non-market-making traders



Panel C: all market makers



*Panel D: non-market-making
traders inventory-managing market makers*

Figure 3
Distribution of By-Trader Inventories

The figure plots the histograms of intra-day cumulative inventories for all traders (Panel A), non-market making traders (Panel B), all market makers (Panel C) and inventory-managing market-makers (Panel D) for the Analysis Sample.

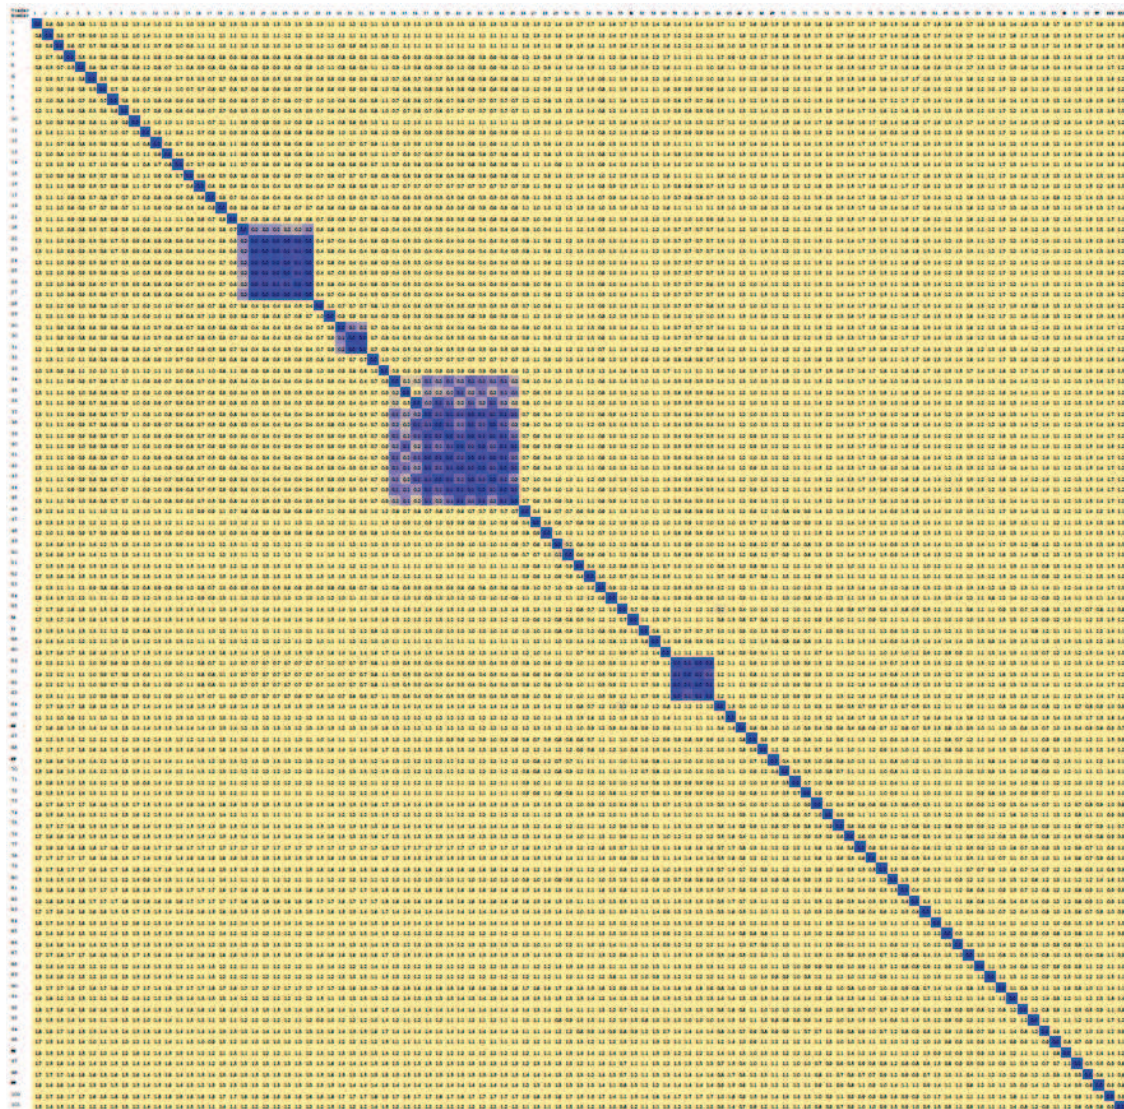


Figure 4
Visualization of Two-way clusters for HFT Groups

The figure represents the table of pairwise distances computed in equation (10) for the 101 trader IDs that are classified as HFT. Smaller numbers are represented by darker cell-shadings.

Table I
Summary Statistics for Trader Type Activities

The table presents summary statistics for trading activities by the different trader types; all numbers are averages per trader in the respective group. Total orders and trades are summed over the entire horizon from January to May 2013 for the Large Sample. Average daily orders are summed across all securities per day.

	All traders	HFT	Buy-side	Retail	non-HFT market makers	HFT market makers	Others
securities traded	14	67	19	42	59	157	9
average daily trades	569	2,246	943	1,549	2,328	15,222	194
total trades	47,434	173,677	84,674	159,970	150,012	1,500,000	11,343
average daily orders	5,130	25,916	2,282	1,396	38,307	318,902	613
total orders	466,299	2,032,000	215,933	144,629	2,509,000	31,160,000	46,415
%orders within 1ms of market event	35	61	37	35	33	82	33
median order to cancel time	274,124	38,494	193,351	550,132	825,670	11,116	274,494
trades at 15:40	71	2,311	38	9	64	3,897	2
orders at 15:40	317	3,353	125	3	690	23,130	14
average daily trade imbalance	91	79	99	92	26	31	92
average daily order imbalance	93	74	99	92	10	11	95

Table II
Summary Statistics Aggressive Orders and Cancellations HFTMMs after Trades

The table provides summary statistics for the total numbers of aggressive order submissions and cancellations after trades. The table below is based on the sum of all observations across the sample for which the local depth is less than the national depth at relevant price (the ask if the trade was a buy, the bid when the trade was a sale); the total is about 1.08M trades. %HFTMM signifies the fraction of orders and cancellations by HFTMMs of all orders and cancellations that were submitted at the time; #HFTMM signifies the number of orders and cancellations. The first four columns collect information for all trades, the last four columns collect the information when an aggressive order occurs within 1 ms (columns 6 and 7) and a cancellation occurs within 1ms (columns 8 and 9).

