Waseda University
Institute for Business and Finance

Working Paper Series

WBF-19-003：June 2019

Trader Classification by Cluster Analysis: Interaction between HFTs and Other Traders

早稲田大学商学部

五島　圭一

早稲田大学ビジネス・ファイナンス研究センター

戸辺玲子

早稲田大学大学院経営管理研究科

宇野　淳

早稲田大学 ビジネス・ファイナンス研究センター

Trader Classification by Cluster Analysis: Interaction between HFTs and Other Traders

Keiichi Goshima,* Reiko Tobe,† Jun Uno ‡

June 7, 2019

Abstract

We classify groups of traders with similar trading characteristics through cluster analysis. This is a first attempt, as far as we know, to use cluster analysis to identify trader groups. We use the inventory ratio, order cancellation ratio, order frequency, and number of stocks per trading server as proxies of behavioral characteristics. The combination of these four variables successfully separates a group of traders matching the key trading characteristics. We identify two groups that satisfy the high-cancellation and low-inventory conditions of high-frequency trading market makers (HFT-MMs). These HFT-MMs are differentiated from other fast traders by their avoidance of market orders and assignment of very small numbers of stocks per trading server to minimize latency. A comparative analysis of trading behaviors between calm and volatile periods using order submission data covering all market participants in ten-seconds intervals finds that HFT-MMs' trading decisions are triggered by marketwide order flow and their inventory changes, but not price information. In the fast-moving market, the magnitude of HFT-MMs' reactions is amplified and the time until a limit order is canceled shortens almost one-third the time under normal market conditions. Thus HFT-MMs remain liquidity providers in the market during both periods.

Keywords: High-frequency trading, cluster analysis, order submission, order flow, inventory

JEL classification: C54; E43; E52; E58; G12; G14

*Waseda University (k1.goshima@gmail.com)

†Waseda University (r.tobe@aoni.waseda.jp)

‡Waseda University (juno@waseda.jp)

1 Introduction

In stock markets worldwide, investors who trade at very high speeds, called high-frequency traders (HFTs), have a dominant presence. Speed has always been of the essence in financial markets. Current HFTs invest in trading infrastructure that ensures their access to fast and high-frequency trading (Biais et al. (2015); Brogaard et al. (2015)).[1] They can respond to events with time stamps on the order of milliseconds, if not microseconds. The latest software and infrastructure developments allow trading firms to use cutting edge algorithmic trading technology at microsecond and millisecond levels. A deep understanding of their trading behavior is therefore crucial in current financial markets, where HFTs provide most of the liquidity (Hosaka (2014)).

Although empirical research on HFTs is rapidly growing, it is largely based on HFT data sets that provide limited coverage of HFT activity and usually not at the account level. This prevents researchers from identifying the series of actions undertaken by individual HFTs (U.S. Securities and Exchange Commission (SEC), 2014). When account-level data are provided, the identification of HFTs is mostly based on the screening of metrics such as the cancellation-to-order ratio (COR), with fairly arbitrary classification thresholds. In this study, we classify traders into groups with hierarchical agglomerative clustering which belongs to unsupervised learning. Cluster analysis classifies traders with similar trading characteristics into the same group. We are not aware of any other studies which use cluster analysis to identify HFTs.[2] This methodology is attractive because it does not require trading metrics thresholds to create subgroups of traders with similar behavioral characteristics.

We select four proxy variables that represent the behavioral characteristics of HFTs well for our cluster analysis: the inventory ratio, the order cancellation ratio, the order frequency, and the number of stocks per server. The first two variables are selected according to the HFT literature, which acknowledges high cancellation and low inventory as key HFT characteristics. The latter two variables indicate the importance of each trader's ability to trade at high frequencies and super-fast speed, respectively. These four variables reveal the underlying principles of investment objectives and order submission strategies that are determined in advance and do not change daily with market conditions. Thus, the combination of these four variables can guide us in finding neighboring entities such as HFTs.

We use order submission data provided by the Tokyo Stock Exchange (TSE). This data set covers all of TSE's electronic messages, including order submission times, types,

[1]Brogaard et al. (2015) show that liquidity providers are willing to pay for premium colocation services that allow them to collocate their servers near exchange matching engines for faster transmission speeds.

[2]Yang et al. (2012) use inverse reinforcement learning to characterize trader behavior and Aït-Sahalia and Saglam (2013) use principal component analysis. However, as far as we know, no attempt has yet been made to use cluster analysis to identify HFTs.

and sizes, and various order execution contingencies, such as "Immediate or Cancel". A key feature of our data set is the inclusion of a virtual server (VS) id and an order identification (OI).[3] The VS is a unique identification number for each TSE customer that enables us to track series of individual traders' actions. We compare trading behaviors in normal and volatile periods because theoretical and empirical studies find that the participation of HFTs declines during times of market stress (Easley et al. (2012); Aït-Sahalia and Saglam (2013); Kirilenko et al. (2017)). This is a crucial issue in terms of market liquidity, since markets rely on high-frequency trading market makers (HFT-MMs) as major liquidity providers. We designate the three days from May 20 to May 22, 2013, as a normal period and the two days from May 23 to 24, 2013, as a volatile period.[4]

We find that two trader groups comprised by 12 trading desks meet the criteria for high-frequently trading liquidity providers, that is, HFT-MMs. In the normal period, HFT-MMs as a group have an average inventory ratio of 9.4% and a cancellation ratio of 83%. These HFT-MMs are differentiated from other fast traders by their exclusive usage of limit orders and assignment of very small numbers of stocks per trading server to minimize latency. They assign seven stocks per VS and submit an average of 122,620 orders per day. Limiting small number of stocks per server is a key feature of HFTs in TSE where the highest level of order/execution link allows a user to submit orders with a maximum of 60 messages per second. In order to submit multiple orders to same stock within a minute, HFTs tend to assign small number of stocks per sever. There are other four groups comprised by 24 trading desks whose order submission frequency is equivalent to that of HFT-MMs. They submit 118,807 orders per day, very close to the order submission frequency of HFT-MMs, but their inventory ratio is 32.7% and their cancellation ratio is 59%. We call these opportunistic high-frequently traders (HFT-OPs).[5] We further define two slower groups as mid-frequency traders (MFTs) and low-frequency traders (LFTs).

The speed advantage of HFTs in terms of information processing and trading is emphasized in the more recent HFT-specific theoretical literature. Higher speeds allow HFTs to react more quickly to public news than other traders can. The other advantages of HFT are the ability to monitor and analyze instantaneously a great amount of information, including the order flow, price/quote information, and changes in inventory.[6] Following the literature, we test whether it is order flow or price/quote information that triggers HFT-MMs' order submissions and whether their decisions change when the market becomes

[3]A VS is a logical communication path established between the TSE's system (matching engine) and a user's system.

[4]See Figure 1. On May 23, 2013, the Nikkei stock average went up and then down by more than 1,000 yen, nearly 7.3%, during the trading session.

[5]e.g., Hagströmer and Nordén (2013); SEC (2014); Kirilenko et al. (2017)

[6]e.g., Cespa and Foucault (2011); Jovanovic and Menkveld (2011); Aït-Sahalia and Saglam (2013); Menkveld (2013); Biais et al. (2015); Gerig (2015); Foucault et al. (2016)

volatile.[7] In the empirical investigation, we analyze HFT-MMs' limit order submissions and cancellations. Information on the sell (buy) orders of other market participants is important in HFT-MMs' buy (sell) limit order decisions, indicating an opportunity to provide liquidity under normal market conditions but an increase in inventory risk in a volatile market. Marketwide order imbalance between buy and sell orders is another risk indicator that can significantly impact their inventory. We use index futures mid-quotes as price/quote information.

We estimate a structural vector autoregression (SVAR) model with 6 lags to analyze HFTs' order submission behavior in conjunction with marketwide order flow and quote changes. We focus on stocks traded by HFT-MMs, totaling 479 stocks, with a data set that includes the entire order submission history (new entries, modifications, order fills, and cancellations) and details of the order contents (number of shares, limit price, and order contingencies such as "Immediate or Cancel"). We obtain intraday data from NEEDS tick data for stock index futures.

The impulse response function (IRF) results indicate that information on the sell (buy) orders of other participants, is one of the most important for HFT-MMs' trading decisions on cancellations. HFT-MMs' inventory change also impacts HFT-MMs' trading decisions, but quote changes of stock index futures market do not have a significant impact. The results indicate HFT-MMs react to order flow information rather than to price/quote information. Their reactions to order flow does not change in the volatile period compared to the normal period except response to inventory. In the volatile period, HFT-MMs reduce the duration in which they leave unfilled limit orders in the market. This shows HFT-MMs adjust their limit order management to maintain proper inventory levels in a fast-moving market.

The next section reviews the literature and Section 3 describes the data. Section 4 explains the cluster analysis classification method, Section 5 describes the hypotheses, and Section 6 reports the results of the empirical analyses. Section 7 concludes the paper.

2 Literature

The literature on HFTs is growing rapidly and is well summarized by Menkveld (2013), Chordia et al. (2013), Jones (2013), Biais and Foucault (2014), and O'Hara (2015). Our study is related to two different strands of the literature: how to identify traders engaging in high-frequency and fast trading and how their trading decisions are related to the behaviors of other market participants. We discuss each of these issues in detail below.

First, our study is motivated by SEC (2010) and SEC (2014). In SEC (2014), it is noted that "the lack of a clear definition of HFT complicated the regulatory review of

[7] The theoretical model of Aït-Sahalia and Saglam (2013) suggests HFT-MMs reduce their activity in volatile periods. Easley et al. (2012) and Kirilenko et al. (2017) find empirical evidences.

market structure. The diversity of HFT highlights the importance of exercising care when using metrics to define HFT and proxies to associate with HFT. In particular, metrics and proxies that are based directly or indirectly on passive order book activity (such as message rates or cancellation rates) may have the effect of excluding a large volume of aggressive HFT activity. In addition, narrow metrics for other aspects of HFT activity, such as intraday and end-of-day inventory levels, can exclude a large volume of HFT activity."

This is also constrained by data sets on which empirical studies on HFT rely. Most HFT studies focus on HFTs as a group because the HFT databases provided by NASDAQ and Euronext Paris do not provide information at the account level.[8] The coverage of HFTs depends solely on how each exchange classifies traders. For example, NASDAQ uses its knowledge about an account and gathers information on the account's business type to determine whether it is an HFT or not. NASDAQ's classification does not include the HFT-type activities conducted by the trading desks of large brokerage firms (SEC (2014)). In the May 2010 Flash Crash, Kirilenko et al. (2017) use CME Group's E-mini data set to identify the unique trading behavior of HFTs based on inventory and volume patterns. Kirilenko et al. (2017) define opportunistic traders who follow a variety of arbitrage trading strategies, including cross-market arbitrage, short securities, statistical arbitrage, and news arbitrage. Hagströmer and Nordén (2013) use trader identifiers and knowledge of member activities provided by NASDAQ OMX Stockholm to manually classify firms as HFT. The authors further divide classified HFTs into "market-making" and "opportunistic."

