LBS Research Online

[N Hirschey](#)

Do High-Frequency Traders Anticipate Buying and Selling Pressure?

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/1458/>

[Hirschey, N](#)

(2021)

Do High-Frequency Traders Anticipate Buying and Selling Pressure?

Management Science, 67 (6). pp. 3321-3345. ISSN 0025-1909

DOI: <https://doi.org/10.1287/mnsc.2020.3608>

INFORMS (Institute for Operations Research and Management Sciences)

<https://pubsonline.informs.org/doi/10.1287/mnsc.20...>

Do High-Frequency Traders Anticipate Buying and Selling Pressure?

Nicholas Hirschey *
London Business School

January 2020

Abstract: This study provides evidence that high-frequency traders (HFTs) identify patterns in past trades and orders that allow them to anticipate and trade ahead of other investors' order flow. Specifically, HFTs' aggressive purchases and sales lead those of other investors, and this effect is stronger at times when it is more difficult for non-HFTs to disguise their order flow. Consistent with some HFTs being more skilled or more focused on anticipatory strategies, I show that trades from a subset of HFTs consistently predict non-HFT order flow the best. The results are not explained by HFTs reacting faster to news or past returns, by contrarian or trend-chasing behavior by non-HFTs, or by trader misclassification. These findings support the existence of an anticipatory trading channel through which HFTs increase non-HFT trading costs.

*Contact information: Department of Finance, London Business School, Regent's Park, London NW1 4SA, United Kingdom; Email: nhirschey@london.edu. I would like to thank NASDAQ OMX for providing the data and Frank Hatheway, Claude Courbois, and Jeff Smith for many useful discussions about the data and market structure. I would also like to thank Andres Almazan, Fernando Anjos, Robert Battalio, Johannes Breckenfelder, Carlos Carvalho, Jonathan Cohn, Shane Corwin, Andres Donangelo, Cesare Fracassi, John Griffin, Terry Hendershott, Paul Irvine, Pete Kyle, Richard Lowery, Tim Loughran, Jeffrey Pontiff, Mark Seasholes, Clemens Sialm, Laura Starks, Sheridan Titman, and seminar participants at Boston College, the European Finance Association meeting, the Georgia Institute of Technology, the Hong Kong University of Science and Technology, the London Business School, the London School of Economics, Rice University, the University at Buffalo, the University of California Irvine, the University of Georgia, Manchester Business School, the University of Melbourne, the University of Miami, the University of Notre Dame, the University of Texas at Austin, and the Western Finance Association meeting for helpful comments and suggestions.

Investors and regulators have immense interest in the automated strategies used by high-frequency traders (HFTs) and how they affect other investors. Liquidity provision by HFTs has clear benefits. Yet HFTs make many informed trades; more than half their dollar volume on NASDAQ comes from marketable trades that must predict returns to be profitable. The welfare consequences of this informed trading are less clear. In particular, a prominent concern is that one way HFTs predict returns is by using past trade and order data to learn which stocks non-HFTs are about to buy or sell. This information is valuable because an HFT who infers a non-HFT is beginning to liquidate a large position can sell shares now and profit when the price subsequently falls.

This paper examines these issues and provides the first evidence that one reason HFTs appear informed is that they anticipate and trade ahead of non-HFT buying and selling pressure. Specifically, I analyze return and trade patterns around periods of intense marketable buying and selling by HFTs. The imbalance between these marketable purchases and sales is a simple measure of HFTs' directional bets.¹ Anticipatory trading implies HFTs trade in the same direction as their expectation of future non-HFT order flow. Consistent with HFTs using such a strategy, I show that when HFTs sell a stock with marketable orders, this predicts future marketable selling by non-HFTs and lower returns.

This strategy has important implications for liquidity and price efficiency. If a non-HFT is selling to fund a liquidity shock, the HFT's selling harms the non-HFT by depressing their liquidation price (Brunnermeier and Pedersen 2005). If instead the non-HFT is selling because the stock is overvalued, the HFT is effectively reverse engineering the non-HFT's signal (Madrigal 1996, Yang and Zhu 2019, Baldauf and Mollner 2020). Competition between the HFT and non-HFT to trade on the signal causes prices to incorporate this information faster.² But it also implies moderated benefits of HFT participation in price discovery, because some information arriving via HFT trades would soon be incorporated into prices by non-HFTs anyway. Additionally, the competition lowers non-HFT profits, which reduces their incentive to acquire information (Grossman and Stiglitz 1980,

¹A marketable limit order is functionally equivalent to a market order: a buy order with a limit price at or above the best ask or a sell order with a limit price at or below the best bid at the time it enters the order book. NASDAQ requires all orders to have a limit price, and a trade is initiated when a marketable limit order crosses either the best bid or best ask.

²Non-HFTs may at first trade less aggressively to hide from HFTs, slowing price adjustment. But once the HFT has learned how the non-HFT is trading, competition between them ultimately pushes prices closer to fundamental value (Madrigal 1996, Yang and Zhu 2019).

Stiglitz 2014, Weller 2018).³ These concerns were the basis for Michael Lewis’s (2014) book, *Flash Boys*, and the founding of the deliberately slowed down IEX trading platform.

The analysis in this paper uses an entire year of unique trade and trader-level data from the NASDAQ Stock Market. It is an important sample for studying HFTs given NASDAQ was the biggest U.S. exchange by volume and HFTs accounted for 40% of its dollar volume.

In tests where stocks are sorted by HFT net marketable buying at the one second horizon, non-HFT net marketable buying for the stocks bought most aggressively by HFTs rises by a cumulative 65% of its one-second standard deviation over the following thirty seconds. For the median stock, this equates to a 100 share HFT imbalance predicting non-HFTs to buy 13 more shares with marketable orders than they sell. The figures for stocks HFTs sell most aggressively are similar, but in the opposite direction. Turning to returns, future returns for stocks HFTs buy aggressively are positive, while returns for stocks they sell aggressively are negative. The returns and non-HFT trade imbalances predicted by these sorts persist for more than five minutes.

The sorts also show who provides liquidity to these marketable HFT trades. For the stocks HFTs buy with marketable trades, their net buying is also positive in the surrounding seconds. Thus at these times HFTs in aggregate take liquidity from non-HFTs.

The evidence is consistent with HFTs recognizing persistent informed non-HFT order flow in real time. Since liquidity providers are slow to update their quotes, HFTs trade ahead of the impending non-HFT order flow and associated price change to earn a profit. I provide two example patterns that illustrate how HFTs could identify stocks with persistent order flow and prices that update slowly. The examples are simple, interpretable, and use only past order flow and returns as their inputs. Though HFTs use more sophisticated algorithms, the example patterns nonetheless succeed in predicting future non-HFT order flow and returns. And HFT order flow in surrounding seconds is consistent with the patterns capturing predictability used by HFTs.

I next examine whether some HFTs are more skilled at anticipating order flow. There is evidence HFTs follow a variety of strategies (Boehmer, Li, and Saar 2018) and fast HFTs are more profitable (Baron, Brogaard, Hagströmer, and Kirilenko 2019). Consistent with HFTs having heterogeneous skills, there is persistence in which individual HFTs’ trades most strongly predict non-HFT order

³These incentive effects arise when HFTs learn from the endogenous order flow of non-HFTs who pay a cost for their signals. This is absent from models in which the speculative HFT’s trades do not affect the signal itself, such as Foucault, Hombert, and Rosu (2016).

flow. I sort HFTs into three groups each month based on the size of coefficients from regressions of non-HFT order flow on the HFTs' lagged trading (excluding predictability from trend-chasing by non-HFTs and serial correlation in non-HFT order flow). For HFTs in the top third the prior month, a one-standard deviation shock to their net marketable buying predicts a cumulative change in non-HFT net marketable buying equal to 14% of its one-second standard deviation over the next 30 seconds. This compares to an estimate of 2% for HFTs in the bottom third. These top HFTs' trades are also more strongly correlated with future returns. This skill persistence result holds regardless whether individual HFTs are sorted based on correlations with their net marketable buying or total net buying (summing marketable and non-marketable trades).