	% HFTMM local depth<NBBO depth		#HFTMM local depth<NBBO depth		#HFTMM when aggressive occurs		#HFTMM when cancellation occur	
	Aggressive	Cancelled	Aggressive	Cancelled	Aggressive	Cancelled	Aggressive	Cancelled
1ms	24.5	92.9	0.042	0.271	0.300	0.435	0.097	1.426
2ms	31.2	92.7	0.024	0.155	0.114	0.256	0.049	0.568
3ms	34.9	92.5	0.016	0.110	0.067	0.190	0.031	0.340
4ms	37.1	92.3	0.012	0.083	0.045	0.149	0.021	0.234
5ms	38.6	92.1	0.009	0.065	0.032	0.119	0.016	0.173
10ms	40.6	91.1	0.003	0.025	0.010	0.048	0.005	0.060
20ms	39.6	90.9	0.001	0.009	0.003	0.017	0.002	0.020
50ms	37.2	90.2	0.000	0.002	0.001	0.004	0.000	0.004

Table III
Probit Regression on Factors determining Aggressive Order Submissions

The table presents the results from the probit estimation of equation (2), which determines the impact of various trading variables on the probability of observing an aggressive order submission by an HFTMM on another marketplace within 1ms or 5ms of the trade. The figures presented here are the marginal effects as estimated by STATA (for 0/1 variables the number is the effect as the variable switches from 0 to 1). The variables are as follows: **takebook_t** is a dummy for whether or not the trade absorbed the local depth; **SOR_t** is a dummy for (multi-venue) SOR trades; **{vol_t > 200}** is a dummy for whether or not the trade exceeded 200 shares; **|Imb_t|** is the absolute value of the cumulative trade imbalance; **early_t** and **late_t** are dummy variables for the first and last half-hour of trading; **momentum_{t-10,t}** is a dummy that is 1 if the trade is in the same direction of the last 10 price movements; $q \times r_{t-10,t}^{\pm}$ is the trade direction multiplied with the cumulative midpoint return over the last 10 trades; **qspread_t** is the NBBO bid-ask spread in cents; **HFTMM passive_t** is a dummy that is 1 if an HFTMM was on the passive side of the trade; **VXX_t** is the prevailing midprice of the exchange traded note VXX, which tracks the U.S. volatility index VIX; **totaltrans_{it}** is the total number of trades for trader *i* up to time *t*; **%totaltrans_{it}** is the fraction that trader *i* has traded of his daily total trades; **ln(cumval_{it})** is the natural logarithm of the total value traded for this trader at the time of the trade; **time since first_{it}** and **voltime since first_{it}** denote the passage time for trader *i* between the first trade and the current trade (the first trade has value 1); Standard errors are clustered at the security level and they are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels.

Panel A: Aggressive HFTMM Orders within 1ms after all trades

takebook	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	
vol > 200	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)
Imb	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
early	-0.00 (0.00)	-0.00 (0.00)	-0.00* (0.00)	-0.00 (0.00)	-0.00** (0.00)	-0.00** (0.00)	-0.00 (0.00)
late	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00*** (0.00)	0.00** (0.00)	0.00 (0.00)
qspread	-0.00** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00** (0.00)
passive HFTMM	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
VXX	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
momentum	-0.00 (0.00)						-0.00 (0.00)
q×return		-1.00*** (0.17)	-0.99*** (0.17)	-1.00*** (0.17)	-1.02*** (0.17)	-1.01*** (0.16)	
ln(cum \$vol)	0.00 (0.00)	0.00 (0.00)					0.00 (0.00)
total trades _i			-0.00*** (0.00)				
%total trades _i				0.00** (0.00)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.05*** (0.01)
takebook×not SOR				44			0.06*** (0.01)
not takeover×SOR							0.01*** (0.00)
Observations	1,085,577	1,082,882	1,082,882	1,082,882	1,082,882	1,082,882	1,085,577
Pseudo R2	0.0525	0.0528	0.0556	0.0527	0.0535	0.0533	0.0538

Table III (cont'd)

This table is Panel B of Table III; it uses the same variable but restricts attention to trades by traders who trade in a single direction for the entire day and who perform at least 10 trades.