Even if account-level data are provided, the identification of HFTs is mostly based on screening with a couple of metrics, such as COR. The thresholds of the metrics used for classification are fairly arbitrary. We therefore use agglomerative cluster analysis to search for hierarchical clusters of traders. We are unaware of any other attempts at using cluster analysis to identify HFTs.

Second, the speed advantage of HFTs in terms of information processing and trading is emphasized in the more recent HFT-specific theoretical literature. The higher speeds allow HFTs to react more quickly to public news than other traders can (e.g., Biais and Woolley (2011); Jovanovic and Menkveld (2011); Foucault et al. (2016)). Jovanovic and Menkveld (2011) have developed a model where information asymmetry can generate both beneficial and deleterious effects. Biais and Woolley (2011) suggest that fast traders increase negative externalities and thus adverse selection, crowding out slower traders. The theoretical paper that is closest to ours is that of Aït-Sahalia and Saglam (2013), who develop a model to study the quote optimization of HFTs that enjoy a latency advantage and trade against many uninformed low-frequency traders (LFTs) and

[8] A number of studies have analyzed the behavior of HFTs classified by NASDAQ (e.g., Brogaard et al. (2014); Brogaard et al. (2016); Brogaard et al. (2018)) and classified by Euronext Paris (e.g., Biais et al. (2016); AMF (2017); Bellia et al. (2017); Colliard and Hoffmann (2017)).

determine why HFTs cancel their orders at a very rapid rate; how HFTs' (endogenous) inventory constraints help shape their order placement and cancellation strategies; how HFTs' provision of liquidity can be expected to change in different market environments, such as high-volatility ones; and how competition among HFTs is expected to affect the provision of liquidity and the welfare of LFTs. We test some of their predictions in this paper. Gerig (2015) differentiates HFTs from other traders in terms of their ability to monitor large numbers of securities contemporaneously and, thus, make better predictions of future order flow. Cespa and Foucault (2011) describe a new mechanism where dealers use the prices of other securities as information that generates spillover effects in terms of both price and liquidity.

Menkveld (2013) analyzes the transactions of a large high-frequently trading firm active on the NYSE Euronext and Chi-X markets right after Chi-X started as an alternative trading venue for European stocks. The author shows that, in 80% of cases, HFTs provided liquidity in both markets during the continuous trading session. In an event study framework, Gomber et al. (2011), Yang et al. (2012), and Menkveld (2013) document the typical behavior of HFTs during the continuous trading session, starting with a zero inventory position at the beginning of the trading day.

These empirical studies reveal specific features of HFTs: HFTs are recognizable by their high frequency of order submissions, high number of trades, high total volume traded, small volume of each trade, and the fact that they carry a low inventory as their primary risk control strategy. HFTs act primarily as market makers, unwilling to make directional bets. HFTs also tend to place many orders, with only some of them being actually executed and many canceled, and they seem to exploit order flow information and generate trading signals on a very short time scale rather than longer-run information about the fundamental value of an asset.

Although HFTs engage in market making, they are not designated market makers. Therefore, their liquidity provision can be expected to decrease as price volatility increases. Since this is precisely when large unexpected orders are likely to hit, markets can become fragile in volatile times, with imbalances arising because of inventories that intermediaries used to but are no longer willing to temporarily hold. These predictions of the models are particularly salient in light of evidence that has emerged regarding flash crashes (Easley et al. (2012); Kirilenko et al. (2017)).

In sum, our paper is related to the HFT literature but differs in several dimensions. First, our study relies on a unique characterization of HFTs that is derived from the specifics of the trading technology rather than relying merely on trading metrics, as described in detail in Section 3. Second, we use the *whole* market sample to identify different trader groups on the TSE, using cluster analysis, which identifies traders with similar trading behaviors as a group from a purely statistical point of view. Most other papers have relied on a flag that differentiates HFTs from non–HFTs or metrics such as the trade-to-cancellation ratio as a threshold. Our reliance on the identification of VSs permits us to get around the problem of limited access to client-specific trading data and still obtain complete data for the whole market.

3 Data

3.1 Sample period

In this study, we use a unique order submission data set provided by the TSE during the eight trading days from May 20 to May 29, 2013. We choose this period to compare HFT-MMs' trading behaviors in normal and volatile periods because theoretical and empirical studies find that the participation of HFTs declines during times of market stress.

[Figure 1 about here]

Figure 1 shows the historical quotes of Nikkei 225 Index Futures on the Osaka Exchange from May 20 through 24, 2013. As shown, the price movement gradually became volatile during the morning session on May 23 and then fell rapidly in the afternoon session. The Nikkei stock average peaked at 15,942 yen in the morning session and dropped rapidly to 14,483 yen by closing time, a 1,143 yen (7.31%) drop from the closing price of the previous day. During the morning session of the next day, May 24, the price recovered to around 15,000 yen and subsided a bit but it then sharply declined again in the afternoon session to a low of 13,982 yen. The fluctuation ranges were larger than 1,000 yen on both May 23 and May 24. We therefore choose the three days from May 20 to May 22, 2013, as the normal period and the two days from May 23 to May 24, 2013, as the volatile period.

3.2 Trading desk formation

Our order submission data set comprises all of TSE's electronic messages, including the date, time, code, limit price, and volume. A key feature of our data set is the inclusion of a virtual server (VS) id and an order identification (OI). The VS is an identification number of the unique connection between a trading desk and the TSE's matching machines by which we can determine which trading desk sent an electronic message. This method enables us to track series of individual trading desk actions. The OI is the identification number of a new order submission. When a trader sets a limit order, for instance, we can track the trader's limit order to execution, cancellation, or modification. Furthermore, we can calculate the time between new limit orders and their cancellation, modification, and execution from the OI. By utilizing this unique information, we attempt to classify trading desks on the TSE into groups and analyze trader behavior at the individual trader level, focusing on HFTs.

We now describe the data preprocessing before the cluster analysis. First, traders handling large numbers of orders use multiple VSs, so we merge VSs used by the same entity and call these merged VSs trading desks (TDs) to distinguish the physical connection equipment of the VS. Figure 2 shows an example of merging VSs.

[Figure 2 about here]

For example, if VS:1 placed a new limit order (OI:1) and VS:2 was used to cancel OI:1 a few seconds later, we consider VS:1 and VS:2 were used by the same trader. If multiple VSs are used to undertake actions with the same OI during the eight trading days from May 20 to May 29, 2013, we merge these VSs into a single TD. We have 4,157 VSs before merging the data and 2,321 TDs after.

Next, because of our focus on HFTs, we extract high-activity TDs, that is, those carrying out more than 100 daily actions per stock. The number of actions per stock per day is defined as

$$action\_per\_stock_{TD,t} = \frac{num\_new_{TD,t} + num\_change_{TD,t} + num\_cancel_{TD,t}}{num\_stock_{TD,t}}, \tag{1}$$

where $num\_new_{TD,t}$ is the number of new orders, $num\_change_{TD,t}$ is the number of order modifications, and $num\_cancel_{TD,t}$ is the number of order cancellations for TD on trading day t. We take the average of the aggregate number of new orders, order modifications, and cancellations per day, $action\_per\_stock_{TD,t}$, over three days, from May 20 to May 22, 2013, when market conditions were normal. A total of 142 TDs satisfy such conditions, which will be used in the cluster analysis described in the next section.

4 Classification with cluster analysis

We classify TDs with cluster analysis.[9] The cluster analysis used in this study is an unsupervised learning, a method for classifying a data set into several groups, or clusters. Cluster analysis classifies a given data set into similar groups, based on the data set. In this study, we employ Ward's method which is one type of the agglomerative hierarchical clustering, since we assume the number of clusters is not obvious and HFTs' strategies have a hierarchical structure in terms of trading style (market making, trend following, statistical arbitrage, etc.)

As described in Section 3.2, our order submission data allow us to distinguish between trading desks but give no information regarding their styles of trading. Therefore, after

[9]Cluster analysis is widely used for various tasks, such as document clustering for document data (Aggarwal and Zhai (2012)), customer segmentation based on purchase history data (Berkhin (2006)) and anomaly detection in sensor data (Chandola et al. (2009)).

grouping together through cluster analysis trading desks that exhibit similar behaviors, we choose groups that match known HFT characteristics.

In our setting, it is important to rely on indicators that represent trading strategy principles and not actions influenced by market conditions. We therefore select the following four proxies for our cluster analysis: the inventory ratio, COR, the number of actions per stock, and the number of stocks per TD.

(i) Inventory ratio

The inventory ratio is a common indicator in empirical studies on HFT behavior. It is the ratio of long or short positions to the total daily trading volume. This indicator represents traders' trading strategy time horizon. Generally, HFTs hold positions for very short periods and try not to carry over positions to the next trading day. The inventory ratio is defined as

$$Inventory\_ratio_{TD,t,i} = \frac{volume\_buy_{TD,t,i} - volume\_sell_{TD,t,i}}{volume\_buy_{TD,t,i} + volume\_sell_{TD,t,i}}, \tag{2}$$

$$Inventory\_ratio_{TD,t} = Median(Inventory\_ratio_{TD,t,i}), \tag{3}$$

where $volume\_buy_{TD,t,i}$ is the total number of shares of buy transactions and $volume\_sell_{TD,t,i}$ is the total number of shares of sell transactions for TD, trading day t, and stock i. We used $Inventory\_ratio_{TD,t}$, that is, the median of each stock i, for the cluster analysis.

(ii) Cancellation-to-order ratio (COR)

The COR is related to traders' limit order management. It also related to the aggressiveness of the limit price in their orders. In the case of HFTs, it indicates traders' awareness of limit order exposure to unfavorable price movement. To reduce intraday adverse selection risk, HFTs try to keep their exposure of limit orders as short as possible. Therefore, it is commonly known that HFTs cancel limit orders within short intervals. The COR is defined as

$$COR_{TD,t} = \frac{num\_cancel_{TD,t}}{num\_new_{TD,t}}. \tag{4}$$

(iii) Number of stocks per TD

The number of stocks traded by a TD reflects the investment strategy adopted by each trader. Although HFTs have an arrangement involving the exclusive usage of multiple VSs because they need the capacity for high-frequency trading, large brokerage houses also subscribe to multiple servers to process the peak-load of orders of millions of customers. Our data cover TDs that use as many as 40 VSs. Two extreme patterns of usage can be distinguished by this variable.