These findings imply that large non-HFTs have an incentive to disguise their order flow, and in reality we see them deploy execution algorithms to achieve this goal. Yet in the competition between HFT and non-HFT algorithms, non-HFTs are at a disadvantage to the extent that they are constrained by their desire to enter or exit their position. Dark pool operator ITG writes, "some traders are more willing to increase fill rate at the expense of execution quality—the risk of information leakage and impact"(ITG 2013). Indeed, I find the correlation between HFT and non-HFT trades is stronger at times non-HFTs are impatient and thus less focused on disguising order flow: at the market open, on days with high volume, and when trading illiquid or small-cap stocks. This is consistent with Goldman Sachs' acknowledgement their algorithms leak more information in small-caps (Traders Magazine 2013).

I consider several explanations for marketable HFT trades leading marketable non-HFT trades. The effect could be caused by HFTs reacting faster to a signal non-HFTs also observe (Von Beschwitz, Keim, and Massa 2015, Chordia, Green, and Kottimukkalur 2018). But the most probable signals they would both utilize in close succession do not fully explain the results (e.g., past returns, news articles, analyst forecasts and recommendations, management guidance, and form 8-K filings). Controls for returns also show the effect is not due to a reverse causality story in which HFT purchases cause trend chasing non-HFTs to purchase shares as well. Another explanation is trade misclassification. Perhaps NASDAQ incorrectly labels some HFT trades as coming from non-HFTs. Then the lead-lag result may be caused by trades from correctly labeled HFTs predicting those of incorrectly labeled HFTs who follow similar strategies. I evaluate this explanation by comparing the way HFT trading forecasts itself to how it forecasts non-HFT trading. Misclassification implies that

non-HFT predictability comes from it being a noisy proxy for HFT trading, so HFT trading should predict itself better than non-HFT trades. Once you look more than seven seconds into the future, HFT marketable trades forecast non-HFT marketable trades more strongly than HFT marketable trades. By 15 seconds later, HFT marketable trading no longer forecasts itself, whereas it continues to forecast non-HFT marketable trading for more than 30 seconds.

Another question is whether the results are more consistent with HFTs trading ahead of order flow to avoid losses on market making inventory positions or to profit from speculation. To evaluate this question, I examine cumulative HFT net buying for stocks in extreme HFT net marketable buying portfolios before and after the sort second. There is no evidence marketable HFT trades are predominately caused by inventory management. It is possible the marketable trades are disposing of inventory acquired on exchanges other than NASDAQ. However, additional evidence is more consistent with a speculative interpretation. Marketable trades are more than half of HFT volume, which is inconsistent with a pure risk management motive. Additionally, HFT trades lead non-HFT trades more strongly early in the trading day and in small-cap stocks, whereas inventory risk is more binding towards the end of the day and HFT market making is more prevalent in large-cap stocks (Yao and Ye 2018).

This paper makes an important contribution to our understanding of predictability in order flow. In a lower frequency context, Barbon, Di Maggio, Franzoni, and Landier (2019) and Di Maggio, Franzoni, Kermani, and Somnavilla (2019) show evidence of non-HFT brokers leaking information about one non-HFT client’s trades to other non-HFT clients. Clark-Joseph (2013) studies a form of anticipatory trading in S&P 500 futures, though his focus is whether HFTs use small trades to learn about price impact. Recent work using data from Canada (Korajczyk and Murphy 2019) and Sweden (Van Kervel and Menkveld 2019) shows HFTs initially providing liquidity at the start of large multi-hour institutional orders before switching to trade in the same direction as the institutional order after 30 or more minutes.⁴ Using U.S. equities data, this paper complements those studies by showing that if you condition on higher frequency second-by-second data, you observe HFTs taking liquidity ahead of non-HFTs at more favorable prices. I also provide evidence that HFTs’ ability to predict order flow is not limited to the latter part of large orders. Specifically,

⁴Reiss and Werner’s (1998) evidence on dealers trading ahead of worked orders is also related, though in their setting the presumption is the customer intends for the dealer to acquire the position beforehand.

HFT marketable trades predict non-HFT marketable trades more strongly the first half hour of trading, which is necessarily early in a multi-hour order. And the second-level spikes in HFT marketable imbalances coincide with spikes in non-HFT marketable imbalances, which shows the patterns in this paper are distinct from longer-term trends.

These findings relate to a broader debate about the merits of modern market structure.⁵ Research shows that faster trading technology and more HFT liquidity provision are generally associated with improvements in common liquidity measures.⁶ There is also evidence a fair portion of price discovery comes from algorithmic traders in general (O’Hara, Yao, and Ye 2014) and HFTs in particular (Brogaard, Hendershott, and Riordan 2014, 2019). But while the net effect of the transition to electronic markets may be positive, there are also channels through which fast trading can have a negative impact. Examples include increased liquidity provider losses to cross-market arbitrage (Budish, Cramton, and Shim 2015) and to news trading (Menkveld and Zoican 2017).

This paper highlights another negative channel. It shows one way HFTs predict returns is by extracting information from informed non-HFT order flow. This is important, because a non-HFT whose trades are anticipated either trades fewer shares than they desire or obtains executions at a worse average price. These effects are hard to detect with standard liquidity measures. Doing so requires conditioning on the magnitude of a non-HFT’s liquidity shock or, if the non-HFT is informed, the magnitude of their signal. I use a structural VAR to provide a rough estimate. The results imply non-HFT price impact is 14% higher when they compete with HFTs using anticipatory strategies.

1 Hypothesis Development

A speculator able to predict marketable trade imbalances could use this information to predict returns and earn a profit (e.g., Madrigal (1996), Brunnermeier and Pedersen (2005), and Yang and

⁵Other issues include the effects of alternative trading venues (Foucault and Menkveld 2008, O’Hara and Ye 2011, Comerton-Forde and Putniņš 2015), tick size (Chao, Yao, and Ye 2016, Yao and Ye 2018, O’Hara, Saar, and Zhong 2019), high-frequency volatility (Kirilenko, Kyle, Samadi, and Tuzun 2017, Hasbrouck 2018, Egginton, Van Ness, and Van Ness 2016), overinvestment in technology (Biais, Foucault, and Moinas 2015), and agency conflicts caused by exchange fee rebates (Battalio, Corwin, and Jennings 2016).

⁶Examples include Hendershott, Jones, and Menkveld (2011), Boehmer, Fong, and Wu (2015), Hendershott and Riordan (2013), Hasbrouck and Saar (2013), Conrad, Wahal, and Xiang (2015), Conrad and Wahal (2019), Brogaard, Hagstromer, Norden, and Riordan (2016), Jovanovic and Menkveld (2015), Menkveld (2013), Breckenfelder (2013), Brogaard and Garriott (2019), Ait-Sahalia and Saglam (2016), Brogaard, Carrion, Moyaert, Riordan, Shkilko, and Sokolov (2017), and Griffith, Van Ness, and Van Ness (2017).

Zhu (2019)). The idea is supported by evidence that marketable purchases lead to price increases and marketable sales lead to price declines (Hasbrouck 1991). This happens because trades may contain information about fundamental values (Glosten and Milgrom 1985, Kyle 1985) or cause liquidity providers to take risky positions (Stoll 1978).