Panel B: Aggressive HFTMM Orders within 1ms after directional trades

takebook	0.04*** (0.00)	0.04*** (0.00)	0.03*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	0.04*** (0.00)	
vol> 200	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)	0.02*** (0.00)
Imb	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
early	-0.00 (0.00)	-0.00 (0.00)	-0.00* (0.00)	-0.00 (0.00)	-0.00** (0.00)	-0.00* (0.00)	-0.00 (0.00)
late	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00** (0.00)	0.00* (0.00)	0.00 (0.00)
qspread	-0.00 (0.00)	-0.00* (0.00)	-0.00* (0.00)	-0.00* (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00* (0.00)
passive HFTMM	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
VXX	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
momentum	-0.00 (0.00)						-0.00 (0.00)
q×return		-0.72*** (0.16)	-0.71*** (0.16)	-0.73*** (0.16)	-0.73*** (0.16)	-0.73*** (0.15)	
ln(cum \$vol)	0.00** (0.00)	0.00** (0.00)					0.00* (0.00)
total trades _i			-0.00*** (0.00)				
%total trades _i				0.00 (0.00)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.06*** (0.00)
takebook×not SOR							0.05*** (0.01)
not takebook×SOR							0.02*** (0.00)
Observations	454,733	454,130	454,130	454,130	454,130	454,130	454,733
Pseudo R2	0.0622	0.0623	0.0641	0.0617	0.0622	0.0621	0.0668

Table III (cont'd)

Panel C: Aggressive HFTMM Orders within 5ms after all trades

takebook	0.15*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	
vol> 200	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)
Imb	-0.00 (0.01)	0.00 (0.01)	-0.01 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.00 (0.01)
early	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)	-0.02*** (0.00)	-0.02*** (0.00)
late	0.00 (0.00)	0.00 (0.00)	-0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
qspread	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
passive HFTMM	-0.00 (0.01)	-0.00 (0.01)	0.00 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.00 (0.01)
VXX	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
momentum	-0.01 (0.01)						-0.01 (0.01)
q×return		-3.59*** (0.64)	-3.60*** (0.64)	-3.60*** (0.64)	-3.65*** (0.64)	-3.64*** (0.64)	
ln(cum \$vol)	0.00 (0.00)	0.00 (0.00)					0.00 (0.00)
total trades _i			-0.00*** (0.00)				
%total trades _i				0.00 (0.00)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.21*** (0.02)
takebook×not SOR							0.18*** (0.02)
not takebook×SOR							0.06*** (0.00)
Observations	1,085,577	1,082,882	1,082,882	1,082,882	1,082,882	1,082,882	1,085,577
Pseudo R2	0.0902	0.0905	0.0949	0.0904	0.0910	0.0910	0.0938

Table III (cont'd)

Panel D: Aggressive HFTMM Orders within 51ms after directional trades

takebook	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	0.14*** (0.01)	
vol > 200	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.10*** (0.01)	0.09*** (0.01)	0.09*** (0.01)	0.09*** (0.01)
Imb	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)
early	-0.01* (0.00)	-0.01* (0.00)	-0.01*** (0.00)	-0.01*** (0.00)	-0.02*** (0.00)	-0.02*** (0.00)	-0.01** (0.00)
late	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.01** (0.00)	0.01 (0.00)	0.00 (0.00)
qspread	-0.00** (0.00)	-0.00** (0.00)	-0.00** (0.00)	-0.00** (0.00)	-0.00** (0.00)	-0.00** (0.00)	-0.00** (0.00)
passive HFTMM	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	-0.00 (0.01)	0.00 (0.01)
VXX	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
momentum	-0.01 (0.01)						-0.01 (0.01)
q × return		-2.58*** (0.60)	-2.58*** (0.61)	-2.60*** (0.60)	-2.60*** (0.60)	-2.62*** (0.60)	
ln(cum \$vol)	0.00* (0.00)	0.00* (0.00)					0.00 (0.00)
total trades _i			-0.00*** (0.00)				
%total trades _i				-0.00 (0.00)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook × SOR							0.22*** (0.02)
takebook × not SOR							0.17*** (0.02)
not takebook × SOR							0.08*** (0.00)
Observations	454,733	454,130	454,130	454,130	454,130	454,130	454,733
Pseudo R ²	0.106	0.106	0.110	0.106	0.106	0.106	0.115

Table IV
Probit Regression on Factors determining HFTMM Cancellations

The table presents the results from the probit estimation of equation (2), which determines the impact of various trading variables on the probability of observing a cancellation by an HFTMM on another marketplace within 1ms or 5ms of the trade. The figures presented here are the marginal effects as estimated by STATA (for 0/1 variables the number is the effect as the variable switches from 0 to 1). The variables are as follows: takebook_t is a dummy for whether or not the trade absorbed the local depth; SOR_t is a dummy for (multi-venue) SOR trades; $\{\text{vol}_t > 200\}$ is a dummy for whether or not the trade exceeded 200 shares; $|\text{Imb}_t|$ is the absolute value of the cumulative trade imbalance; early_t and late_t are dummy variables for the first and last half-hour of trading; $\text{momentum}_{t-10,t}$ is a dummy that is 1 if the trade is in the same direction of the last 10 price movements; $q \times r_{t-10,t}^\pm$ is the trade direction multiplied with the cumulative midpoint return over the last 10 trades; qspread_t is the NBBO bid-ask spread in cents; HFTMM passive_t is a dummy that is 1 if an HFTMM was on the passive side of the trade; VXX_t is the prevailing midprice of the exchange traded note VXX, which tracks the U.S. volatility index VIX; totaltrans_{it} is the total number of trades for trader i up to time t ; $\% \text{totaltrans}_{it}$ is the fraction that trader i has traded of his daily total trades; $\ln(\text{cumval}_{it})$ is the natural logarithm of the total value traded for this trader at the time of the trade; $\text{time since first}_{it}$ and $\text{votime since first}_{it}$ denote the passage time for trader i between the first trade and the current trade (the first trade has value 1); Standard errors are clustered at the security level and they are in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels.