The number of traded stocks per VS is defined as

$$stocks\_per\_VS_{TD,t} = \frac{num\_stocks_{TD,t}}{num\_VS_{TD,t}}, \tag{5}$$

where $num\_stocks_{TD,t}$ is the total number of stocks traded by the TD and $num\_VS_{TD,t}$ is the number of VSs used by the TD on trading day t. Because of TSE's restrictions on the number of electronic messages per second, HFTs assign small numbers of stocks per VS. HFTs can cancel a limit order immediately after its submission; therefore, limited numbers of stocks are assigned to each VS.[10] On the other hand, a large brokerage house handles large numbers of different stocks (CUSIP numbers) so that each of their VSs handles orders for many different stocks.

(iv) Number of actions per stock
The number of actions per stock is our proxy for the latency requirement of each TD. Latency, here, refers to the elapsed time needed to send a message to the TSE. Outside researchers do not have access to such information. We suppose that those who need to conduct high-frequency trading for one stock have access to trading facilities that enable low-latency order submissions. Later, we measure the actual time between a new order's submission and its cancellation, which could be a closer measure of latency; however, such measures are available only for HFTs. The number of messages per stock is a good proxy for all TDs' latency desires.

We classify 141 high-activity TDs into similar trading strategy groups, using Ward's methods based on the four above-mentioned proxies, computed as averages over the period from May 20 to May 22, which is considered the normal period. Hierarchical agglomerative clustering creates an initial state with N clusters for data of sample size N. Starting from an initial state, we calculate the distance $d(C_1, C_2)$ between the clusters C_1 and C_2 including x_1 and x_2, respectively, and merge the two closest clusters one by one. By repeating this merging procedure until all clusters are merged into one, we can acquire a hierarchical structure (Hastie et al. (2001)). The distance function $d(C_i, C_j)$ between clusters in Ward's method is defined as

$$\begin{aligned} d(C_i, C_j) &= E(C_i \cap C_j) - E(C_i) - E(C_j), && (6) \\ E(C_i) &= \sum_{x \in C_i} (x - \overline{x_i})^2. && (7) \end{aligned}$$

Two clusters are merged to minimize the difference between the cluster dispersion after merging and the sum of the variance of each cluster before merging one by one. We

[10] The TSE provides three levels of service, with a maximum of 60, 40, and 20 messages per second, respectively. We refer to all messages, including cancellations, in our definition of quotes. According to a prominent HFT, for a trader to remain truly anonymous, at least 20 VSs are necessary to implement a strategy of trading 1,500 stocks all at once. If the HFT also needs to cancel several orders immediately after submitting new ones, an additional 20 VSs can be required, for a total of 40 VSs necessary to support intensive HFT activity across multiple stocks. Using multiple VSs, each trader optimizes the performance of the trading operations for subsets of stocks. Some traders operate within a specific group of stocks every day, in which case they may establish the allocation of stocks to each VS. Other traders may change part of their allocation on a day-by-day basis.

preprocess the input data by taking the natural logarithms of the number of actions per stock and the number of stocks per VS. Furthermore, to match the scale between the four indicators, each indicator is standardized by its average value and standard deviation. Figure 3 shows the results of the classification of TDs with Ward's method as a dendrogram. The dendrogram is the tree diagram obtained by hierarchical clustering and the vertical axis represents the distance between individual TDs or clusters.

[Figure 3 about here]

Table 1 summarizes the descriptive statistics of 10 groups classified by Ward's method. This table describes the averages over the period from May 20 to May 22 for each group. The numbers 1 to 10 on the horizontal axis in Figure 3 correspond to the group numbers in Table 1. To grasp the characteristics of each group, in addition to the four proxies used in the cluster analysis, we use seven additional indicators: the number of new orders, the number of total actions, the number of filled order and the rate at which they were filled, COR, the order quantity, the number of market orders, and the number of VSs used by a TD.

[Table 1 about here]

Groups 9 and 10 satisfy the conditions of a high COR and a low inventory ratio, a characteristic of HFT-MMs; they have the highest frequency of $action\_per\_stock_{TD,t}$ among the 10 groups; and their $stocks\_per\_VS_{TD,t}$ values are very low, at 13 and two, respectively. These behavioral characteristics match those of HFT-MMs. Another distinguishing behavioral feature is the complete avoidance of market order. Although market orders can be executed immediately, they can be executed at unexpected prices. Given the increasing numbers of low-latency traders on the TSE, the best quotes at the time of order submission may no longer be available when the order reaches the exchange.[11] In an order-driven market such as the TSE, limit orders are a vehicle to supply liquidity to the market. To earn on the bid–ask spread, HFT-MMs must trade through limit orders.

Among groups 1 through 8, there are some satisfy either the high COR criterion or the low inventory ratio criterion. For example, group 8 has the second lowest inventory ratio, 8.0%, but COR is 46%, ranking sixth among the 10 groups, which is not as high as COR of groups 9 and 10, which are 91% and 78%, respectively. Groups 2 and 7, respectively, have high COR, 85% and 80%, but relatively high inventory ratios of 63.0% and 30.2%.

[11] Marketable limit orders are 14.8% and 18.2% of the total number of limit orders submitted by HFT-MMs in the normal and volatile periods, respectively. These ratios are much lower compared to other groups

The 10 groups identified by cluster analysis can be categorized into three distinct styles, mainly based on frequency: groups 9 and 10 are identified as HFT-MMs; groups 2, 5, 7, and 8 are HFT-OPs, and groups 1, 3, 4, and 6 are middle-frequency traders (MFTs). Trading desks excluded from the cluster analysis are LFTs. Table 2 shows the descriptive statistics for the four groups.

[Table 2 about here]

HFT-OPs have equivalent numbers of actions per stock and higher numbers of new order submissions as HFT-MMs but differ in terms of their inventory ratio (32.7%) and COR (59%).[12] In terms of order size, they place orders that are almost six times as large (1,851 shares) as those of HFT-MMs (347 shares), about 21% of their new orders are market orders, and they have much higher fill rates (50%) than HFT-MMs do (19%).[13] Clearly the immediate execution of orders has a higher priority among HFT-OPs. This meets the characteristics of HFT-OPs who engage a variety of arbitrage trading strategies such as cross-market arbitrage, statistical arbitrage, and news arbitrage (Hagströmer and Nordén (2013); SEC (2014); Kirilenko et al. (2017)). MFTs exhibit the lowest activity (new orders and actions) among all the groups except for LFTs. The fill rate is 76%, the inventory ratio is 85.0%, and only 13.9% of new orders are market orders, less than for HFT-OPs. The higher fill rate and much lower market order usage are due to MFTs trading with cost-cautious manner. LFTs submit orders less frequently because of their high fill rate (71%) and inventory ratio (94.9%). These features are consistent with a buy-and-hold strategy. The four groups have the highest average order quantities, which suggests the presence of institutional investors.

The number of stocks per VS used by each group varies substantially. HFT-MMs assign the smallest number of stocks among the four groups which is seven stocks in average. HFT-OPs assign 17 stocks per VS and they use 14 VSs which is the largest operation in terms of multiple VSs usage.[14] Whereas LFTs assign 347 stocks per VS and handle orders of the broad-range of stocks submitted by many different types of investors.

In sum, our cluster analysis successfully separates out four groups with many distinguishing trading style characteristics. In addition to frequency of order submissions

[12] Our evidence that HFT-MMs and HFT-OPs show the same order submission level is different from Hagströmer et al. (2014) which analyze NASDAQ-OMX market in August 2011 and February 2012. The level of order submissions from market making HFTs represented 86% of HFT orders submissions, while opportunistic HFTs represented 14% of them.

[13] Marketable limit orders are 28.9% and 32.3% of the total number of limit orders submitted by HFT-OPs in the normal and volatile periods, respectively. Their ratios are almost double of those for HFT-MMs.

[14] The combination of 14 VSs and 17 stocks per VS allows them to trade 238 stocks simultaneously. This number is close to 225 stocks in Nikkei Stock Average which is a popular index arbitrage strategy.

(total as well as per stock), the fill rate, inventory and COR are substantially different.[15] The difference of the number of stocks per server indicates that the latency requirement of the trading infrastructure is very different among four trader groups.

5 Hypotheses

In the following empirical section, we investigate the behavior of HFT-MMs under different market conditions. HFT-MMs have become the main liquidity provider in the TSE but are not designated market makers. Therefore, their provision of liquidity can be expected to decrease as price volatility increases, as suggested by the literature (Easley et al. (2012), Aït-Sahalia and Saglam (2013), and Kirilenko et al. (2017)). This is an important issue which requires empirical examination. Our data set allows us to examine the relation between HFT-MM's actions and rapidly changing market conditions in our sample periods. As we inspected in the previous section, HFT-MM group satisfies many important features of HFT-MM behaviors. Guided by the following recent theoretical models, we attempt to find triggers of their decision-making with respect to information available from the market.

The literature suggests that there are two possibilities that HFTs utilize their speed advantages. One possibility is that HFT-MMs behaviors are driven by price changes of individual and an entire stock market. Biais et al. (2015) consider a setting in which a HFT enjoys a speed advantage when gathering information. As Foucault et al. (2016) suggest HFT can react to news faster than others. News reflect to prices of financial assets. The other possibility is that they monitor order flow in the market. Aït-Sahalia and Saglam (2013) suggest that HFT-MM engages liquidity provision with superior ability of order flow monitoring and inventory risk management. We test the following four hypotheses.

Hypothesis 1 *High volatility leads to HFT-MMs' reduction in liquidity supply.*

Kirilenko et al. (2017) have analyzed the 2010 Flash Crash in the United States and find sell orders from HFT-MMs were the result of increased inventory. Aït-Sahalia and Saglam (2013) suggest that higher volatility reduces the liquidity provision of HFT-MMs.

Hypothesis 2 *HFT-MMs monitor and react to the order flow of stocks in their universe.*

Aït-Sahalia and Saglam (2013) model HFT-MMs' quote posting/cancellation behavior. However, HFT-MMs act as a signal, with orders from other investors, especially LFTs. In this case, order flow information is more important than price information.

[15] Hagströmer et al. (2014) assess order types from different trader groups and they find, with respect to order cancellation rates, there are very small economic differences between HFTs and non-HFTs in NASDAQ-OMX market.

Hypothesis 3 *HFT-MMs monitor the price movements of stocks to be traded and submit orders accordingly.*

Hypothesis 4 *HFT-MMs monitor inventory fluctuations due to trading.*

The model of Aït-Sahalia and Saglam (2013) shows that HFT-MMs manage the risk exposure of their inventory levels through high-frequency trading. This hypothesis can complement Hypothesis 1.

The testing period covers the five days from May 20 to 24, 2013. We choose this period because, in the first three days, stock prices are stable but, in the latter two days (on May 23 and May 24), stock prices become very volatile as described in Section 3.1. These periods are suitable to verify the hypotheses.