HFTs are good candidates for speculators who learn about order flow. Doing so requires fast trading technology and automated analysis of quote updates and trades. And HFTs specifically invest in these technologies.

This anticipatory trading hypothesis can be tested by comparing the timing of HFT marketable trades to non-HFT marketable trades. Assume an HFT identifies a signal indicating there will be non-HFT marketable selling in the near future. If the liquidity provider has not conditioned on this signal, then the liquidity provider's bid will be too high. The HFT should sell the stock with a marketable limit order that hits the bid. Afterwards, non-HFTs will similarly sell with marketable orders, driving the price lower and creating a profit for the HFT.

Hypothesis 1 *HFT net marketable buying will be positively correlated with future non-HFT net marketable buying and returns.*

In principle, the HFT could also sell with a non-marketable order, but much of the variation in HFTs' non-marketable orders is caused by liquidity provision. By focusing on HFTs' marketable orders, I instead focus on times when their trades are more likely to be caused by information.

The second main hypothesis seeks to determine whether HFTs are anticipating non-HFTs trading for information or liquidity reasons. Price changes after HFTs' trades provide insight into which of these explanations is more typical.

Hypothesis 2 *If HFTs anticipate informed order flow, price changes predicted by their marketable trades will be permanent, whereas if they anticipate uninformed order flow, the price changes will be temporary.*

There are many plausible ways HFTs could learn about order flow. The one I emphasize is insufficiently disguised positive serial correlation in order flow. Marketable buying is known to have positive serial correlation (Hasbrouck 1991, Chordia, Roll, and Subrahmanyam 2002). This can arise from splitting a large order into small trades executed over time. Order splitting is

generally optimal for a trader who has discretion over when to trade, and it is useful for both informed traders (Kyle 1985) and liquidity traders (Vayanos 2001). But this practice leaves open the possibility that in executing the first part of an order, a large trader reveals more than he intends. A similar type of information leakage can occur when a trader simultaneously sends orders for a stock to different exchanges (e.g., NYSE and NASDAQ), but routing complications cause the orders to arrive at those venues at different times (Lewis 2014). These signals may be profitable to use if liquidity providers do not condition on them when setting bid-ask quotes.

Hypothesis 3 *One way HFTs profit from anticipatory trading is by identifying times when prices have not updated to account for positive serial correlation in non-HFT order flow.*

The final main hypothesis relates to understanding differences between HFTs. Anticipatory trading requires an HFT to be more skilled than some liquidity providers at learning from order flow. Additionally, we do not expect different HFTs to follow the same strategies (Boehmer, Li, and Saar 2018). It follows that anticipatory trading may be concentrated in a subset of the HFTs.

Hypothesis 4 *If certain HFTs are better at forecasting order flow or focus more on the strategy, then trades from these HFTs will be consistently more strongly correlated with future non-HFT trades than trades from other HFTs.*

These hypotheses are not biased by using exclusively NASDAQ data. A finding that HFTs make marketable purchases and sales immediately before non-HFTs on NASDAQ is true regardless of what is happening on other exchanges and regardless of the HFTs' net position across all exchanges. To reverse this finding across all markets, it would be necessary to find the reverse pattern elsewhere. But there is no reason to believe NASDAQ is not representative. It was the biggest exchange, market structure has converged to electronic limit order books everywhere, and mostly the same HFTs and non-HFTs trade on all exchanges.

2 Data

2.1 Sample construction

This study uses intra-day transactions data from the NASDAQ Stock Market. NASDAQ was the largest U.S. exchange by volume during the January to December 2009 sample period,⁷ with a 21.3% share of total volume (NASDAQ 10K 2009). This includes 33.0% of NASDAQ-listed volume and 15.7% of NYSE-listed volume. Like all exchanges, it is structured as an electronic limit-order book and executions primarily come from professional traders.⁸

The 96 sample stocks (48 NASDAQ-listed and 48 NYSE-listed) are randomly selected from a sample universe representative of the stocks traded by actively managed investment funds. I use a random sample instead of all common stocks because of computational constraints. The sample is constructed from single-class CRSP common stocks (share code 10 or 11) at the end of 2008. I exclude the bottom two NYSE size deciles to approximate common definitions of active funds' investable universe (e.g., Russell 3000 or MSCI Investable Market 2500). To ensure sample stocks are fairly liquid, I require greater than \$1 million average daily dollar volume in December 2008 and a stock price greater than \$5 at the end of 2008. From this sample universe of 1,882 stocks, I then randomly select 6 NASDAQ-listed and 6 NYSE-listed stocks from each of the eight remaining size deciles.

Table 1 reports summary statistics for all stock days in the sample. The sample averages 93 stocks per trading day.⁹ Market capitalization ranges from \$22 million to \$125,331 million. On average, 27.2% of the stocks' dollar volume trades on NASDAQ, and this varies little across size portfolios. Comparing small-caps to mid- and large-caps, as market capitalization increases, there are increases in median price (\$14.77, \$25.04, and \$31.37), median dollar volume (\$1.9 million, \$15.5 million, and \$120.2 million), and median HFT share of dollar volume (14.8%, 29.2%, and 40.9%).

⁷I exclude January 27th from the sample, because quote data for NYSE-listed stocks is missing.

⁸Most retail brokerages have contracts with market making firms who pay for the right to fill retail orders. For example, in Q3 2009 E*Trade routed nearly all its customers' market orders and over half its customers' limit orders to either Citadel or E*Trade's market making arms (E*Trade 2009). See Battalio and Loughran (2008) for a discussion of these relationships.

⁹ The number of stocks varies for two reasons: first, to ensure tick size does not constrain price movements, stocks are temporarily removed any day the prior close is less than \$1; second, to remove stocks that stop trading due to bankruptcy or acquisition, stocks are permanently removed from the sample after a day with dollar volume less than \$100,000. These filters use past rather than future data to avoid look-ahead biases.

2.2 Background on HFTs

The firms typically described as HFTs are trading firms using high-turnover automated strategies. Examples include Tradebot Systems, Inc., and GETCO. These firms are remarkably active traders. On their websites, Tradebot says they often account for more than 5% of total U.S. equity trading volume, and GETCO says they are “among the top 5 participants by volume on many venues” (Tradebot 2010, GETCO 2010). They are a large part of the market. The TABB Group estimated HFTs accounted for 61% of U.S. Equity share volume in 2009 (Tabb 2009).

2.3 HFT identification

The NASDAQ trade data classifies firms as either an HFT or a non-HFT. NASDAQ does not disclose the identities of the HFTs or exact criteria used to determine the classification. This HFT versus non-HFT classification is the same as in Brogaard, Hendershott, and Riordan (2014), though their dataset is different. NASDAQ identifies HFTs using a variety of qualitative and quantitative criteria. The firms classified as HFTs typically use low-latency connections and trade more actively than other investors. Their orders have shorter durations than other investors, and they show a greater tendency to flip between long and short positions in a stock during a day.

A disadvantage of the classification scheme is that it is largely done at the broker level. As a result, HFT orders routed through the standard trading system of a broker that handles non-HFT orders, such as Goldman Sachs or Morgan Stanley, will be misclassified as non-HFT (Brogaard, Hendershott, and Riordan 2014). This is a concern primarily for HFTs that are internal divisions of larger investment banks. Independent HFTs who act as their own broker or who use specialist brokers to co-locate servers and retain direct control of their routing infrastructure should be classified correctly. Section 4.3 evaluates whether potential misclassification has an important effect on the empirical results.