Panel A: HFTMM cancellations within 1ms all trades

takebook	0.07*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	
vol > 200	0.06*** (0.01)	0.06*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.06*** (0.01)
Imb	-0.03** (0.01)	-0.02** (0.01)	-0.03*** (0.01)	-0.03*** (0.01)	-0.03*** (0.01)	-0.03*** (0.01)	-0.02** (0.01)
early	0.00 (0.01)	0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)
late	-0.00 (0.01)	-0.00 (0.01)	0.00 (0.01)	-0.00 (0.01)	0.01 (0.01)	0.00 (0.01)	0.00 (0.01)
qspread	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
passive HFTMM	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.06*** (0.01)	0.06*** (0.01)
VXX	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00* (0.00)
momentum	0.01 (0.01)						0.01 (0.01)
q×return		-0.75 (0.64)	-0.80 (0.65)	-0.80 (0.65)	-0.90 (0.63)	-0.85 (0.63)	
ln(cum \$vol)	0.01*** (0.00)	0.01*** (0.00)					0.01*** (0.00)
total trades _i			0.00 (0.00)				
%total trades _i				0.01*** (0.01)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.13*** (0.03)
takebook×not SOR			48				0.08*** (0.02)
not takebook×SOR							0.11*** (0.01)
Observations	1,085,577	1,082,882	1,082,882	1,082,882	1,082,882	1,082,882	1,085,577
pseudo R2	0.0318	0.0314	0.0296	0.0297	0.0307	0.0309	0.0407

Table IV (cont'd)

This table is Panel B of Table IV; it uses the same variable but restricts attention to trades by traders who trade in a single direction for the entire day and who perform at least 10 trades.

Panel B: HFTMM cancellations within 1ms directional trades

takebook	0.07*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	0.08*** (0.02)	0.07*** (0.02)	0.08*** (0.02)	
vol > 200	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.06*** (0.01)
Imb	-0.03 (0.02)	-0.03 (0.02)	-0.03* (0.01)	-0.03* (0.02)	-0.03* (0.02)	-0.03* (0.02)	-0.03* (0.01)
early	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.02** (0.01)	-0.02* (0.01)	-0.01 (0.01)
late	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.01)	0.02** (0.01)	0.01 (0.01)	0.00 (0.01)
qspread	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
passive HFTMM	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.07*** (0.01)	0.05*** (0.01)	0.07*** (0.01)
VXX	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)	0.01 (0.00)
momentum	0.01 (0.01)						0.01 (0.01)
q×return		-0.81 (0.57)	-0.88 (0.57)	-0.86 (0.57)	-0.86 (0.55)	-0.90 (0.55)	
ln(cum \$vol)	0.01*** (0.00)	0.01** (0.00)					0.00* (0.00)
total trades _i			-0.00*** (0.00)				
%total trades _i				0.01*** (0.00)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.13*** (0.03)
takebook×not SOR							0.08*** (0.02)
not takebook×SOR							0.12*** (0.01)
Observations	454,733	454,130	454,130	454,130	454,130	454,130	454,733
Pseudo R2	0.0375	0.0372	0.0372	0.0364	0.0375	0.0383	0.0477

Table IV (cont'd)

Panel C: HFTMM cancellations within 5ms all trades

takebook	0.19*** (0.04)	0.19*** (0.05)	0.19*** (0.05)	0.20*** (0.05)	0.20*** (0.05)	0.19*** (0.05)	
vol> 200	0.17*** (0.02)	0.17*** (0.02)	0.17*** (0.02)	0.17*** (0.02)	0.17*** (0.02)	0.17*** (0.02)	0.16*** (0.02)
Imb	-0.05** (0.03)	-0.05* (0.03)	-0.06** (0.03)	-0.05** (0.03)	-0.05** (0.03)	-0.06** (0.03)	-0.05** (0.03)
early	-0.01 (0.02)	-0.01 (0.02)	-0.03 (0.02)	-0.02 (0.02)	-0.03* (0.02)	-0.04** (0.02)	-0.02 (0.02)
late	-0.02 (0.01)	-0.02 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	-0.00 (0.01)	-0.01 (0.01)
qspread	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
passive HFTMM	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)	0.15*** (0.02)	0.13*** (0.03)	0.14*** (0.02)
VXX	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)
momentum	0.04 (0.02)						0.04 (0.02)
q×return		-0.95 (1.31)	-1.04 (1.32)	-1.01 (1.32)	-1.16 (1.30)	-1.14 (1.26)	
ln(cum \$vol)	0.02*** (0.01)	0.01*** (0.01)					0.01 (0.00)
total trades _i			-0.00 (0.00)				
%total trades _i				0.01 (0.01)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.31*** (0.05)
takebook×not SOR							0.16*** (0.05)
not takebook×SOR							0.20*** (0.02)
Observations	1,085,577	1,082,882	1,082,882	1,082,882	1,082,882	1,082,882	1,085,577
Pseudo R2	0.0668	0.0659	0.0643	0.0642	0.0657	0.0672	0.0854

Table IV (cont'd)

Panel D: HFTMM cancellations within 5ms directional trades

takebook	0.20*** (0.04)	0.20*** (0.05)	0.20*** (0.05)	0.21*** (0.05)	0.20*** (0.05)	0.20*** (0.05)	
Imb	0.17*** (0.02)	0.17*** (0.02)	0.18*** (0.02)	0.18*** (0.02)	0.18*** (0.02)	0.18*** (0.02)	0.16*** (0.02)
Imb	-0.05 (0.03)	-0.04 (0.03)	-0.04 (0.03)	-0.05 (0.03)	-0.05 (0.03)	-0.05* (0.03)	-0.05* (0.03)
early	-0.02 (0.02)	-0.02 (0.02)	-0.04* (0.02)	-0.03 (0.02)	-0.06*** (0.02)	-0.06*** (0.02)	-0.03 (0.02)
late	-0.01 (0.01)	-0.01 (0.01)	-0.00 (0.01)	0.00 (0.01)	0.03** (0.01)	0.01 (0.01)	-0.00 (0.01)
qspread	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
passive HFTMM	0.16*** (0.03)	0.16*** (0.03)	0.16*** (0.03)	0.16*** (0.03)	0.16*** (0.03)	0.11*** (0.03)	0.15*** (0.03)
VXX	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)	-0.01 (0.00)
momentum	0.04 (0.02)						0.04* (0.02)
q×return		-1.01 (1.43)	-1.13 (1.44)	-1.10 (1.43)	-1.09 (1.37)	-1.23 (1.32)	
ln(cum \$vol)	0.01** (0.01)	0.01** (0.01)					0.01 (0.01)
total trades _i			-0.00*** (0.00)				
%total trades _i				0.00 (0.01)			
volume time _i					-0.00*** (0.00)		
time _i						-0.00*** (0.00)	
takebook×SOR							0.32*** (0.05)
takebook×not SOR							0.15*** (0.05)
not takebook×SOR							0.21*** (0.02)
Observations	454,733	454,130	454,130	454,130	454,130	454,130	454,733
Pseudo R2	0.0749	0.0739	0.0753	0.0728	0.0745	0.0779	0.0952