6 Empirical analysis

6.1 Comparison of the normal and volatile periods

First, we provide the results of the test for Hypothesis 1. Table 3 summarizes the descriptive statistics of the four groups during the normal and volatile periods.

[Table 3 about here]

In both Panels A and B of Table 3, the number of new order submissions is larger than in Table 2. The share of new orders from HFT-MMs is higher in the volatile period than in the normal period. The order execution rate is slightly higher but the cancellation ratio remains at the same level as in the normal period. We do not find any measures in Table 3 to support Hypothesis 1. The high volatility of May 23 and May 24 was probably caused by the halt in trading of the government bonds futures market due to the introduction of the Bank of Japan's aggressive monetary policy. The source and nature of the price swings was not stock market specific and differed from the situation in the Flash Crash in 2010. HFT-MMs remain not only actively engaged in market making, but also behave the same as during the normal period in terms of the cancellation ratio, the avoidance of market orders, and so forth.

[Figure 4 about here]

It is interesting to note that the simultaneous increase in order submissions and cancellation suggest that HFT-MMs' limit orders sit in the book for shorter times than before. As Figure 4 shows, the limit order's elapsed time between new order entry and cancellation is 32.4 seconds in the normal period and 12.1 seconds in the volatile period. In response to the fast-moving market in our sample, HFT-MMs increased their frequency of order management (new submission/cancellation cycle) to adapt to volatile market conditions instead of withdrawing from the market.

6.2 Variable definitions for SVAR estimation

To further investigate HFT-MMs' decision making, we estimate a SVAR model with aggregated data every ten-seconds. Current super-fast trading infrastructures make fast and high-frequent trading possible, with responses to event at the millisecond level. Our order submission data therefore needs to be obtained at proper time intervals. The best way to measure latency is the elapsed time between order entry and cancellation. Orders cancelled within ten seconds comprise 74% of all orders under the normal period. Therefore, we pursue the following analysis of HFT-MMs' trading decisions using intraday data in ten-seconds interval. The universe of stocks comprises 479 firms traded by HFT-MMs in the period from May 20 to May 22.

We first define the trading activity variables for HFT-MMs. For each stock, we define the following variables:

$buy^n_{s,MM}$: The number of buy limit orders submitted by HFT-MMs at time s.

$buy^c_{s,MM}$: The number of buy limit order cancellations submitted by HFT-MMs at time s.

The liquidity provision though HFT-MMs' buy limit orders can be triggered by the sell orders of other investor groups. To examine this relation, we define the following:

$sell^n_{s,other}$: The number of sell orders submitted by all investors other than HFT-MMs at time s.

We next define the order imbalance and inventory measures. Our imbalance measure is defined as the difference between the numbers of buy and sell orders from all traders:

$$imb_s = buy^n_s - sell^n_s. \tag{8}$$

We also define an inventory measure, as follows:

$$inv_{s,MM} = volume\_buy_{s,MM} - volume\_sell_{s,MM}, \tag{9}$$

where $volume\_buy_{s,MM}$ ($volume\_sell_{s,MM}$) is the number of purchases (sales) carried out by HFT-MMs at time s. Therefore, $inv_{s,MM}$ indicates the inventory change of HFT-MMs per second. We further define future price to examine the effect of price changes in the entire stock market. We define the following:

$fprice_s$: The rate of change of the futures mid-quote (average of the bid and ask prices for Nikkei stock index futures) at time s.

We use the nearest contracts of Nikkei stock average futures for our futures prices traded on the Osaka Exchange. We aggregate all variables into ten-second intervals when we estimate a SVAR.

6.3 Empirical approach

In this section, we look to our trading activity and price variables and construct our SVAR model.

[Table 4 about here]

We first check stationarity. As summarized in Table 4, although the future price, $fprice$, cannot reject the null hypothesis of a unit root, the first difference, $\Delta fprice$, can be rejected according to the Dickey–Fuller test. We also look at the correlations between these six variables, which are not tabulated here. Since two trade activity variables buy^n_{MM} and buy^c_{MM} and HFT-MMs' inventory inv_{MM} are highly correlated (71.3% and 58.3%), we use the first differences Δbuy^n_{MM} and Δbuy^c_{MM} instead of buy^n_{MM} and buy^c_{MM}, respectively, for our empirical model.

[Table 5 about here]

As shown in Table 5, the correlation between Δbuy^n_{MM} (Δbuy^c_{MM}) and $sell^n_{other}$ is 0.294 (0.373). HFT-MMs increase their limit buy orders (cancellations) when the other investor groups increase their sell orders. We also find that innovations in price and order imbalance are positively correlated. Two sets of variables $\{\Delta buy^n_{MM} \;\; \Delta buy^c_{MM}\}$ and $\{sell^n_{other} \;\; inv_{MM}\}$ are highly correlated (70.9% and 76.0%). We set four different SVAR specifications so that these variables are not included in the same estimating model in our empirical analyses.

We now construct our empirical model using the SVAR specification. Based on the unit root tests and correlations above, we set the following vectors of state variables

$$\begin{aligned}
Y^{n1}_s &= (\Delta buy^n_{MM} \;\; sell^n_{other} \;\; imb \;\; \Delta fprice)' \\
Y^{c1}_s &= (\Delta buy^c_{MM} \;\; sell^n_{other} \;\; imb \;\; \Delta fprice)' \\
Y^{n2}_s &= (\Delta buy^n_{MM} \;\; imb \;\; inv_{MM} \;\; \Delta fprice)' \\
Y^{c2}_s &= (\Delta buy^c_{MM} \;\; imb \;\; inv_{MM} \;\; \Delta fprice)', \qquad (10)
\end{aligned}$$

and the SVAR takes the following form:

$$Y_s^{model} = \beta_0 + \beta_1 Y_{s-1} + \beta_2 Y_{s-2} + \ldots + \beta_l Y_{s-l} + \eta_s, \tag{11}$$

where Y^{model} is a vector of state variables, where $model$ is either $n1$, $c1$, $n2$, or $c2$. η_s is the reduced-form residual with $E[\eta_s \eta_s'] = \Sigma$. We denote the mapping from structural shocks ν_s to the residual η_s by the matrix $\mathbf{A}$, with $\eta_s = \mathbf{A}\nu_s$. The identification of the structure requires a set of constraints on $\mathbf{A}$. Different orderings of the endogenous variables and different constraints on $\mathbf{A}$ yield different IRFs. Chordia et al. (2005) chooses an ordering based on microstructure theory, where information or endowment shocks affect prices and liquidity through trading. Accordingly, we first order the four trading variables Δbuy^n_{MM}, Δbuy^c_{MM}, $sell^n_{other}$, and imb; then inventory inv_{MM}; and, finally, set futures price, $\Delta fprice$. We further impose the following restrictions on the matrix $\mathbf{A}$,

$$\mathbf{A} = \begin{pmatrix} 1 & * & * & 0 \\ * & 1 & * & 0 \\ * & * & 1 & 0 \\ * & * & * & 1 \end{pmatrix}. \tag{12}$$

The restriction implies that we identify an idiosyncratic shock to future price assuming that on impact, the effects of the shock to the other state variables such as the number of new limit buy orders by HFT-MMs, those of cancellations of buy orders by HFT-MMs, those of new limit sell orders by groups other than HFT-MMs, marketwide imbalance, and inventory held by HFT-MMs are close to zero. In the empirical estimation, we choose l to be 6, the number of lags in Eq. (11) by the Schwarz Bayesian information criterion.[16]

[Table 6 about here]

Table 6 presents pairwise Granger causality tests between the endogenous variables of the SVAR. The cell associated with the variable in row i and column j shows the chi-squared statistics and the significance of the test whose null hypothesis is that variable i does not Granger-cause variable j. In Table 6, many endogenous variables Granger-cause other endogenous variables. There is two-way causation between each pair of the variables Δbuy^n_{MM}, Δbuy^c_{MM}, $sell^n_{other}$, imb, inv_{MM}, and $\Delta fprice$ in most cases. On the other hand, for marketwide imbalance and $sell^n_{other}$ (or inv_{MM}), the null hypotheses cannot be rejected, indicating $sell^n_{other}$ (inv_{MM}) does not Granger-cause imb.

To determine whether each variable is endogenous or exogenous, we further perform block exogeneity tests to determine whether the lag of any variables Granger-causes any

[16]Since the slopes of the information criterion (as a function of lags) on the basis of the other criteria are flat for greater lags, we use the Schwarz information criterion, which indicates shorter lags.

other variable in the system. The results, not tabulated, indicate no endogenous variable to be excluded from the SVAR system. We therefore decide to use the four vectors of state variables in Eq. (10) in our empirical model.

[Table 7 about here]

Table 7 reports summary statistics for the endogenous variables of VAR system. We present the statistics separately for the normal period (May 20 to May 22) and the volatile period (during the afternoon trading sessions of May 23 and May 24). During the period of large fluctuations, changes in the numbers of HFT-MM orders and cancellations have a larger standard deviation (214.05 and 233.58, respectively) than during a normal period (120.32 and 129.27, respectively). The numbers of sell orders by all groups except for HFT-MMs also has a larger mean and standard deviation during the period of wide fluctuations. The mean order imbalance is positive (buy orders predominate) during the normal period but becomes negative (sell orders predominate) during the volatile period, when the variation in order imbalance also increases.

6.4 IRFs

We now proceed with additional tests for Hypotheses 2 to 4, described in Section 5, using the data for the normal and volatile periods. We compute the cumulative IRFs for the endogenous variables of the SVAR system. A cumulative IRF shows the cumulative impact of the positive shock of one unit of an endogenous variable. We orthogonalize the impulse using Cholesky decomposition of the residual covariance matrix. The orderings of the endogenous variables and constraints on rotate matrix are described in Section 6.3. We follow the order of Δbuy^{n}_{MM}, Δbuy^{c}_{MM}, $sell^{n}_{other}$, imb, inv_{MM}, and $\Delta fprice$.

[Figure 5 about here]

Figure 5 shows the cumulative IRFs of the numbers of new limit buy orders and buy order cancellations submitted by HFT-MMs to one unit of the sell orders shock from other trader groups for a period of 18 intervals (ten seconds $\times$ 18 intervals, i.e. three minutes). Two standard error confidence bands from 500 runs of bootstrapping are provided to gauge the statistical significance of the responses. In response to new sell orders from others during the normal period, specified above, the number of HFT-MMs' new buy orders increases by 0.3 unit of the variable on the first ten seconds and then rapidly decay from day one to day three and more gradually after that. The number of HFT-MMs' buy order cancellations increases by 0.35 on the first interval and then diminishes the impact on second interval and remains around 0.1 thereafter. An innovation in sell orders shock from other trader groups forecasts an increase in new submission as well as cancelation of buy limit orders by HFT-MM. The results indicate higher participation of HFT-MM. In the case of a volatile period, such as on May 23 and May 24, an innovation in sell orders shock from other trader groups forecasts similar pattern and magnitude of responses from HFT-MM, except that the response of the cancellation gradually decays after 80 seconds.