2.4 Construction of returns and trade imbalances

I calculate intra-day returns using cleaned bid-ask midpoints from the National Best Bid and Best Offer (NBBO).¹⁰ The NBBO aggregates quotes from all exchanges, so it is the best measure of a

¹⁰ The cleaning procedure follows guidelines in Hasbrouck (2010). First, I remove quote updates where the bid on one exchange is greater than the ask on another. This can be a reporting artifact due to latency or a temporary

stock’s quoted price.

I use all trades during regular continuous trading hours (no open or closing cross) to construct net marketable buying and net buying imbalances. Marketable trades consist of all liquidity-removing sides of transactions. Non-marketable trades include all liquidity-providing sides of transactions (both visible and hidden orders). The NASDAQ matching engine makes this marketable/non-marketable designation when a trade is executed. All orders are limit orders, and a trade is initiated when someone submits a limit order that crosses the spread. The matching engine marks the order from the party who initiated the transaction as a marketable order and the other party’s order as non-marketable.

A net marketable buying imbalance, defined as shares in buyer-initiated trades minus shares in seller-initiated trades, is a common measure of buying and selling pressure from the existing literature (e.g., Chordia, Roll, and Subrahmanyam 2002). The net buying imbalance is shares bought minus shares sold and is used to measure position changes of different investor groups. To put trade imbalances on a similar scale across stocks, I normalize all imbalance measures by a stock’s 20-day trailing volume from CRSP.

Panel B in Table 1 summarizes trade imbalances for the sample stocks prior to their being normalized by trailing volume. The table describes the distribution of the stock-day standard deviations of HFTs’ net buying, their net marketable buying, and non-HFTs’ net marketable buying. The mean standard deviation of HFTs’ net buying (HFT_{NB}) for all stock days is 83 shares, compared to 80 shares for their net marketable buying (HFT_{NMB}). These figures are slightly smaller than the 100 shares that O’Hara, Yao, and Ye (2014) report as the median trade size on NASDAQ in 2008 and 2009. The average standard deviation of non-HFTs’ net marketable buying ($non-HFT_{NMB}$), at 125 shares, is somewhat higher than that of HFTs. The wide variation in im-

arbitrage that gets corrected within fractions of a second. Second, I remove quotes where the bid-ask spread is more than 20% greater than the bid-ask midpoint. This 20% filter removes rarely occurring stub quotes (e.g., bid at \$0.01 and ask at \$20.00) that would cause large fluctuations in the bid-ask midpoint that is not reflective of a true price change. The 20% relative spread cutoff for what constitutes reasonable quotes comes from the SEC’s initial single-stock circuit breaker rule. The rule halted stocks after a 10% return in 5 minutes to prevent “clearly erroneous transactions” (SEC 2010). This is consistent with a 20% relative spread limit, because a stock with successive trades first at a negligible spread and then at a 20% relative spread (i.e., bid/ask 10% away from prior trade) would trip the circuit breaker. When an invalid quote is removed due to either filter, I use the last valid midpoint for the price until a new valid midpoint is available. A new midpoint is typically available within a few milliseconds. To remedy bad pre-market quotes in the NYSE data, the last of which is used to proxy for the opening price, I throw out the last price before the open if there is more than a 20% difference between the last pre-open bid-ask midpoint and the first post-open bid-ask midpoint.

balance standard deviations between size portfolios motivates the normalization by trailing volume in later results.

3 Tests of the anticipatory trading hypotheses

This section tests whether HFT trades forecast non-HFT trades and then examines additional empirical implications of the pattern being caused by anticipatory trading.

Throughout the paper, I use a method similar to Fama and MacBeth (1973) to account for cross-sectional correlation across stocks and autocorrelation through time. Observations in nearby seconds are potentially correlated, so they should not be treated as independent. To be conservative, I collapse all estimates on a day to a single average estimate for that day. This gives me a time-series of the average effect each day. I then base hypothesis tests on the mean of that daily time series. In effect, the question I ask is whether the estimate is consistently different from zero throughout the days in the sample. I do not assume observations from the same day are independent. I only assume different days are independent. Of course, there is likely some residual correlation between nearby days, so I use Newey and West t -statistics with the optimal lag length for the daily series chosen following Newey and West (1994). Despite the conservative approach to standard errors, many of the tests show high t -statistics. This happens when the estimated effect is consistent across the days in the sample.

3.1 Do trades from HFTs lead trades from non-HFTs?

Anticipatory trading implies an increase in HFT net marketable buying should forecast an increase in non-HFT net marketable buying and an increase in returns (Hypothesis 1). The first tests use portfolio sorts to examine the hypothesized relationships in a simple and easy to interpret form. Afterward, vector autoregressions (VARs) evaluate whether HFT trading has any predictive power beyond that of past returns and serial correlation in non-HFT trading.

The sorts identify stocks with large HFT net marketable buying imbalances in a given second. I then look at non-HFT net marketable buying and returns afterward.

Every stock is assigned to one of ten portfolios each second based on the stock's HFT net marketable buying imbalance. Portfolio breakpoints are calculated from non-zero observations

during the prior trading day. Then, a daily mean of the variable of interest for each portfolio is calculated by taking an average among all stock-second observations for the portfolio that day:

$$V_{d,p} = \frac{1}{N_p} \sum_{i,t} V_{d,p,i,t},$$

where V is the variable of interest, d indexes days, p indexes portfolios, N_p is the number of stock-second observations in the portfolio that day, i indexes stocks, and t indexes seconds. Hypothesis tests are based on the means of these daily time series.

Figure 1 plots cumulative net marketable buying for sort deciles ten (stocks HFTs bought) and one (stocks HFTs sold). Because the figure plots cumulative values, positive net marketable buying causes the line in that second to rise, while negative net marketable buying causes the line to fall. The dark red lines indicate HFTs' cumulative net marketable buying plotted on the left y-axis. Looking at the stocks HFTs buy, HFT net marketable buying is slightly positive before the sort, spikes in the sort period, and then is relatively flat afterwards. The sort-period spike is expected, because the line is the cumulative value of the variable used to form the portfolios. The real question is what happens to non-HFT net marketable buying afterward.

Consistent with Hypothesis 1, Figure 1 shows that in the seconds after these sorts on HFT net marketable buying, non-HFT net marketable buying is positive for the stocks HFTs bought and negative for the stocks HFTs sold. Cumulative non-HFT net marketable buying is depicted by dashed blue lines plotted on the right y-axis. In the seconds after the sort, cumulative non-HFT net marketable buying for the stocks HFTs bought steadily increases. This indicates consistently more marketable purchases than sales—non-HFTs buy the stocks HFTs previously bought. HFTs also buy shares after the sort, but in comparison to non-HFTs, they buy relatively more in the sort second than afterward.

Table 2 presents this sort data in a form conducive to hypothesis tests. Panel A shows that in the first thirty seconds after the sort, cumulative non-HFT net marketable buying rises to 0.65 times the one-second standard deviation (or 27 shares for the median stock). Over the next four and a half minutes, it continues rising to 1.15 (or 48 shares for the median stock). These values are both significantly different from zero, with t -statistics of 14.35 and 2.41. The picture for the stocks HFTs sold most aggressively is symmetric; non-HFT net marketable buying in decile one is

−0.69 and −1.89 times the one-second standard deviation in the first 30 seconds and the first five minutes after the sort respectively.