Table V
Panel Regression for Technology Change: Liquidity

The table presents our results for our panel regression of equation (3) which estimates the effect of the technology change on liquidity variables. We consider four liquidity variables of interest: %HFTMM_{*ti*} is the percentage of trades per stock per day that involve an HFTMM on the passive side; ln(\$Depth) is the time-weighted marketplace depth at the bid and ask; at best is the percentage of the day that the marketplace was at both the best bid and ask price; qspread is the marketplace level time-weighted quoted spread, measured in basis points of the midquote. The explanatory variables are dummies for the marketplaces interacted with a dummy for whether the observation is for the time after the technology change. VXX_{*t*} is the trade-weighted average of the volatility ETF VXX. All regressions include stock fixed effects. Standard errors are double-clustered by security and date, and they are presented in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels.

	%HFTMM _{<i>ti</i>}	ln(\$Depth)	at best	qspread
market A	-7.56** (3.11)	-0.25*** (0.06)	-0.04* (0.02)	73.31 (54.75)
market B	3.44* (2.04)	-0.12** (0.06)	0.02 (0.03)	69.37 (54.45)
market D	-1.96 (4.47)	-0.24** (0.11)	-0.06*** (0.02)	-43.98 (358.99)
market C	-6.92*** (2.34)	-0.14** (0.07)	-0.03 (0.03)	93.77 (60.23)
market E	-2.29 (4.19)	-0.08* (0.05)	0.01 (0.01)	47.96 (140.44)
market F	-3.92 (2.55)	-0.11*** (0.04)	0.02** (0.01)	63.31 (58.16)
VXX	-0.10 (0.42)	-0.03*** (0.01)	0.00* (0.00)	22.87 (19.07)
Observations	4,088	4,088	4,088	4,088

Table VI
Panel Regression for by-Market Aggressive Orders

The table presents our results for our panel regression of equation (3) which estimates the effect of the technology change on after-trade aggressive orders by HFTMMs. Panel A presents the estimates for aggressive orders on other markets within 1ms of the trade, Panel B for 5ms reactions. We consider four variables of interest: $\#HFTMM$ counts the number of trades with HFTMM aggressive orders after the trade; $\#HFTMM$ orders counts the total number of HFTMM aggressive orders after the trade; $\%HFTMM$ is the percentage of HFTMM aggressive orders of all aggressive orders; $\#nHFTMM$ orders is the total number of non-HFTMMs aggressive orders. The explanatory variables are dummies for the marketplaces interacted with a dummy for whether the observation is for the time after the technology change. VXX_t is the trade-weighted average of the volatility ETF VXX. We further control for the fraction of all trades that involve an HFTMM on the passive side. All regressions include stock fixed effects. Standard errors are double-clustered by security and date, and they are presented in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels.

	<i>Panel A: within 1ms</i>				<i>Panel B: within 5ms</i>			
	$\#HFTMM$	$\#HFTMM$ orders	$\%HFTMM$	$\#nHFTMM$ orders	$\#HFTMM$	$\#HFTMM$ orders	$\%HFTMM$	$\#nHFTMM$ orders
market A	2.11 (1.98)	1.49 (2.49)	2.79 (3.29)	16.71* (9.76)	13.65*** (4.76)	2.89 (1.85)	9.89*** (3.41)	-1.82 (5.96)
market B	-0.99 (4.75)	-7.14 (6.81)	7.74 (5.76)	-13.12 (13.02)	2.83 (11.17)	-6.10 (5.02)	8.52** (3.97)	-10.07 (7.63)
market D	-1.38 (1.36)	-2.35 (1.90)	8.62** (4.00)	-1.77 (5.38)	-3.43 (3.15)	-1.67 (1.55)	11.65*** (3.85)	-1.91 (2.86)
market C	-10.14** (3.95)	-13.60*** (5.11)	-16.90*** (4.72)	11.86 (10.19)	-16.72* (8.97)	-7.46** (3.56)	-14.35*** (4.03)	8.02 (6.01)
market E	-1.64 (1.14)	-2.75* (1.65)	-12.50** (5.02)	5.72 (4.02)	-2.68 (2.72)	-1.50 (1.29)	-10.63** (4.70)	6.19** (3.15)
market F	-1.43 (1.85)	-2.58 (2.70)	-20.91*** (4.07)	10.18** (4.39)	2.24 (3.46)	0.84 (1.74)	-14.94*** (3.87)	10.62*** (2.33)
%trades with passive HFTMM	0.01 (0.01)	0.01 (0.01)	-0.01 (0.03)	0.04 (0.04)	0.03 (0.03)	0.01 (0.01)	0.02 (0.02)	0.01 (0.02)
VXX	-0.40 (0.27)	-0.67* (0.39)	0.12 (0.52)	0.94 (1.20)	-0.88 (0.64)	-0.32 (0.26)	-0.23 (0.66)	0.50 (0.52)
Observations	4,088	4,088	3,770	4,088	4,088	4,088	3,957	4,088