[Figure 6 about here]

Figure 6 shows the IRFs for the numbers of new limit buy orders and buy order cancellations by HFT-MMs to an order imbalance shock. HFT-MMs' new buy limit orders show a significant positive response to an order imbalance shock on the first interval and disappear quickly while their cancellation show significant positive response after ten seconds. During the volatile period, the response of new buy orders becomes insignificant, but that of cancellations is initially significantly negative and thereafter insignificant. An innovation in order imbalance shock forecasts an immediate increases in new submission and delayed increases in cancelation of buy limit orders by HFT-MM, but the magnitude of the responses are smaller than those for other trader groups' order submission shock.

[Figure 7 about here]

Figure 7 show the IRFs of the numbers of new limit buy orders and buy order cancellations by HFT-MMs to a futures price shock. An innovation in futures market returns does not forecasts an increases/decrease in new submission, but it forecasts negative response in cancelation of buy limit orders by HFT-MMs. HFT-MMs reduce their buy order cancellations when futures market prices rise.

[Figure 8 about here]

Figure 8 shows the IRFs of the numbers of new limit buy orders and buy order cancellations by HFT-MMs to an inventory shock. The inventory response to new buy orders from HFT-MM is strikingly high compared to other three shocks examined above. In response to inventory shock during the normal period, the number of HFT-MMs' new buy orders increases by closer to five units on the first ten seconds and then rapidly decay from interval of ten to 20 seconds and more gradually decline thereafter. The number of HFT-MMs' buy order cancellations increases by 4 on the first ten seconds and decay thereafter. The response during the volatile period is similar, but the magnitude is milder. These results might appear counter-intuitive to the notion that HFT-MMs are sensitive to the level of an inventory. Because our definition of an inventory variable is the difference between the numbers of buy and sell in a ten-second interval. This corresponds to a change in inventory instead of the level. When a change in inventory is positive, it means buy limit order submitted by HFT-MM is executed so that HFT-MM resubmit new limit buy order to engage market making. The results suggest that HFT-MMs monitor their inventory change and keep engaging liquidity provision through limit orders.

We conclude from the evidence of the IRFs that HFT-MMs react to their inventory and order flow from other market participants rather than to price movements when deciding on their trading actions.

7 Conclusion

In this research, we use virtual server identifications provided by the TSE to identify HFTs with similar behavioral characteristics through cluster analysis. We select four proxy variables: the inventory ratio, the order cancellation ratio, the order frequency, and the number of stocks per server. The first two variables, such as high cancellation and low inventory ratios, are the most commonly acknowledged features of HFT-MMs. The latter two variables highlight the importance of the ability to trade at high frequencies and quickly. The combination of these four variables successfully finds a group of traders matching the key HFT-MM characteristics. These are differentiated from other fast traders by their exclusive usage of limit orders. We conduct a comparative analysis of trading behaviors between the calm and volatile periods. In the volatile period, HFT-MMs place more orders than in the normal period and keep equivalent order execution rate. This suggests the simultaneous increase in order submissions and cancellation. It means that HFT-MMs' limit orders sit in the book for far-shorter times than before. The limit order's elapsed time between new order entry and cancellation is 32.4 seconds in the normal period and 12.1 seconds in the volatile period. Orders cancelled within one second increases to 22.1

Upon investigation of HFT-MM trading decision makings, we estimate a SVAR model with 6 lags with intraday aggregated data every ten-seconds interval and generate IRFs to identify whether order flow or price/quote information triggers HFT-MMs' trading actions. Our results suggest that HFT-MMs respond to the order flow of other participants and their own inventory changes than to the price/quote movement of stock index futures. During times of high volatility, HFT-MMs respond with greater magnitude and more quickly to such information. HFT-MMs remain liquidity providers in the market during both periods as oppose to prior literature.

References

Aggarwal, C. C. and C. Zhai (2012). A survey of text clustering algorithms. In C. C. Aggarwal and C. Zhai (Eds.), *Mining Text Data*, pp. 77–128. Springer.

Aït-Sahalia, Y. and M. Saglam (2013, October). High frequency traders: Taking advantage of speed. Working paper 19531, National Bureau of Economic Research.

AMF (2017). Study of the behaviour of high-frequency traders on euronext paris. Autorite des Marches Financiers (AMF).

Bellia, M., L. Pelizzon, M. G. Subrahmanyam, J. Uno, and D. Yuferova (2017). Coming early to the party. Working paper 182, SAFE. Available at SSRN: https://ssrn.com/abstract=3038699.

Berkhin, P. (2006). A survey of clustering data mining techniques. In J. Kogan, C. Nicholas, and M. Teboulle (Eds.), *Grouping Multidimensional Data*, pp. 25–71.

Biais, B., F. Declerck, and S. Moinas (2016). Who supplies liquidity, how and when? BIS working paper 563. Available at SSRN: https://ssrn.com/abstract=2789736.

Biais, B. and T. Foucault (2014). Hft and market quality. *Bankers, Markets and Investors 128*, 5–19.

Biais, B., T. Foucault, and S. Moinas (2015). Equilibrium fast trading. *Journal of Financial Economics 116*(2), 292–313.

Biais, B. and P. Woolley (2011). High frequency trading. Working paper, Toulouse School of Economics and London School of Economics.

Brogaard, J., B. Hagstromer, L. Nordén, and R. Riordan (2015). Trading fast and slow: Colocation and liquidity. *The Review of Financial Studies 28*(12), 3407–3443.

Brogaard, J., T. Hendershott, and R. Riordan (2014). High-frequency trading and price discovery. *Review of Financial Studies 27*(8), 2267–2306.

Brogaard, J., T. Hendershott, and R. Riordan (2016). High-frequency trading and the 2008 short sale ban. Working paper, University of Washington.

Brogaard, J., T. Hendershott, and R. Riordan (2018). Price discovery without trading: Evidence from limit orders. Available at SSRN: https://ssrn.com/abstract=2655927.

Cespa, G. and T. Foucault (2011). Learning from prices, liquidity spillovers, and market segmentation. Discussion paper DP8350, CEPR.

Chandola, V., A. Banerjee, and V. Kumar (2009, July). Anomaly detection: A survey. *ACM Comput. Surv. 41*(3), 15:1–15:58.

Chordia, T., A. Goyal, B. Lehmann, and G. Saar (2013). High-frequency trading. Research paper series 20-2013, Johnson School. Available at SSRN: https://ssrn.com/abstract=2278347.

Chordia, T., A. Sarkar, and A. Subrahmanyam (2005). An empirical analysis of stock and bond market liquidity. *The Review of Financial Studies 18*(1), 85 – 129.

Colliard, J.-E. and P. Hoffmann (2017). Financial transaction taxes, market composition, and liquidity. *The Journal of Finance 72*(6), 2685–2716.

Easley, D., M. López de Prado, and M. O'Hara (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies 25*(5), 1457–1493.

Foucault, T., J. Hombert, and I. Roşu (2016). News trading and speed. *Journal of Finance 71*(1), 335–382.

Gerig, A. (2015). High-frequency trading synchronizes prices in financial markets. Available at SSRN: https://ssrn.com/abstract=2173247.

Gomber, P., B. Arndt, M. Lutat, and T. E. Uhle (2011). High-frequency trading. Available at SSRN: https://ssrn.com/abstract=1858626.

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

Hagströmer, B., L. Nordén, and D. Zhang (2014). How aggressive are high-frequency traders? *Financial Review 49*(2), 395–419.

Hastie, T., R. Tibshirani, and J. Friedman (2001). *The Elements of Statistical Learning*. Springer Series in Statistics. New York, NY, USA: Springer New York Inc.

Hosaka, G. (2014). Analysis of high-frequency trading at tokyo stock exchange.

Jones, C. M. (2013). What do we know about high-frequency trading? Research paper 13-11, Columbia Business School. Available at SSRN: https://ssrn.com/abstract=2236201.

Jovanovic, B. and A. J. Menkveld (2011). Middlemen in limit-order markets. Working paper, WFA.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun (2017). The flash crash: High-frequency trading in an electronic market. *Journal of Finance 72*(3), 967–998.

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

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

SEC (2010). Concept release on equity market structure. U.S. Securities and Exchange Commission (SEC).

SEC (2014). Equity market structure literature review part ii: High frequency trading. U.S. Securities and Exchange Commission (SEC).

Yang, S. Y., M. E. Paddrik, R. Hayes, A. Todd, A. Kirilenko, P. A. Beling, and W. T. Scherer (2012). Behavior based learning in identifying high frequency trading strategies. *2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr)*, 1–8.

Figure 1: Nikkei 225 index futures on the Osaka Exchange from May 20 to May 24, 2013