If the order flow patterns in Figure 1 and Table 2 are due to HFTs anticipating non-HFT buying and selling, then we should also see that the stocks HFTs bought aggressively have positive future returns and the stocks they sold have negative future returns. Figure 2 shows that in the 30 seconds after the sort, prices of stocks in HFT net marketable buying decile ten increase by 1.23 basis points, while the prices of decile one stocks decline by −1.07 basis points. Table 2B shows both portfolios’ returns are significantly different from zero, with t -statistics of 11.78 and −13.48. These price changes will be permanent if HFTs anticipate order flow from informed traders and temporary if they anticipate liquidity traders (Hypothesis 2). The last two columns in Table 2B show partial reversal of these returns over the next four and a half minutes. Decile ten reverses from 1.23 basis points 30 seconds after the sort to 0.57 basis points five minutes afterwards; decile one reverses from −1.07 to −0.58 basis points. Nonetheless, the return spread between these two portfolios remains positive. This is consistent with Brogaard, Hendershott, and Riordan’s (2014) findings, using a state-space model, that HFT marketable trades forecast permanent price changes. Thus, the evidence is more consistent with HFTs anticipating informed rather than uninformed order flow.

Figure 3 examines who provides liquidity to these marketable HFT trades. In Panel A, stocks are sorted into deciles by HFT net marketable buying the same as in Figure 1. The only difference is that I now plot cumulative HFT net buying, which aggregates both marketable and non-marketable trades. A comparison of net marketable buying to net buying reveals who provides liquidity. For example, if one HFT uses a marketable order to purchase 100 shares from another HFT, then net marketable buying across the two HFTs is 100 shares. But net buying is 0, because the first HFT’s purchase and the second’s sale cancel each other out.

Panel A shows that HFTs in aggregate take liquidity from non-HFTs in these sorts. In the sort second, net buying for decile ten rises by 5.77, while it falls by 5.83 for decile one. Averaging the two and multiplying by its median one-second standard deviation (28 shares) gives an average 162 share net position change for either decile. The equivalent calculation for HFTs’ net marketable buying in Figure 1 gives a 213 share net marketable imbalance ($\frac{8.19+8.21}{2} \times 26$). Therefore, 24% of the net liquidity provision comes from HFTs ($\frac{213-162}{213}$) and 76% comes from non-HFTs. This compares to

HFTs' 37% share of liquidity provision by dollar volume in the sample. The evidence is consistent with HFT liquidity providers stepping away, leaving less-skilled non-HFT liquidity providers to fill informed HFT orders. Panel B explores differences between HFTs, which I defer until Section 3.4.

The sorts provide a few additional insights. First, the order flow and return results are stronger in small-cap stocks. In the first 30 seconds after the sort, the difference between non-HFT net marketable buying for the stocks HFTs bought (decile ten) versus sold (decile one) is 2.66 for small-cap stocks, compared to 1.34 for all stocks. Similarly, the decile ten minus decile one return spread is 5.18 basis points for small-cap stocks, compared to 2.30 for all stocks. Both effects persist through five minutes after the sort. The larger post-sort spread in smaller stocks could be due to non-HFTs having a harder time disguising order flow when trading relatively illiquid stocks. Section 3.5 examines this illiquidity hypothesis in more detail. Second, stocks HFTs buy with marketable orders have higher returns the prior 30 seconds (4.43 basis points) than the stocks they sell (−4.30 basis points). The trend chasing is notable, because HFTs are often characterized as liquidity providers. But this pattern is more consistent with these trades demanding liquidity than providing it. Third, that HFT trades predict non-HFT trades at horizons much longer than a second shows the result is not an artifact of cross-market arbitrage. This type of arbitrage is understood to occur at sub-second horizons (see Budish, Cramton, and Shim 2015).

Table 3 builds on the sort results using a vector autoregression (VAR) similar to Hasbrouck (1991). The portfolio sorts show evidence of positive serial correlation and trend-chasing in non-HFT net marketable buying. These effects provide a simple way for an HFT to predict non-HFT trades. The VAR isolates the predictive ability of HFTs' marketable trades in excess of anything coming from these simple signals.

The VAR is a system of three equations in which lags of returns, HFT net marketable buying, and non-HFT net marketable buying are all used to explain each other:

$$R_t = \alpha_1 + \sum_{i=1}^{10} \gamma_{1,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{1,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{1,i} R_{t-i} + \epsilon_{1,t} \quad (1)$$

$$HFT_{NMB,t} = \alpha_2 + \sum_{i=1}^{10} \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \quad (2)$$

$$non-HFT_{NMB,t} = \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t} \quad (3)$$

where R_t is the one-second NBBO midpoint return, $HFT_{NMB,t}$ is the one-second HFT net mar-

ketable buying, and $non-HFT_{NMB,t}$ is the one-second non-HFT net marketable buying. The equations are estimated for each stock every day. All variables are standardized by their standard deviation among all stocks that day. Table 3 averages all stocks' coefficients on the same day and then reports significance tests based on the mean of the resulting daily time series (Fama and MacBeth 1973).

The main question is whether HFT net marketable buying is still positively correlated with future non-HFT net marketable buying and returns. Table 3 shows both positive correlations still exists in the VAR. In the equation forecasting non-HFT net marketable buying, the lag one coefficient on HFT net marketable buying is 0.0007, rising to 0.0021 at lag two and then declining slowly to 0.0016 at lag ten. Though the first lag coefficient is the smallest, they are all significantly different from zero, and the decay in predictability from lags two to ten is slow. Similarly, in the equation forecasting returns, a one-standard deviation increase in HFT net marketable buying leads to a 0.0185 standard deviation increase in the next second's return, with a t -statistic of 10.92. Coefficients remain positive and significantly different from zero up through the ninth lag.

A related question is whether HFT trades predict non-HFT trades better than non-HFT trades predict HFT trades. If HFTs learn about order flow from past trades, then there will be some correlation between past non-HFT trades and future HFT trades. However, the evidence for HFTs trading ahead of predicted order flow is strongest if HFT trading predicts non-HFT trading more strongly than non-HFT trading predicts HFT trading.¹¹

Table 3 shows HFT net marketable buying does in fact predict non-HFT net marketable buying more strongly than the reverse. The lag one coefficients at first seem to contradict this statement; in the HFT forecasting equation, the coefficient on non-HFT net marketable buying is bigger (0.0028) than the HFT coefficient in the non-HFT forecasting equation (0.0007). However, this flips at higher lags. Putting all ten lags together, the non-HFT coefficients sum to -0.005 when predicting HFT net marketable buying, compared to a sum of 0.016 for the HFT coefficients when predicting non-HFT net marketable buying.

This section evaluated lead-lag relationships between HFT and non-HFT net marketable buying

¹¹Note that HFTs do not need to trade earlier than non-HFTs to profit from figuring out what the non-HFTs will be buying or selling (Brunnermeier and Pedersen 2005). For instance, if an HFT learns that informed non-HFTs are buying a stock, the HFT can buy the stock at the same time as the informed non-HFTs. In this scenario, the HFT is profiting from anticipating informed non-HFT order flow, but they are not necessarily trading before the non-HFTs.

using portfolio sorts and VARs. Both sets of tests show an increase in HFT net marketable buying is followed by an increase in non-HFT net marketable buying and returns. The effects are strongest in small-cap stocks and consistent with HFTs anticipating informed order flow.

3.2 Example patterns HFTs could use to predict order flow

This section provides insight into the types of patterns that allow HFTs to infer non-HFT order flow. There are many possibilities, and an exhaustive search is not feasible. However, Table 4 provides two illustrative examples that predict non-HFT trading and returns. In practice, I expect HFTs to use sophisticated machine-learning algorithms. But I use simple rule-based sorts so that they are easy to interpret.

