How do High-Frequency Traders Trade?

Nupur Pavan Bang and Ramabhadran S. Thirumalai[1]

1. Introduction

High-frequency traders (HFTs) account for a large proportion of the trading volume in security markets today.[2] Despite this, there is very little understanding of how and why they trade. Are they more likely to demand or supply liquidity? Do they exit the market or increase participation when there is an exogenous information shock? Does their participation increase because they are informed or do they exit because they do not want to suffer losses to informed traders? HFTs typically end the trading day with very low inventory positions. Are they able to manage their inventories better when there is an exogenous information shock or do they end up with positions that are farther away from their preferred inventory positions? Our paper attempts to address these questions by comparing and contrasting their trading behavior around unexpected macroeconomic shocks as well as unexpected firm-level earnings surprises.

Recent studies focus on the impact of algorithmic trading and high frequency trading on various dimensions of market quality. There are studies that find that algorithmic trading improves liquidity and quote informativeness. On the other hand, there are studies that show that algorithmic traders consume liquidity, which leads to wider spreads and worse market quality. Other studies have studied the impact of HFTs, a subset of algorithmic traders, on market quality and efficiency measures. Again there are contrasting results.

While prior studies focus on the impact of HFTs on market quality, there is hardly any information on how HFTs trade. Our study examines how HFTs trade in the BSE (Bombay Stock Exchange) 200 stocks around earnings surprises and macroeconomic shocks in 2011.

---

[1] The authors are with the Indian School of Business, India. Nupur Pavan Bang is at the Hyderabad campus; email: nupur_pavan@isb.edu; phone: +91 (40) 2318 7337 and Ramabhadran Thirumalai is at the Mohali campus; email: ram_thirumalai@isb.edu; phone: +91 (172) 459 1722. The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of NSE, NYU or ISB. We are grateful to Viral Acharya and Sayee Srinivasan for useful comments. We would also like to thank the Bombay Stock Exchange for providing the data through an agreement with the Centre for Analytical Finance at the Indian School of Business. All errors are our own.

[2] HFTs account for around one-third of volume in India and two-third of volume in the U.S. and in Europe (see http://articles.economictimes.indiatimes.com/2014-04-11/news/49058847_1_high-frequency-trading-hft-algorithmic-trading).

2. Our study

To better understand how HFTs trade, we investigate how they trade around exogenous information shocks. Specifically, we examine their trading around firm-specific earnings announcement surprises and unexpected interest rate changes by the RBI. We conjecture that if HFTs demand liquidity, they may have an advantage in processing firm-specific information and hence may trade profitably around earnings surprises. On the other hand, if they largely supply liquidity, they may exit the market around earnings surprises to avoid losing to informed traders. Macroeconomic surprises affect the entire market. HFTs are less likely to have an information advantage around such surprises and hence more likely to withdraw from the market to avoid losses to informed traders.

Using proprietary data from the BSE, we examine the trading behavior of three categories of traders, namely, normal, buy-side algo, and HFT. We identify 153 earnings surprises involving 102 firms and 2 macroeconomic shocks. We use 112 (86) stocks to examine the impact of a larger-than-expected increase in repurchase and reverse repurchase rates announced by the Reserve Bank of India (RBI) on May 3 (July 26), 2011 on HFTs' and buy-side algos' trading behavior. Using a market model with the BSE Sensex as the market index, we estimate the cumulative abnormal returns (CAR) over day 0 and +1 (Day 0 is the announcement date).

We compare the HFTs order handling and trading behavior on the event date (Day 0) to those over a control period, which we define as Days -5 through -1 before the event. This helps us determine how HFTs react when there is an exogenous shock. Do they stop or reduce their participation in markets or are they more likely to demand or supply liquidity?

3. Orders and Trades

We find that HFTs tend to submit more orders (around 150 orders) and larger orders (around 35,000 shares per order) than other types of traders. However, there is no significant difference in the number of new orders and order size between event and control periods around earnings surprises. This suggests that even though there is likely to be higher informed trading during the event window, since he is able to average his profits and losses across his orders, he does not change his order submission strategy. We also find that HFTs do not increase the number of order deletions around earnings surprises. There is some evidence that HFTs marginally reduce the number of modifications around earnings surprises. This is likely due to a larger proportion of his orders getting executed.