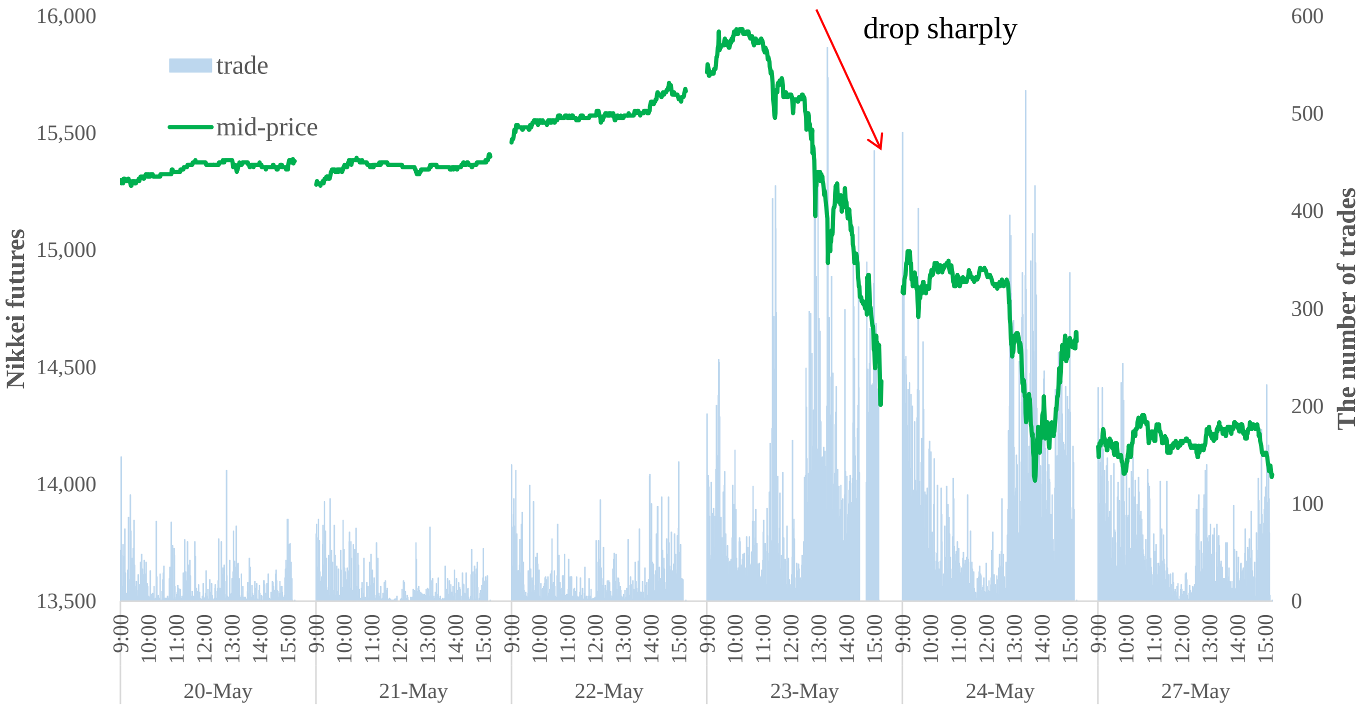

Note: The minute-by-minute evolution of the mid-quote of Nikkei 225 futures and the number of trades from May 20 through May 24, 2013, excluding times before opening in the morning and afternoon. The data set is from Osaka Exchange.

Figure 2: Example of merging VSs into one trading desk

Action	**OI**	**VS**
Set a limit order	10	1
Change a limit price	10	1
Set a limit order	11	7
Execute a limit order	10	5
Cancel a limit order	11	9

VS:1 and VS:5
are used by the same trader.

VS:7 and VS:9
are used by the same trader.

Note: This figure shows an example of merging VSs into a single TD.

Figure 3: Dendrogram

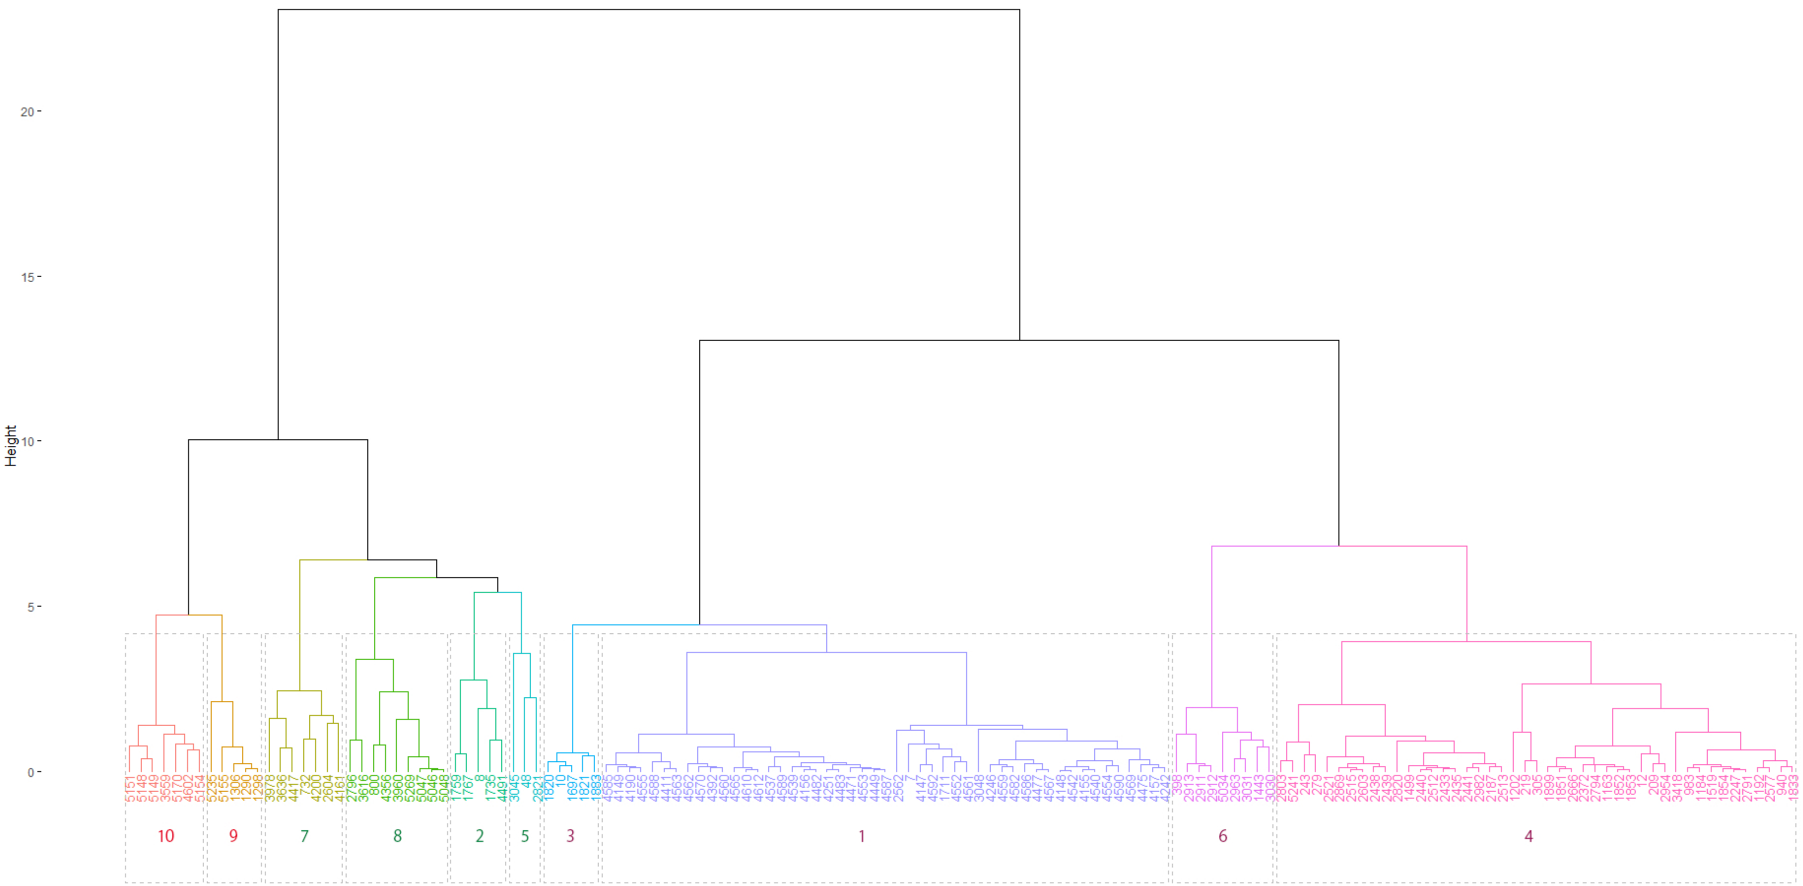

Note: This figure shows the dendrogram for the classification of TDs using Ward's method. The dendrogram is a tree diagram obtained by the hierarchical clustering. Trading desks are split into 10 groups, depicted by boxes in dashed lines. The numbers 1 to 10 on the horizontal axis represent the trader group being classified. The data set cover the period from May 20 through May 22, 2013.

Figure 4: Elapsed times until execution and cancellation, May 22 to May 23, 2013

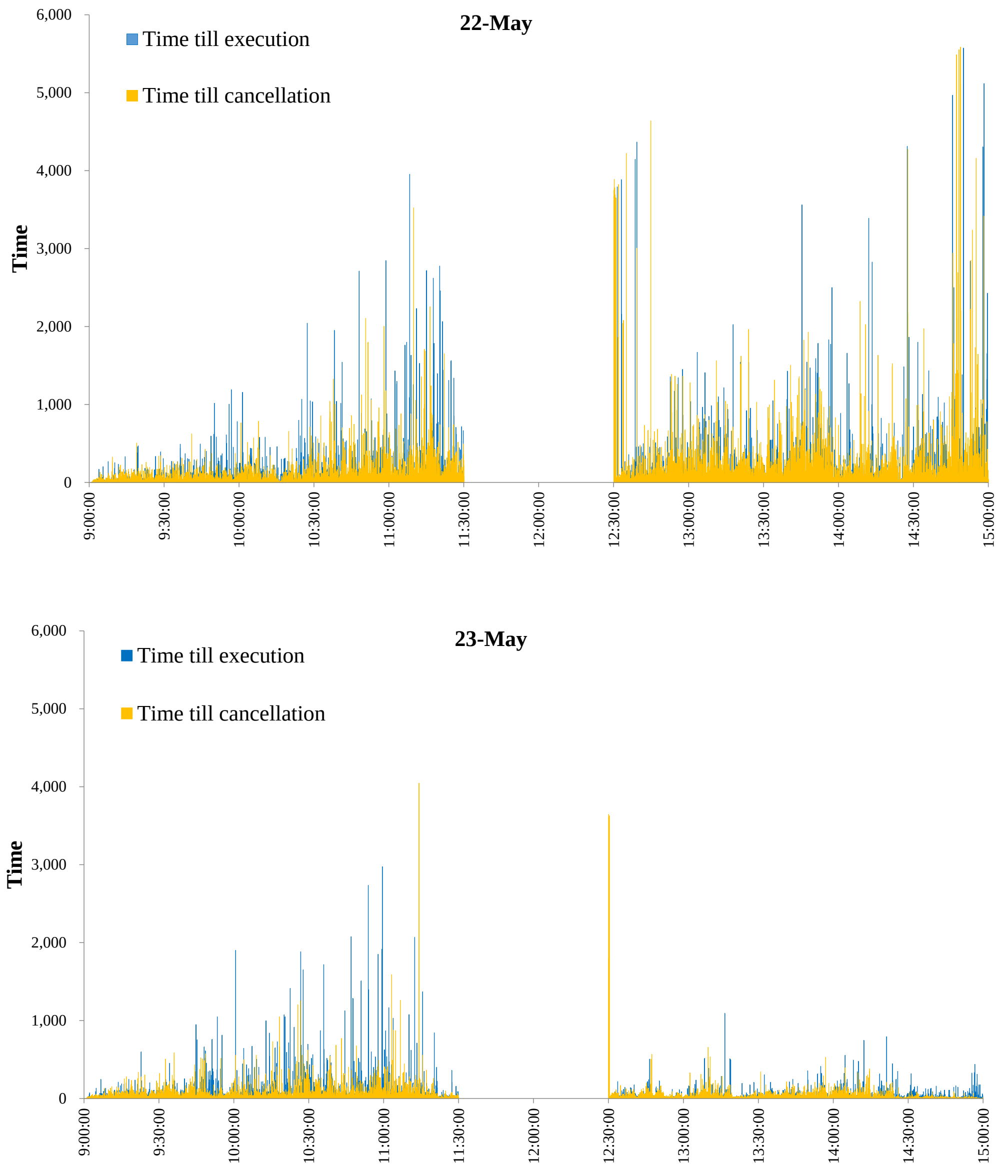

Note: This figure shows the elapsed times of limit orders (in units of seconds) between new order entry and execution (blue line) and between new order entry and cancellation (grey line). The horizontal axis shows the time and the data set cover May 22 and May 23, 2013.