Both patterns are based on the idea of finding times when a stock's price has not fully adjusted to persistent non-HFT order flow (Hypothesis 3). This type of pattern is consistent with Figure 1, which showed HFT marketable buying spike near the start of a period of persistent non-HFT marketable buying. To implement the test, stocks are sorted into *Buy* and *Sell* portfolios at the end of second t based on non-HFT net marketable buying ($non-HFT_{NMB}$) and returns (R) from seconds $t-5$ to $t-1$. Waiting until the end of second t to do the sort ensures an HFT would have up to a second between observing the sort variables and making a trading decision. The only difference between the two patterns is the thresholds used to identify serial correlation in lag non-HFT order flow. To form the *Buy* portfolio, the first pattern requires $non-HFT_{NMB,t-1} > 0$ and $\sum_{l=1}^5 non-HFT_{NMB,t-l} > 2$. The second requires $non-HFT_{NMB,t-1} > 1$ and $\sum_{l=1}^5 non-HFT_{NMB,t-l} > HFT_{NMB,t-1}$. They both then require $\prod_{l=1}^5 (1 + R_{t-l}) - 1 \geq 0$ and $R_{t-1} = 0$ to mark times the price is slow to update to non-HFT order flow at $t-1$. Restrictions for identifying the *Sell* portfolio are the same but with signs and inequalities flipped.

Table 4 shows the two patterns are both successful in predicting higher non-HFT net marketable buying and higher returns for the *Buy* portfolio than the *Sell* portfolio. For the first pattern, Panel A shows that the first 30 seconds after the sort, the *Buy*–*Sell* portfolio difference is 5.43 for non-HFT net marketable buying and 3.32 basis points for returns. Panel B shows these differences for the second pattern are 8.44 for non-HFT net marketable buying and 3.35 basis points for returns. In both cases, the differences are significantly different from zero and widen through five minutes after the sort.

These findings indicate they are feasible patterns HFTs could use to predict order flow and returns. Additionally, in both panels HFT net marketable buying is higher for the *Buy* portfolio than for the *Sell* portfolio in the seconds surrounding the sort. This is consistent with the patterns capturing non-HFT predictability that HFTs actually use.

3.3 How is non-HFT price impact affected by anticipatory trading?

An HFT able to anticipate a non-HFT's trades will compete with that non-HFT to obtain executions. This increases the non-HFT's price impact, because for every 100 shares they trade, there is the direct impact of those 100 shares plus the impact of competing HFT trades. I estimate the effect of this competition using a structural VAR (Hasbrouck 1991). This allows me to estimate the price impact of a 100 share shock to non-HFT net marketable buying. I then provide a comparison to price impact at times when there is no competition from HFTs. The comparison gives an estimate of the additional price impact from HFTs' anticipatory trades. Albeit with the caveat that identification comes from the structural specification rather than from use of an instrument that specifically identifies times HFTs are not engaged in anticipatory trading. The estimate implies competition from HFTs raises non-HFT price impact by about 14%. The results are in Figure 4.

Like the VAR used elsewhere in the paper, the price impact VAR has an equation for returns (R), for non-HFT net marketable buying ($non-HFT_{NMB}$), and for HFT net marketable buying (HFT_{NMB}). And the regressions include lags of all three as explanatory variables. But the price impact VAR includes additional terms to capture contemporaneous price impact and competition. First, all equations include lags of a competition variable (C). When HFT and non-HFT net marketable buying are both positive or both negative, they are competing for the same liquidity. At these times, the competition variable is $C_t = sign(HFT_{NMB,t}) \times |HFT_{NMB,t}| \times |non-HFT_{NMB,t}|$ (the *sign* term makes it positive for purchases and negative for sales). Otherwise $C_t = 0$. Second, the return equation includes contemporaneous values of HFT trading, non-HFT trading, and the competition variable. This captures the effect of time t trades on time t returns. Third, the HFT equation contains contemporaneous non-HFT trading as an explanatory variable. The intuition for this timing is that HFTs submit anticipatory trades in second t based on their expectation of what non-HFT order flow will be that second. It is important to account for this within-period correlation if HFTs and non-HFTs compete to trade on the same signal.

The impulse response function for returns measures price impact. It shows the cumulative effect on returns of a 100 share shock to non-HFT net marketable buying. The lagged variables in the VAR control for heterogeneity before the shock. I estimate two impulse response functions. The first shows non-HFT price impact using the full VAR, which incorporates the effect of competing with HFTs. The second is an estimate for a counterfactual world with no HFT competition. It sets HFT_{NMB} in every second to zero and measures price impact if *non-HFT* $_{NMB}$ each second after the shock is the same as in the first scenario. With a 10-lag VAR, Figure 4 shows the estimated non-HFT price impact when competing with HFTs is 0.281 basis points the second of the shock, accumulating to 0.846 basis points after 30 seconds. The estimate from the counterfactual with no HFTs is 0.274 basis points the second of the shock and 0.743 basis points 30 seconds later. The difference between the two estimates is statistically significant by the first second after the shock. It implies HFT competition raises the price impact of non-HFT net marketable buying by 14% ($= 0.846/0.743 - 1$). I also extended the VAR to 30 lags, and the implied effect of competition was still around 14% ($1.130/0.994 - 1$).

3.4 Is predicting order flow related to HFT skill?

This section tests first whether some HFTs are better at predicting order flow (Hypothesis 4). It then examines if the HFTs who predict order flow the best are also better at predicting returns. The first task is to identify the HFTs who are the best at predicting order flow.

I use regressions similar to the non-HFT forecasting regression in Table 3. The difference here is the regression is estimated separately for each HFT in the sample. There are two regression estimates for each HFT every day. In the first, the HFT variable is the HFT’s net marketable buying. In the second, the HFT variable is the HFT’s net buying. HFT net buying is more applicable here than in the aggregate results, because these estimates are calculated for each HFT separately. As a result, we do not have to worry that variation in net buying by HFTs anticipating order flow is obscured by variation in net buying caused by different HFTs providing liquidity. The heading for Table 5 contains the regression equation.

High-frequency traders’ ability to predict buying and selling pressure is measured in two ways: first, by the average coefficient on the first lag of the HFT’s net marketable buying or net buying, and second by the average sum of the coefficients on all ten lags of their net marketable buying or

net buying. I take the mean of each ability measure across all days in a month for each HFT and sort the sample into three groups based on the magnitude of the HFTs' ability measures. I define the HFTs who are skilled at predicting order flow as those with the largest positive coefficients.

One simple way to look at consistency is to look at the probability an HFT in the highest correlation group remains in that group in future months. Figure 5 plots the probability an HFT who is in the highest-correlation group will again be in the highest-correlation group one, two, and three months later. There are three groups, so under the null hypothesis of no persistent difference among HFTs, only 33.3% of the HFTs should still be in the highest-correlation group one month after the sort. In fact, for all sort methods, between 57% and 72% of the HFTs are still in the highest-correlation group one month later. Similarly, in months two and three, more HFTs are still in the high group than would be the expectation under the null hypothesis of no persistence.

Another way to examine persistence is to compare post-sort month coefficients for the three HFT groups. Table 5 reports post-sort month coefficients for these groups. The left set of columns is for when the HFT variable is net marketable buying, and the right set is for when the HFT variable is net buying. The top group of rows examines the persistence of the coefficient on the first lag of the HFT variable, $\overline{\gamma_{d,t,1}}$. When the HFT variable is net marketable buying, the average $\overline{\gamma_{d,t,1}}$ for the highest-correlation group is 0.024, compared to -0.005 for the lowest-correlation group. Both the Newey and West (1994) p-value and non-parametric rank sum p-value indicate this difference is significantly different from zero. The left group of columns in the middle rows of Table 5 show results using ten lags of HFTs' net marketable buying. As was the case for the test using just the first lag, the difference between the highest and lowest-correlation groups is significantly different from zero. The right group of columns in Table 5 report tests using HFTs' net buying. Results from these tests are the same as for the net marketable buying tests. Finally, the bottom rows show the sum of HFT coefficients when the VAR is extended to 30 lags, and the results are similar. Whether one looks at coefficients on the first lag, first ten lags, or first 30 lags of HFT net marketable buying or net buying, there are persistent differences between the highest and lowest-correlation HFT groups. These results indicate that some HFTs' trades are consistently more strongly correlated with future non-HFT order flow.