HFTs' order handling behavior is starkly different around macroeconomic shocks. We find that they submit a third fewer orders in response to a macroeconomic shock and their order size falls by half. Order deletions do not change significantly but modifications fall by half. Taken in conjunction, these results suggest that HFTs reduce their trading in response to macroeconomic shocks. This is likely because they feel that they are at a disadvantage relative to informed traders and hence reduce the number of orders and order sizes to reduce their losses to informed traders.

Buy-side algos, on the other hand, trade differently from HFTs. They submit more orders and larger orders around earnings surprises. This suggests that they are informed and attempt to profit from their information advantage. Overall, they do submit fewer and smaller orders than HFTs but make far more order modifications. They do not change their order handling behavior much around macroeconomic shocks.

4. Types of orders HFTs use

When there is an exogenous event, do HFTs trade more aggressively by submitting more market orders and aggressively priced limit orders? They may be trading on short-term information and hence trade aggressively. On the other hand, they may act purely as market makers and submit more orders during the event to provide more liquidity and capture more of the bid-ask spread. We find that HFTs do not change the number of market, limit, or stop-loss orders they submit around earnings surprises. On the other hand, they reduce the number of limit orders they submit around macroeconomic announcements by a third. If their limit orders largely supply liquidity, this is consistent with them withdrawing from the market to avoid losing to informed traders. Algos increase the number of both market as well as limit orders around earnings surprises but do not change the number of different types of orders around macroeconomic shocks.

It is possible traders simply make their limit orders more or less aggressive around information shocks without actually changing the number of limit orders. To examine this, we determine if limit orders are liquidity demanding or liquidity supplying. We use the trade data to identify which order triggers the transaction, which we call liquidity-demanding orders. We compare the time stamp of the transaction to the latest time stamp of the two orders involved in the transaction. This latest time stamp could be from when the order was submitted with no subsequent modifications or from the last modification to the order. The order with time stamp closest to the trade is the liquidity-demanding order and the other order is the liquidity-supplying order. For example, a trade occurs at 10:00:00. The buy order was submitted at 9:45:00 with no further modifications. The sell order was first submitted at 9:30:00 and was last modified at 9:59:59. Our rule categorizes the sell order as the liquidity-demanding order and the buy order as the liquidity-supplying one.

We find that HFTs use 51 more liquidity-demanding orders after an earnings surprise than before. This suggests that HFTs do want to trade quickly right after an earnings surprise, which is consistent with them trading on information. Interestingly, we find that HFTs reduce the number of liquidity-supplying orders by more than 50 percent, from 84 orders per day before a macroeconomic shock to 31 orders per day after. They also reduce the number of liquidity-demanding orders by over 40 percent, from 109 orders per day to 62 orders per day after. These results suggest that HFTs are not informed after a macroeconomic shock and hence withdraw from both sides of the market.

We also find that buy-side algos trade more on both sides of the market after an earnings surprise but do not change their order submissions after a macroeconomic shock. Like HFTs, buy-side algos also appear to be informed after earnings surprises.

5. Price and Order Aggressivness

We categorize limit orders by the aggressiveness of their prices. We do this in two different ways. In Panel A, we determine the distance of an order's limit price from the best quote on the same side as the number of ticks.[3] This is calculated as the difference between best bid (limit)

[3]Same side means the bid side of the book for buy orders and the ask side of the book for sell orders.

price and the limit (best ask) price divided by Rs. 0.05 for buy (sell) orders.[4] A negative number implies that the order is priced aggressively resulting in either a partial or full execution or simply improving the best quote on the same side. A zero indicates that the limit price is adding additional depth at the best quote on the same side. A positive number of ticks indicate that the limit order adds additional depth in the book behind the best quote on the same side. We find an increase, though marginally insignificant, in some of the HFTs' aggressively-priced orders after an earnings surprise. Further, consistent with our earlier results, we find that HFTs reduce order submissions at all levels of price aggressiveness after macroeconomic shocks. On the other hand, buy-side algos appear to increase the number of orders at all levels of price aggressiveness both after earnings surprises as well as after macroeconomic shocks.

An alternate measure of order aggressiveness is as follows. Buy (sell) orders with limit price greater (less) than or equal to the best ask (bid) price are the most aggressively priced orders (Category 1) as they result in at least a partial execution. Category 2 is for buy (sell) orders whose limit price is greater (less) than the bid (ask) price but less (greater) than the ask (bid) price. These orders simply improve the current best quoted price on the same side without any execution. Category 3 is for buy (sell) orders whose limit price is equal to the bid (ask) price. They add additional depth to the best quote. Category 4 is for buy (sell) orders whose limit price is at the four lower (higher) prices behind the best bid (ask) price. These orders add additional depth behind the best quotes to the publicly disseminated part of the order book. Finally, buy (sell) orders with limit price less (greater) than the five best prices on the buy (sell) side of the order book are in Category 5 (least aggressive). Our results are similar to those of Price aggressiveness.