Figure 5: Response to a sell order shock from others

(a) Response of Δbuy^{n}_{MM} to $sell^{n}_{other}$

(b) Response of Δbuy^{c}_{MM} to $sell^{n}_{other}$

20–22

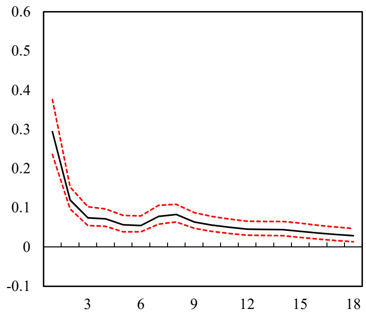

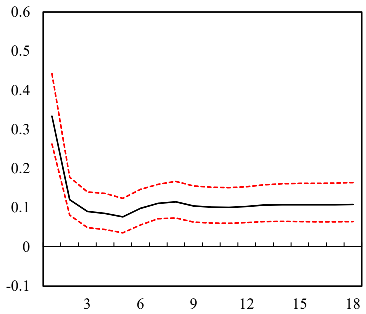

23–24

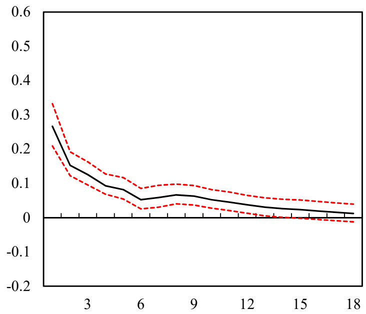

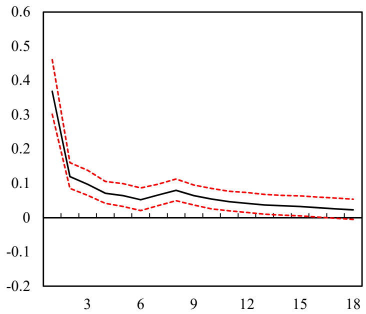

Note: This figure shows the cumulative IRFs of (a) the number of new limit buy orders by HFT-MMs and (b) the number of buy order cancellations by HFT-MMs to a sell order shock from other traders. The upper panels are for May 20 through May 22, 2013, and the lower panels are for during the afternoon trading sessions of May 23 and May 24, 2013. The mean impulse response is shown in black. The dashed lines mark a pointwise 95% credible interval around the median based on 500 runs of bootstrapping.

Figure 6: Response to an order imbalance shock

(a) Response of Δbuy^{n}_{MM} to imb (b) Response of Δbuy^{c}_{MM} to imb

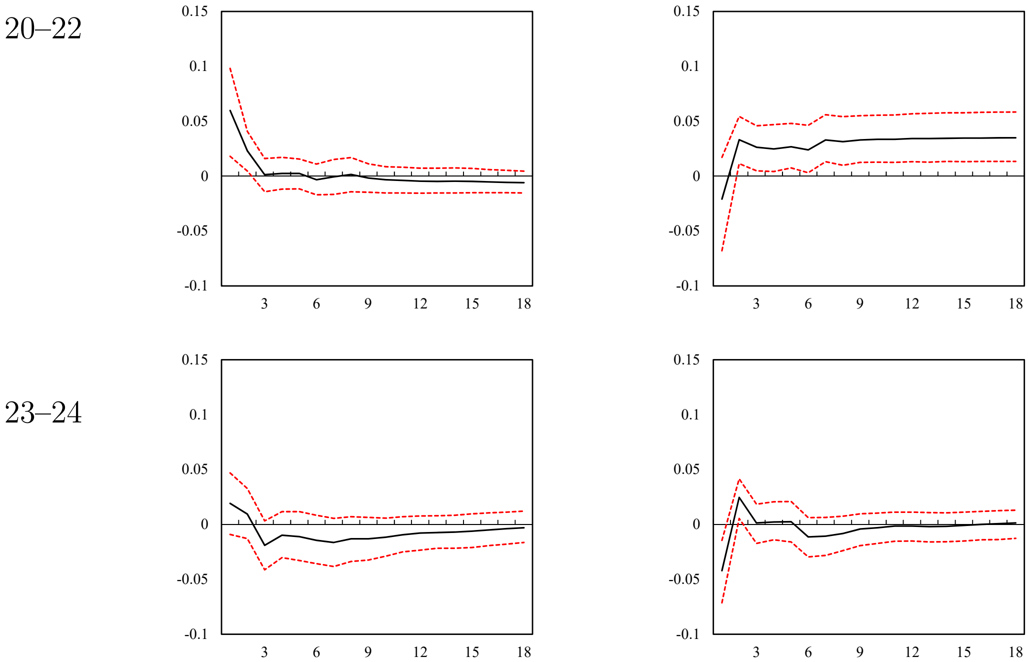

Note: This figure shows the cumulative IRFs of (a) the number of new limit buy orders by HFT-MMs and (b) the number of buy order cancellations by HFT-MMs to an order imbalance shock. The upper panels are for May 20 to May 22, 2013, and the lower panels are for during the afternoon trading sessions of May 23 and May 24, 2013. The mean impulse response is shown in black. The dashed lines mark a pointwise 95% credible interval around the median based on 500 runs of bootstrapping.

Figure 7: Response to a futures price shock

(a) Response of Δbuy^{n}_{MM} to $\Delta fprice$

(b) Response of Δbuy^{c}_{MM} to $\Delta fprice$

20–22

23–24

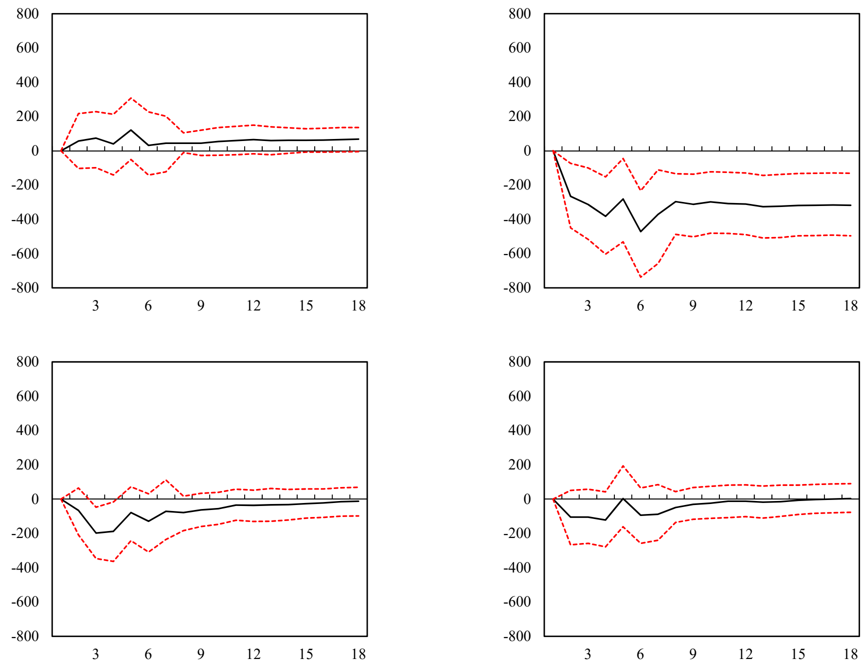

Note: This figure shows the cumulative IRFs of (a) the number of new limit buy orders by HFT-MMs and (b) the number of buy order cancellations by HFT-MMs to a futures price shock. The upper panels are for May 20 to May 22, 2013, and the lower panels are for during the afternoon trading sessions of May 23 and May 24, 2013. The mean impulse response is shown in black. The dashed lines mark a pointwise 95% credible interval around the median based on 500 runs of bootstrapping.

Figure 8: Response to an inventory shock

(a) Response of Δbuy^{n}_{MM} to inv_{MM}

(b) Response of Δbuy^{c}_{MM} to inv_{MM}

20–22

23–24

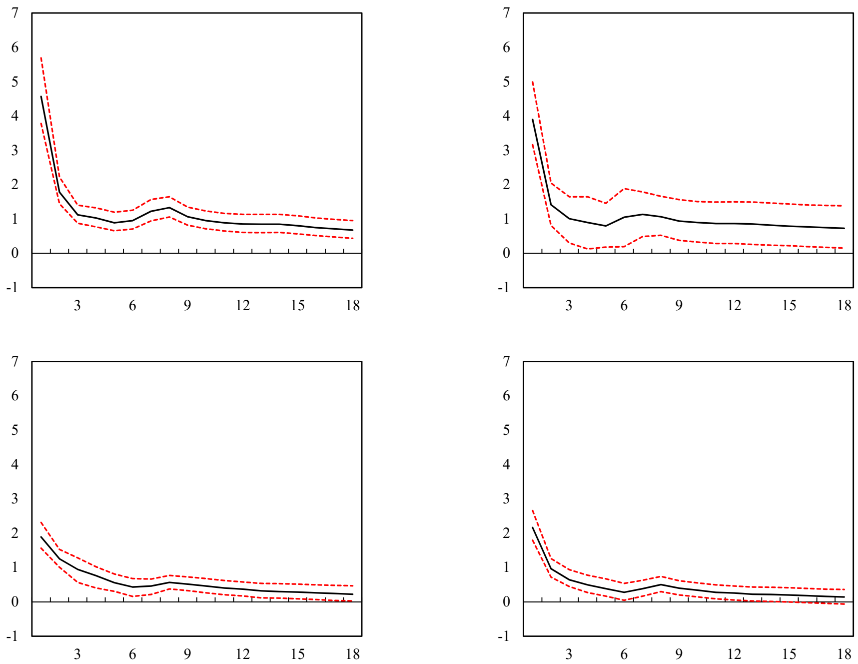

Note: This figure shows the cumulative IRFs of (a) the number of new limit buy orders by HFT-MMs and (b) the number of buy order cancellations by HFT-MMs to an inventory shock. The upper panels are for May 20 to May 22, 2013, and the lower panels are for during the afternoon trading sessions of May 23 and May 24, 2013. The mean impulse response is shown in black. The dashed lines mark a pointwise 95% credible interval around the median based on 500 runs of bootstrapping.