Having split HFTs into groups, I revisit the question of who provides liquidity to HFTs' marketable trades when they buy or sell a stock aggressively. Prior results aggregating all HFTs showed

that liquidity mostly came from non-HFTs. But perhaps certain HFTs are consistently on the other side of these trades. Panel B in Figure 3 examines cumulative HFT net buying with HFTs split by the coefficient on the first lag of their net marketable buying in the 10-lag regression. The sort is now on net marketable buying aggregated among HFTs in the *High* group. If HFTs in the *Mid* or *Low* groups consistently provide liquidity to the *High* HFTs, then their net buying should be negative in the second when *High* HFTs are buying stocks with marketable orders. However, *Mid* and *Low* group HFTs are actually slight net buyers in the sort second when *High* HFTs buy aggressively (and vice versa for the *Sell* portfolio).

Next, Table 6 examines whether marketable trades from HFTs who are skilled at predicting order flow (those in the *High* group) also forecast larger returns. HFTs are again split into three groups using the methodology from Table 5. Then, trades are aggregated among HFTs in each group, resulting in one time-series of aggregated trades for each HFT group. Returns are then alternately regressed on ten lags of each aggregate HFT series, controlling for ten lags of returns and non-HFT net marketable buying. Thus, the regressions identify the HFT groups' ability to forecast returns that is independent of information in past returns and non-HFT order flow. For brevity, I only report results for the split on the first lag of HFT net marketable buying in the 10-lag regression.

The results indicate marketable trades from HFTs who are skilled at predicting order flow are better at predicting returns. Coefficients on the first three lags of HFT net marketable buying in the *High* HFT regressions are 0.0328, 0.0210, and 0.0160, compared to 0.0268, 0.0160, and 0.0125 in the *Low* group. Both the means and medians of the two coefficient time-series are significantly different from each other, as are the differences in the sum of all 10 lags. These findings are consistent with HFTs in the *High* group being more skilled at predicting short-term returns.

3.5 Are non-HFTs' trades more predictable when they are impatient?

Non-HFTs have an incentive to disguise their order flow from HFTs. But as the ITG quote on page 3 highlighted, there is a tradeoff between disguising order flow and trading a large position quickly. It follows that non-HFT order flow may be more predictable when they are impatient. This section tests the hypothesis by comparing VAR estimates from normal times to estimates from times when non-HFTs are hypothesized to be impatient. The VAR is the same as in Section 3.1. I focus

on the equation where the dependent variable is non-HFT net marketable buying. If HFTs are doing relatively more anticipatory trading, then I expect larger positive coefficients on the lagged HFT imbalance variables. I examine several times when I expect non-HFTs to be impatient: at the market open and close, on days with high volume or high trade imbalances, and when trading illiquid stocks.

A non-HFT may be impatient at the open if they received an overnight signal (e.g., news) and are concerned other investors either already have or will soon receive the same signal. Similarly, investors near the close may have private information about an end-of-day news announcement. Or they may be facing a liquidity shock that needs to be funded before the close. I use the first and last half hour of the trading day as the opening and close periods. The control period is the rest of the day.

High volume or imbalance days are likely days when certain investors are trading large positions. And it is harder to hide with noise traders when trading a large position. Thus I compare stock days with high volume or a high value of the absolute net marketable buying imbalance to other days. Inspired by Gervais, Kaniel, and Mingelgrin (2001), I classify a day as high volume or high imbalance if the volume/imbalance is among the top 10% of days the past month (i.e., rank relative to prior 19 trading days is 19 or 20).

The motivation for examining illiquid stocks is similar to the intuition for volume and imbalances. If non-HFTs do not perfectly scale position sizes relative to liquidity, then in illiquid stocks, their trades will be proportionately larger and thus harder to hide. I identify illiquid stocks as those with wide bid-ask spreads or wide relative spreads ($= \text{Bid-Ask Spread} \div \text{Bid-Ask Midpoint}$). Relative spreads account for a \$0.01 spread implying different levels of liquidity for a \$10 stock and a \$100 stock.

Table 7 shows that non-HFT trade predictability is reliably stronger at the market open, on high volume days, and in illiquid stocks. Both the coefficient on the first lag of HFT net marketable buying and the sum of all ten lag coefficients are significantly higher at these times non-HFTs are expected to be impatient. For comparison, the sum of all ten lag coefficients is 0.0076 higher at the open than during the middle of the day, 0.0086 higher on high-volume days than other days, 0.0087 higher in stocks with wide versus narrow quoted spreads, and 0.0046 higher in stocks with wide versus narrow relative spreads. On high imbalance days, the first and sum of all ten coefficients are

also higher than other days, but the t -statistics for the differences with other days are not quite as strong (2.13 and 2.10). The one test that does not go as predicted is the test examining trading near the close. Predictability near the close is weaker than the middle of the day, perhaps because HFTs do not want to build an inventory position they cannot unwind before the end of the day. Overall, the evidence is consistent with HFTs having an easier time anticipating order flow when non-HFTs are more constrained when trying to disguise their trades.

4 Alternative explanations

While the anticipatory trading hypothesis predicts that HFT liquidity demand will be positively correlated with future non-HFT liquidity demand, there are other possible explanations. Any signal that HFTs and non-HFTs separately observe and then trade on could result in HFT liquidity demand positively forecasting non-HFT liquidity demand, particularly if HFTs react faster. This contrasts with anticipatory trading, in which the HFT does not directly observe the same signal as the non-HFT but instead infers it from the non-HFT's trading activity.

It is difficult to completely rule out this alternative, because it is not possible to fully examine all inputs to HFT and non-HFT trading processes. Nevertheless, if the relation between HFT and non-HFT trading is still strong after accounting for the most likely signals they would both utilize in close succession, then it makes this correlated signal explanation less probable. To examine the issue, this section explores the two most likely signals HFTs and non-HFTs could both be reacting to: news and past returns. It then evaluates whether the results could be caused by another correlated signal explanation—trader misclassification.

4.1 Are HFTs simply reacting faster to news?

Reacting to news faster than other investors is an important HFT skill and plausible explanation for HFT marketable trades leading non-HFT marketable trades (Ye, Yao, and Gai 2013, Foucault, Hombert, and Rosu 2016).

Alternative Hypothesis 1 *If HFT trades lead non-HFT trades because HFTs are reacting faster to firm news, then the lead-lag relationship should go away after excluding days with firm news.*

News days are identified using a comprehensive set of events: days when an article about the firm is published, an analyst covering it changes their earnings forecast or investment recommendation, management provides new guidance, the firm files an SEC Form 8-K indicating a material event, or the absolute value of the stock’s market-adjusted return is large. The news articles come from Factiva, data for analysts and guidance from Thomson IBES unadjusted detail files, and 8-Ks from the SEC EDGAR website.¹² The news article, analyst, guidance, and 8-K events cover the most plausible public news signals HFTs and non-HFTs would both observe. Factiva for example contains articles from over 35,000 sources, including most major newswires, newspapers, and magazines such as *The New York Times*, *Wall Street Journal*, *Dow Jones Newswire*, and *Reuters*.¹³ Nonetheless, it is possible HFTs and non-HFTs respond to a signal that is not in one of these databases. The return-based news day measure addresses this concern. If there is an unobserved news event, trading on the signal will affect the stock price. Thus, examining trading after excluding days with absolute market-adjusted returns greater than relatively small thresholds (1%, 0.5%, and 0.25%) minimizes the chance any observed pattern is caused by unobserved news events.