6. Trade executions

Since some HFTs submit larger orders, we expect them to trade more. We find that daily gross traded value and number of shares traded is not different for HFTs after an earnings surprise. However, since HFTs submit larger orders, we find that they end up with 67 more trades per day after an earnings surprise, which is a statistically significant increase at the 5 percent level. Buy-side algos increase their trading significantly after an earnings surprise but no change after a

[4]The minimum price variation on the BSE is Rs. 0.05.

macroeconomic shock. This further supports our previous findings that buy-side algos trade on firm-specific information right after an earnings surprise.

7. Inventory Management

Given the larger amount of trading due to the exogenous shock, do HFTs manage their inventories less or more aggressively? The larger volume may make it easier for HFTs to keep their inventory closer to zero at all times of the day. Alternatively, they may see profitable opportunities and be willing to move away from their optimal inventory levels

We use two measures of trader inventory. The first is an intraday inventory balance measure over each day of the control and event period. Lower values of the intraday inventory balance measure show that HFTs aggressively manage their inventory intraday while higher values show that they do not manage their inventory very aggressively. The second measure is the end-of-day inventory balance, which is the ratio of the absolute value of the end-of-day net position to the gross volume traded for that day.

We find that HFTs have the highest intraday inventory imbalance, on average, during the control as well as event period for earnings surprises as well as macroeconomic shocks when compared to the other types of traders. While this suggests that they let their positions deviate substantially during the trading day, there is another explanation for the large level of the intraday inventory balance measure. The BSE 200 stocks constitute the largest stocks in the Indian markets and are cross-listed on the NSE. It is likely that HFTs, given their superior execution systems, manage their inventories close to zero across the two markets. We find that there is no significant change in the intraday inventory measures for HFTs after earnings surprises. However, the measure halves after a macroeconomic shock. This reduction in the intraday inventory measure could be due to their reduced participation in markets after macroeconomic shocks. As they trade smaller quantities, they are less likely to deviate from zero. We find similar results when we study the end-of-day inventory balances. HFTs' end-of-day inventory balance does not change significantly after an earnings surprise, but they halve after a macroeconomic shock.

For buy-side algos, we find that their intraday inventory balance increases after an earnings surprise. This suggests that either they are worse off or are managing their positions across the

BSE and the NSE. Similar to HFTs, we find a reduction in inventory positions after a macroeconomic shock. This may be related to the reduced participation in markets by buy-side algos.

8. Realized Spreads

Next, we examine how much money HFTs make from their trading around earnings surprises and macroeconomic shocks. We measure the information of HFTs by computing the realized spread after each transaction. Overall, the losses of none of the different trader types changes around earnings surprises. Around macroeconomic shocks, HFTs appear to go from negative realized spreads during the control period to marginally positive realized spreads during the event period, though this change is significant only for the 60-minute realized spread measure. This is consistent with them not having any information advantage around macroeconomic shocks. HFTs do not change their trading activity around earnings surprises. However, around macroeconomic shocks, they reduce their trading activity but still make some losses.

Buy-side algos tend to reduce their losses around macroeconomic shocks. They make losses prior to the shock as well as after the shock. This is consistent with them being long-term traders. In the short-run, they tend to bear trading costs (or make losses).

Conclusions

We examine how HFTs respond to exogenous information shocks. Specifically, we examine their order handling and trading behavior, their inventory management, and the profitability around earnings surprises and macroeconomic shocks. We find that HFTs do not significantly change their order handling and trading behavior around earnings surprises but do reduce their participation in the market after a macroeconomic shock. HFTs also do not change the mix of order types that they use around earnings surprises, although they appear to increase the number of liquidity-demanding orders. On the other hand, they use fewer aggressively priced limit orders around macroeconomic shocks. The profitability of their orders does not change around earnings announcements. However, despite reducing their market participation after a macroeconomic shock, they still make losses to informed traders.

We also examine the trading behavior of buy-side algos. We find that they increase their market participation around earnings surprises and do not change it around macroeconomic surprises.