Table 1. Descriptive statistics of the trading behaviors for groups classified by Ward's method

Group number	New order	Action	Action/ stock	Filled	Filled/ new order	COR	Order quantity	Market order	VS	Stocks/ VS	Inventory	trading desk
1	2,733	9,198	138	2,125	103%	8%	803	278	2	38	70.5%	49
2	10,880	21,981	428	1,001	23%	85%	3,567	2	17	5	62.9%	5
3	20,442	25,183	115	13,462	66%	6%	840	3,143	1	224	72.2%	5
4	19,675	35,783	184	9,712	51%	55%	176	371	5	54	99.5%	45
5	393,034	453,318	865	98,095	41%	10%	201	12,886	12	21	62.7%	3
6	3,523	6,504	153	2,524	56%	42%	534	191	7	7	98.4%	9
7	87,499	161,954	229	24,500	26%	80%	910	541	18	38	30.1%	7
8	20,056	27,538	250	17,369	87%	46%	2,181	2	11	6	8.0%	9
9	148,792	274,893	1,262	25,388	10%	91%	407	0	11	13	6.2%	5
10	7,607	13,854	2,537	1,519	25%	78%	305	0	3	2	11.6%	7

Note: This table shows the trading characteristics for 10 trader groups identified by cluster analysis. All variables are expressed as averages per TD per day during the period from May 20 through May 29, 2013. Here, new order is the number of new order submissions, action is the total number of order messages, action/stock is the number of order messages per stock traded by TD, filled is the number of executed orders, filled/new order is the fraction of filled orders to new orders as a percentage, COR is the fraction of orders cancelled to new orders as a percentage, order quantity is the average number of shares per order, market order is the number of market orders, VS is the average number of VSs used by a TD, stocks/VS is the number of stocks traded through one VS used by a TD, inventory is the inventory at the end of the day, and trading desk is the number of TDs in each group. The computational details are given in the main text.

Table 2. Descriptive statistics of HFT-MMs, HFT-OPs, MFTs, and LFTs (May 20 to May 22)

	New order	Action	Action/ stock	Filled	Filled/ new order	COR	Average order quantity	Market order	VS	Stocks/ VS	Inventory	trading desk
HFT-MM	66,434	122,620	1,768	11,464	19%	83%	347	0	6	7	9.4%	12
HFT-OP	84,438	118,807	373	26,130	50%	59%	1,851	1,770	14	17	32.7%	24
MFT	10,677	20,791	168	5,845	76%	30%	521	442	3	50	85.0%	108
LFT	4,159	7,298	19	2,121	71%	31%	1,695	214	2	347	94.9%	2,177

Note: This table shows the trading characteristics for the four groups: HFT-MMs, HFT-OPs, MFTs, and LFTs. All variables are expressed as averages of the numbers of members in each group during the period from May 20 through 22, 2013. Here, new order is the number of new order submissions, action is the total number of order messages, action/stock is the number of order messages per stock traded by TD, filled is the number of executed orders, filled/new order is the fraction of filled orders to new orders as a percentage, COR is the fraction of orders canceled to new orders as a percentage, order quantity is the average number of shares per order; market order is the number of market orders, VS is the average number of VSs used by a TD, stocks/VS is the number of stocks traded through one VS used by a TD, inventory is inventory at the end of the day, and trading desk is the number of TDs in each group. The computational details are given in the main text.

Table 3. Descriptive statistics of HFT-MMs, HFT-OPs, MFTs, and LFTs in volatile periods

Panel A: May 23

	New order	Action	Action/ stock	Filled	Filled/ new order	COR	Average order quantity	Market order	VS	Stocks/ VS	Inventory	trading desk
HFT-MM	139,449	253,256	3,639	27,899	22%	79%	280	0	6	6	6.6%	12
HFT-OP	129,769	201,705	615	51,077	47%	59%	1,305	2,795	15	17	37.7%	24
MFT	19,439	39,571	253	11,745	83%	28%	543	663	4	52	61.8%	108
LFT	7,222	12,751	30	3,467	80%	32%	2,124	279	2	383	91.3%	2,177

Panel B: May 24

	New order	Action	Action/ Stock	Filled	Filled/ new order	COR	Average order quantity	Market order	VS	Stocks/ VS	Inventory	trading desk
HFT-MM	208,109	388,672	5,568	29,878	21%	81%	290	0	6	7	6.7%	12
HFT-OP	125,908	197,296	661	45,065	73%	51%	1,796	2,458	15	15	35.3%	24
MFT	17,053	37,253	259	9,762	84%	26%	479	574	4	56	58.7%	108
LFT	7,931	14,799	35	3,141	77%	32%	1,888	282	2	384	91.0%	2,177

Note: This table shows the trading characteristics for four groups: HFT-MMs, HFT-OPs, MFTs, and LFTs. All variables are expressed as the averages of the members in each group. Panel A shows the statistics for May 23 and Panel B those for May 24, 2013. Here new order stands for the number of new order submissions, action is the total number of order messages, action/stock is the number of order messages per stock traded by TD, filled is the number of executed orders, filled/new order is the fraction of filled orders to new orders as a percentage, COR is the fraction of canceled orders to new orders as a percentage, order quantity is the average of number of shares per order, the market order is the number of the market order, VS is the average number of VSs used by the TD, stocks/VS is the number of stocks traded through one VS used by a TD, inventory is the inventory at the end of the day, and trading desk is the number of TDs in each group. The computational details are given in the main text.

Table 4. Unit root test

	buy^n_{MM}	buy^c_{MM}	$sell^n_{other}$	imb	inv_{MM}	$fprice$
Level	-7.895^{**}	-8.2037^{**}	-8.3053^{**}	-18.645^{**}	-8.4411^{**}	-3.1052
First difference						-19.199^{**}

Note: This table shows the Dickey–Fuller statistics and p values of augmented Dickey–Fuller tests. The null hypothesis is the presence of a unit root. The variables for this tests are the number of new limit buy orders by HFT-MMs, buy^n_{MM}; the number of cancellations of buy orders by HFT-MMs, buy^c_{MM}; the number of new limit buy orders by groups other than that of HFT-MMs, $sell^n_{other}$; marketwide imbalance, imb; the inventory held by HFT-MMs, inv_{MM}; and the rate of change of the futures mid-quote, $fprice$. All variables are computed for observations aggregated in ten-seconds interval. The sample period is from May 20 to May 24th, 2013. Asterisks indicate statistical significance at the 10% (*) and 5% (**) levels.

Table 5. Correlation matrix of SVAR innovations

	Δbuy^n_{MM}	Δbuy^c_{MM}	$sell^n_{other}$	imb	inv_{MM}	$\Delta fprice$
Δbuy^n_{MM}	1					
Δbuy^c_{MM}	0.709	1				
$sell^n_{other}$	0.294	0.373	1			
imb	0.095	0.004	−0.224	1		
inv_{MM}	0.340	0.315	0.760	0.021	1	
$\Delta fprice$	0.078	−0.044	−0.089	0.413	0.030	1

Note: This table shows the correlations between SVAR variables: the change in the number of new limit buy orders by HFT-MMs, Δbuy^n_{MM}; the change in the number of buy order cancellations by HFT-MMs, Δbuy^c_{MM}; the number of new buy orders by groups other than the HFT-MMs, $sell^n_{other}$; marketwide order imbalance, imb; change in inventory held by HFT-MMs, inv_{MM}; and change in futures prices, $\Delta fprice$. All variables are computed for observations aggregated in ten-seconds interval.

Table 6. Granger causality test

	Δbuy^n_{MM}	Δbuy^c_{MM}	$sell^n_{other}$	imb	inv_{MM}	$\Delta fprice$
Δbuy^n_{MM}		28.97**	7.37**	1.91*	8.8**	3.42**
Δbuy^c_{MM}	2.81**		7.98**	1.87*	6.26**	4.40**
$sell^n_{other}$	14.67**	35.7**		1.50	5.46**	5.58**
imb	8.62**	23.96**	9.42**		4.70**	28.77**
inv_{MM}	25.83**	20.33**	21.58**	1.28		2.36**
$\Delta fprice$	7.26**	7.83**	6.34**	16.14**	4.57**	

Note: This table shows the chi-squared statistics and p values of pairwise Granger causality tests between the endogenous variables of the SVAR. The null hypothesis is that the row variable does not Granger-cause the column variable. The endogenous variables are the change in the number of new limit buy orders by HFT-MMs, Δbuy^n_{MM}; the change in the number of cancellations of buy orders by HFT-MMs, Δbuy^c_{MM}; the number of new limit buy orders by groups other than the HFT-MMs, $sell^n_{other}$; marketwide imbalance, imb; the inventory held by HFT-MMs, inv_{MM}; and change in futures prices, $\Delta fprice$. All variables are computed for observations aggregated in ten-seconds interval. The sample period is from May 20 to May 24, 2013. Asterisks indicate statistical significance at the 10% (*) and 5% (**) levels.

Table 7. Statistics of endogenous variables

	May 20–22 Number of observations: 5,394				Afternoon trading of May 23 and May 24 Number of observations: 1,798			
	Mean	SD	Max	Min	Mean	SD	Max	Min
Δbuy^n_{MM}	-0.3085	120.32	1295	-1329	-0.701	214.05	1072	-1082
Δbuy^c_{MM}	1.722	129.27	1963	-1248	0.378	233.58	1188	-1100
$sell^n_{other}$	1652.73	731.37	12637	686	3489.50	1677.87	17911	1057
imb	11.79	432.94	5894	-6348	-90.05	1085.08	4708	-6683
inv_{MM}	75.45	43.42	635	7	179.57	125.52	999	1
$\Delta fprice$	0.00037	0.01930	0.1280	-0.0655	-0.00470	0.09760	0.4210	-0.4670

Note: This table reports the mean, standard deviation (SD), maximum value (Max), and minimum value (Min) for the endogenous variables of the SVAR system: the change in the number of new limit buy orders by HFT-MMs, Δbuy^n_{MM}; the change in the number of cancellations of buy orders by HFT-MMs, Δbuy^c_{MM}; the number of new limit buy orders by groups other than the HFT-MMs, $sell^n_{other}$; marketwide imbalance, imb; the inventory held by HFT-MMs, inv_{MM}; and change in futures prices, $\Delta fprice$. The four most columns present statistics for May 20 through May 22, 2013, between 9:00:10 and 11:30:00 and between 12:30:10 and 15:00:00. The four rightmost columns present statistics for May 23 and May 24, 2013, between 12:30:10 and 15:00:00. All variables are computed for observations aggregated in ten-seconds interval.