Table VII
Panel Regression for by-Market Cancellations

The table presents our results for our panel regression of equation (3) which estimates the effect of the technology change on after-trade cancellations by HFTMMs. Panel A presents the estimates for cancellations on other markets within 1ms of the trade, Panel B for 5ms reactions. We consider four variables of interest: $\#HFTMM$ counts the number of trades with HFTMM cancellations after the trade; $\#HFTMM$ orders counts the total number of HFTMM cancellations after the trade; $\%HFTMM$ is the percentage of HFTMM cancellations of all cancellations; $\#nHFTMM$ orders is the total number of non-HFTMMs cancellations. The explanatory variables are dummies for the marketplaces interacted with a dummy for whether the observation is for the time after the technology change. VXX_t is the trade-weighted average of the volatility ETF VXX. We further control for the fraction of all trades that involve an HFTMM on the passive side. All regressions include stock fixed effects. Standard errors are double-clustered by security and date, and they are presented in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels.

	<i>Panel A: within 1ms</i>				<i>Panel B: within 5ms</i>			
	$\#HFTMM$	$\#HFTMM$ <i>cancels</i>	$\%HFTMM$	$\#nHFTMM$ <i>cancels</i>	$\#HFTMM$	$\#HFTMM$ <i>cancels</i>	$\%HFTMM$	$\#nHFTMM$ <i>cancels</i>
market A	23.43* (12.74)	33.75 (22.85)	-1.29 (3.34)	-1.42 (2.61)	81.84*** (26.80)	43.11* (22.78)	0.43 (2.56)	-0.97 (2.07)
market B	-10.84 (19.10)	-24.70 (36.06)	-10.11*** (2.10)	8.98*** (2.84)	-19.83 (31.06)	-43.78 (39.48)	-7.49*** (2.22)	8.14** (3.72)
market D	8.90 (7.68)	17.29 (14.84)	-1.05 (2.91)	1.80** (0.82)	15.46 (11.59)	13.34 (10.87)	3.64 (2.45)	0.03 (0.92)
market C	-29.61* (16.77)	-48.99 (30.69)	-8.46** (3.63)	-4.03 (2.46)	0.18 (25.04)	-12.96 (24.02)	-1.67 (2.71)	-2.73 (2.38)
market E	-0.43 (6.07)	0.83 (11.91)	-4.97** (2.49)	1.49*** (0.26)	7.04 (7.16)	3.48 (6.21)	-1.26 (2.55)	1.50** (0.63)
market F	-3.56 (7.49)	-9.17 (15.85)	1.71 (2.56)	0.81*** (0.31)	12.42 (8.77)	3.37 (9.42)	3.48* (2.01)	1.35** (0.66)
%trades with	0.13***	0.21***	0.06***	0.01*	0.27**	0.24***	0.08***	0.02
passive HFTMM	(0.05)	(0.07)	(0.02)	(0.01)	(0.11)	(0.07)	(0.02)	(0.01)
VXX	0.53 (1.81)	2.03 (3.55)	-1.17** (0.54)	0.55*** (0.11)	2.05 (2.22)	1.78 (1.46)	-0.80* (0.47)	0.57*** (0.20)
Observations	4,088	4,088	3,721	4,088	4,088	4,088	3,988	4,088

Table VIII
Order Posting Relative to Best Prices

The table depicts the order posting behavior of market makers relative to locally best posted prices that is plotted in Figure 2. Negative columns labels stand for “worse” than the best prices, positive numbers are for improvements of the best prices. Each figure represents the average fraction of their orders that inventory-managing market makers have submitted relative to the best locally posted price on the marketplace where their order was posted. The number of observations are per ticker, date, and volume-interval.

Panel A: Order posting vs. trade imbalance

Imbalance	≤ -5	-4	-3	-2	-1	at best	+1	+2	+3	+4	$\geq +5$	avg. imbalance	Obs
0-10%	8.5	0.7	2.0	1.9	3.5	52.8	7.4	1.2	2.2	0.9	5.7	4.8	33422
10-20%	9.3	0.9	2.4	2.1	7.8	50.4	7.7	1.4	2.5	1.0	7.2	14.3	19195
20-30%	10.6	1.2	2.4	2.3	7.3	47.6	8.1	1.5	2.8	1.2	7.8	24.0	8354
30-40%	10.6	1.5	3.5	2.8	3.5	46.6	7.3	1.6	3.2	1.3	8.9	34.1	2787
40-50%	11.5	2.1	4.3	3.0	8.5	44.6	7.5	1.5	3.1	1.3	9.7	44.4	1108
50-60%	11.2	2.0	4.1	3.7	7.1	43.1	6.7	1.7	3.0	1.4	11.6	54.3	512
60-70%	10.3	2.0	4.3	3.3	7.1	38.5	7.3	1.2	2.6	1.2	14.3	64.2	213
70-80%	11.7	1.9	5.1	3.9	8.1	37.0	7.8	1.6	3.4	1.3	12.8	74.4	131
80-90%	10.9	2.0	4.4	4.1	5.4	31.8	7.0	1.6	2.3	1.5	22.2	84.4	80
90-100%	10.0	2.0	4.4	2.2	4.5	23.7	10.2	2.8	3.6	1.2	21.6	97.1	103

Panel B: Order posting vs. market maker inventory

Inventory	≤ -5	-4	-3	-2	-1	at best	+1	+2	+3	+4	$\geq +5$	avg. MM inventory	Obs
0-10%	8.8	0.8	2.2	2.0	8.4	52.2	7.5	1.3	2.3	1.0	6.3	2	60103
10-20%	11.8	1.7	3.7	2.6	6.4	39.6	8.0	1.6	3.2	1.6	9.8	14	3361
20-30%	13.4	1.8	4.1	2.6	6.7	34.9	8.8	1.4	3.3	1.3	10.0	24	1103
30-40%	14.7	2.2	4.2	2.8	6.9	30.7	8.7	1.6	3.4	1.8	13.8	34	486
40-50%	18.5	2.3	3.4	2.7	7.6	32.2	3.3	1.5	3.3	1.7	11.4	44	215
50-60%	19.5	2.3	5.0	3.0	6.5	27.1	7.9	1.4	4.2	1.4	16.5	53	141
60-70%	18.1	2.5	4.8	3.2	6.2	27.4	7.2	1.3	3.0	2.0	18.0	63	113
70-80%	17.5	2.4	4.1	3.3	5.6	23.1	3.7	3.6	4.0	1.6	13.0	74	54
80-90%	8.9	1.1	3.5	3.1	6.4	29.2	6.8	2.3	4.0	1.9	16.1	84	30
90-100%	11.8	1.5	2.9	1.6	3.4	17.8	8.6	1.5	2.4	1.3	18.5	100	328

Table IX
Order Submission by Market Makers in Response to Demand/Supply Pressure

The table estimates the effect of trade imbalances on the order submission behavior of market makers depicted by equation (8). *Order submission imbalance_t* is the difference of buy and sell order volume relative to all order volume in volume interval t , *trade imbalance_{t-1}* is the difference of buyer- and seller-initiated volume relative to total volume in volume interval $t - 1$, *cumulative trade imbalance_{t-1}* is the cumulative trade imbalance since the beginning of the trading day. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	order submission imbalance _t			
	Inventory-managing market makers		all market makers	
trade imbalance _{t-1}	0.073*** (0.022)		0.036* (0.020)	
cumulative trade imbalance _{t-1}		0.122*** (0.022)		0.073*** (0.021)
Observations	61,107	61,274	61,348	61,574

Table X
Prices Submitted by Market Makers in Response to Demand/Supply Pressure

The table estimates the effect of trade imbalances on the price of orders that market makers submit; the underlying equation is (5) Variable $\Delta^{\text{bid}}\text{vwap}_{t,b}^o$ is the distance (in bps) of the volume weighted average order price for buy orders relative to the last bid price in volume interval $t - 1$, and $\Delta^{\text{ask}}\text{vwap}_{t,s}^o$ is the distance (in bps) of the volume weighted average order price for sell orders relative to the last ask price in volume interval $t - 1$. *Trade imbalance* $_{t-1}$ is the difference of buyer- and seller-initiated volume relative to total volume in volume interval $t - 1$, *cumulative trade imbalance* $_{t-1}$ is the cumulative trade imbalance since the beginning of the trading day. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	bid prices: $\Delta^{\text{bid}}\text{vwap}_{t,b}^o$			offer prices: $\Delta^{\text{ask}}\text{vwap}_{t,s}^o$		
<i>Panel A: Inventory-managing market makers</i>						
trade imbalance $_{t-1} \times$ buy	-0.000 (0.000)		0.000 (0.000)	-0.000 (0.002)		-0.000 (0.002)
trade imbalance $_{t-1} \times$ sell	-0.001*** (0.000)		-0.000 (0.000)	0.000 (0.001)		0.001 (0.001)
cumulative trade imbalance $_{t-1} \times$ buy		-0.007*** (0.002)	-0.007*** (0.002)		-0.001 (0.006)	-0.001 (0.006)
cumulative trade imbalance $_{t-1} \times$ sell		-0.008*** (0.002)	-0.007*** (0.002)		-0.007 (0.009)	-0.008 (0.009)
Int $_t$	0.002*** (0.000)	0.002*** (0.000)	0.002*** (0.000)	0.001 (0.002)	0.001 (0.002)	0.001 (0.002)
Observations	60,626	60,761	60,626	60,537	60,662	60,537
<i>Panel B: All market makers</i>						
trade imbalance $_{t-1} \times$ buy	-0.000 (0.000)		0.000 (0.000)	0.000 (0.001)		0.000 (0.001)
trade imbalance $_{t-1} \times$ sell	-0.001** (0.000)		-0.000* (0.000)	0.000 (0.001)		0.000 (0.001)
cumulative trade imbalance $_{t-1} \times$ buy		-0.004*** (0.001)	-0.004*** (0.001)		0.001 (0.003)	0.001 (0.002)
cumulative trade imbalance $_{t-1} \times$ sell		-0.003*** (0.001)	-0.003*** (0.001)		-0.003 (0.004)	-0.003 (0.004)
Int $_t$	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)
Observations	61,113	61,293	61,113	61,122	61,294	61,122

Table XI
Order Location in the Book in Response to Demand/Supply Pressure

The table estimates the effect of trade imbalances on the location in the order book of market maker orders. $|trade\ imbalance_{t-1}|$ is the absolute value of the trade imbalance, i.e. the absolute value of the difference of buyer- and seller-initiated volume relative to total volume in volume interval $t-1$, $|cumulative\ trade\ imbalance_{t-1}|$ is the absolute value of the cumulative trade imbalance since the beginning of the trading day. $\geq +2, +1$ depict improvements of the local BBO, $\leq -2, -1$ depict order submissions at worse prices than the local BBO. The explanatory variables are the fractions of market maker orders of all market maker orders that are submitted at the respective prices. All orders considered are passive and visible, i.e. they do not include immediate-or-cancel, dark, or marketable orders. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	$\geq +2$	$+1$	at best	-1	≤ -2	$\geq +2$	$+1$	at best	-1	≤ -2
$ imbalance_{t-1} $	-0.019*** (0.006)	-0.007** (0.003)	-0.037*** (0.005)	-0.003* (0.002)	-0.008 (0.005)					
$ cum\ imbalance_{t-1} $						0.093*** (0.022)	0.017*** (0.005)	-0.156*** (0.026)	-0.003 (0.008)	0.068*** (0.023)
Int_t	-0.035*** (0.006)	-0.002 (0.004)	0.043*** (0.012)	-0.004* (0.003)	-0.034*** (0.009)	-0.041*** (0.008)	-0.010* (0.006)	-0.019 (0.017)	-0.012*** (0.002)	-0.044*** (0.012)
Observations	62,581	62,581	62,531	62,581	62,581	65,909	65,909	65,909	65,909	65,909

Table XII
Do Market Makers Make Markets?

The table estimates the relation of market-wide order imbalances and market market inventories. If market makers indeed make markets, their inventories should be negatively related to aggregate order flow (e.g., if the market as a whole is dominated by buyers, then liquidity-providing market makers would be the sellers); the underlying equation is (4). The market makers inventory is computed as the difference of the cumulative buy volume less the sell volume relative to the market makers' total volume. We estimate the relationship both contemporaneously and for lagged trade imbalances. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	inventory _t			
imbalance _t	-0.830*			
	(0.499)			
cumulative imbalance _t		-0.061***		
		(0.021)		
imbalance _{t-1}			-0.006	
			(0.004)	
cumulative imbalance _{t-1}				-0.045***
				(0.016)
Observations	62989	66822	62,542	65,802

Table XIII
Order Submission by Market Makers in Response to Inventory Accumulation

The table estimates the effect of market maker inventories on the order submission behavior of market makers depicted by equation (7). *Order submission imbalance_t* is the difference of buy and sell order volume relative to all order volume in volume interval t , *inventory_{t-1}* is the difference of the market makers cumulative (since the beginning of the trading day) buying- and selling-volume relative to their total volume by volume interval $t - 1$. *Int_t* is a control for the volume interval. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	order imbalance _t	
inventory _{t-1}	-0.248*** (0.030)	-0.248*** (0.030)
Int _t		-0.001 (0.006)
Observations	61,256	61,652

Table XIV
Prices Submitted by Market Makers in Response to Inventory Accumulation

The table estimates the effect of market makers' inventories on the price of orders that market makers submit. Variable $\Delta^{\text{bid}}\text{vwap}_{t,b}^o$ is the distance (in bps) of the volume weighted average order price for buy orders relative to the last bid price in volume interval $t - 1$, and $\Delta^{\text{ask}}\text{vwap}_{t,s}^o$ is the distance (in bps) of the volume weighted average order price for sell orders relative to the last ask price in volume interval $t - 1$. Inventory_{t-1} is the difference of market makers' cumulative buying- and selling volume relative to their total volume since the beginning of the trading day. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	bid prices: $\Delta^{\text{bid}}\text{vwap}_{t,b}^o$	offer prices: $\Delta^{\text{ask}}\text{vwap}_{t,s}^o$
inventory _{t-1} × buy	-0.016*** (0.003)	0.005 (0.004)
inventory _{t-1} × sell	-0.011*** (0.004)	0.013*** (0.005)
Observations	60,709	60,610

Table XV
Order Location in Response to Inventory Accumulation

The table estimates the effect of market makers' inventories on the location of their orders in the order book. $|inventory_{t-1}|$ is the absolute value of the market makers' cumulative inventory by the end of volume interval $t - 1$. $\geq +2, +1$ depict improvements of the local BBO, $\leq -2, -1$ depict order submissions at worse prices than the local BBO. The explanatory variables are the fractions of market maker orders of all market maker orders that are submitted at the respective prices. All orders considered are passive and visible, i.e. they do not include immediate-or-cancel, dark, or marketable orders. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	$\geq +2$	$+1$	at best	-1	≤ -2
$ inventory_{t-1} $	0.085*** (0.019)	-0.016** (0.008)	-0.356*** (0.050)	-0.034*** (0.008)	0.044* (0.024)
Int_t	-0.040*** (0.008)	-0.012* (0.006)	-0.028 (0.013)	-0.014*** (0.002)	-0.044*** (0.012)
Observations	65,934	65,934	65,934	65,934	65,934

Table XVI
Inventory Management and Liquidity

The table estimates the effect of lagged market maker inventories (and the associated contemporaneous inventory management) on liquidity, measured by time-weighted bid-ask spreads (both in cents and in basis points of the prevailing midquote), the range measure (the highest less the lowest midquote per volume interval relative to the time-weighted average midquote), and the absolute value of the midquote return. $|Inventory_{t-1}|$ is the absolute value of the market makers' inventory. Panel A estimates the effect for the pooled sample of all securities, Panel B estimated the effect when the sample is split into terciles (low, medium and high) of competition for liquidity provision. All specifications include stock fixed effects. Standard errors are in parentheses and are clustered by date and security. * indicates significance at the 10% level, **at the 5% level, and *** at the 1% level.

	quoted cents	bid-ask spread bps	quote range	return $ r_{t-1,t} $
<i>Panel A: Full Sample</i>				
$ inventory_{t-1} $	-0.811 (0.646)	-5.682 (5.186)	0.216*** (0.032)	-2.192 (4.147)
Observations	61,852	65,881	65,802	65,873
<i>Panel B: Split Sample</i>				
$ inventory_{t-1} \times \text{low competition}$	0.077 (0.052)	0.135 (0.102)	0.155*** (0.019)	0.066 (0.175)
$ inventory_{t-1} \times \text{medium competition}$	-3.600 (3.067)	-23.788 (20.148)	0.353*** (0.086)	-11.017 (18.182)
$ inventory_{t-1} \times \text{high competition}$	-0.702** (0.348)	-1.618** (0.779)	0.217*** (0.037)	0.566 (1.518)
Observations	65,881	65,881	65,802	65,873