Table 8 tests the news hypothesis using the Table 3 VAR estimated on days when there is no news about the firm. The table shows the positive correlation between non-HFT and lagged HFT net marketable buying still exists after excluding news days.

The first row shows estimates for the full sample, while subsequent rows show the effects of excluding various types of news days. For the news article, analyst, guidance, and 8-K events, the table shows results excluding the news day and also from five days before to five days after the event. The extended ± 5 -day window excludes days that might be contaminated if the news leaks through private channels before or after the official announcement. The extended window is unnecessary for the return tests, because the return filter already excludes days with information leakage regardless the proximity to public news announcements.

In the Factiva exclusion tests, all coefficients on lagged HFT trading are positive and significant. Additionally, the sum of the first ten lag coefficients is not significantly different from the full sample.

In the IBES, 8-K, and return exclusion tests, the first lag coefficients are not significantly different from zero. However, coefficients on lags two and three and the sum of the first 10 are all

¹²See the [SEC website](#) for events that trigger 8-Ks. Data is from Phua, Tham, and Wei (2018).

¹³ Sample firms are matched to articles using IDs Factiva embeds in articles. See Tetlock, Saar-Tsechansky, and Macskassy (2008) and Gurun and Butler (2012) for evidence on the information content of Factiva articles.

positive and significant. The sum of the first 10 lag coefficients is smaller than the full sample for some of these tests. But the test with the biggest difference is still within 25% of the full-sample estimate.

The final test is the most restrictive. It excludes any day meeting any of the news day definitions. When only the actual news day is removed, coefficients on all three individual lags and the sum of the first ten are all positive and significant. The version using the ± 5 -day window is extremely restrictive, removing all but 835 of the 23,071 stock days in the full sample. With the extended window, the coefficients on the first three lags are positive, though only the second is significantly different from zero. Looking at the sum of the first ten coefficients provides more power in the limited sample. This sum is positive and significantly different from zero, and the difference with the full sample is not significantly different from zero.

These results show HFT net marketable buying still positively forecasts non-HFT net marketable buying on days without news.

4.2 Is the explanation non-HFTs and HFTs both chase price trends?

Another hypothesis for why HFT trading leads non-HFT trading is that non-HFT trading has a simple predictable relationship with past returns. Section 3.1 showed HFTs and non-HFTs both chase short-term price trends with their marketable trades. It is possible HFT trades predict non-HFT trades because they are both reacting to lagged returns, except HFTs react faster. Additionally, with non-HFTs following trend-chasing strategies, another possibility is purchases by HFTs could cause future non-HFT buying through their effect on returns.

Alternative Hypothesis 2 *If HFT net marketable buying leads non-HFT net marketable buying because non-HFTs chase past price trends, then there should be no correlation between HFT and non-HFT trading after controlling for lagged returns.*

This alternative hypothesis is addressed by the VAR in Table 3. The coefficients in regressions of non-HFT net marketable buying on HFT net marketable buying are positive controlling for lagged returns. HFT trades may cause some non-HFT trend chasing through their effect on returns, but it is not sufficient to explain the lead-lag correlation between HFT and non-HFT trading.

4.3 Are HFT trades forecasting misclassified HFT trades rather than true non-HFT trades?

Finally, it is possible HFT trading predicts non-HFT trading because some HFT trades are misclassified. Misclassification can occur if an HFT routes orders through brokers that handle non-HFT orders, such as Goldman Sachs and Morgan Stanley (Brogaard, Hendershott, and Riordan 2014). The problem with misclassification is that HFT trading is serially correlated. This serial correlation could be caused by some HFTs being slower or by trade splitting. Either way, if you take some of these later HFT trades and throw them into the non-HFT bucket, then current HFT trading will forecast those later misclassified trades. Consequently, misclassification could lead to an erroneous conclusion that HFT trades were forecasting non-HFT trades.

To see how HFT trades would forecast misclassified trades, you can look at how HFT trading forecasts itself. If correctly classified HFT trades are representative of misclassified trades, then the predictability of correctly classified trades should be proportional to that of misclassified trades. For example, a 100 share purchase by correctly classified HFTs might forecast that correctly classified HFTs purchase 20 shares 5 seconds from now and 10 shares 10 seconds from now. If misclassified trade predictability is proportional, then misclassified HFTs might purchase half those amounts 5 and 10 seconds from now. The essential idea is that the decay in predictability will be approximately the same. Importantly, the above hypothesis can be true even if correctly classified trades are not entirely representative of misclassified ones. It is only necessary that the later correctly classified HFT trades that generate the serial correlation are representative of later misclassified trades.

We can make additional predictions under the assumption that most HFT trades are correctly classified. This is a reasonable assumption given NASDAQ attributes roughly 40% of its dollar volume to the HFTs they identified. If correctly classified HFT trades account for the majority of HFT trades, then HFT trading should predict itself more than it predicts non-HFT trading. This is because misclassification implies non-HFT trading is a noisy HFT proxy, so lead-lag correlations should be attenuated relative to a sample with no noise (i.e., the correctly classified HFT trades).

Alternative Hypothesis 3a *Misclassification implies the decay in the correlation between non-HFT trading and lagged HFT trading is proportional to the decay in the correlation between HFT trading and lagged values of itself.*

Alternative Hypothesis 3b *Misclassification implies HFT trading should forecast itself more strongly than non-HFT trading.*

Figure 6 tests the misclassification hypotheses by comparing the VAR coefficients on HFT net marketable buying in Table 3. The dark red line plots coefficients from the regression where the dependent variable is HFT net marketable buying, while the light blue line is for the regression where the dependent variable is non-HFT net marketable buying. The VAR regression coefficients are standardized, so to facilitate comparison they are transformed to show the number of shares change in the dependent variable for a 100 share change in HFT net marketable buying.¹⁴ The y-axis is capped to zoom in on higher lags.

First, the figure shows the decay of the HFT to lagged HFT correlation is faster than the decay of the non-HFT to lagged HFT correlation. Using the VAR with 10 lags, Panel A shows coefficients for the first few lags are much higher when HFT net marketable buying is the dependent variable (dark red lines) than when non-HFT net marketable buying is the dependent variable (light blue lines). However, as the lags increase, the coefficients when HFT net marketable buying is the dependent variable drop steadily, while the coefficients when non-HFT net marketable buying is the dependent variable remain comparatively stable. This finding contradicts Alternative Hypotheses 3a.

Second, HFT net marketable buying predicts non-HFT net marketable buying better than itself at higher lags. This conflicts with Alternative Hypothesis 3b. Specifically, after seven lags, the HFT coefficient is actually bigger in the regression where the dependent variable is non-HFT net marketable buying. The coefficients in Panel A imply a 100 share increase in HFT net marketable buying ten seconds ago is associated with a 0.066 share increase in HFT net marketable buying, compared to a 0.244 share increase in non-HFT net marketable buying. Using the VAR with 30 lags, Panel B shows the coefficient in the HFT regression is roughly zero by the fifteenth lag, while the coefficient in the non-HFT regression remains positive and significantly different from zero even after 30 lags.

These findings are inconsistent with the misclassification explanation. The simpler explanation is that HFT marketable trades are mostly forecasting correctly classified non-HFT marketable

¹⁴ For example, the lag 1 HFT coefficient in the non-HFT regression in Table 3 is 0.0007. Table 1 shows the mean standard deviations of HFT_{NMB} (80) and $non-HFT_{NMB}$ (125). The coefficient in terms of a 100 share change in HFT net marketable buying is $0.0007 \times 125 \times 100 / 80 = 0.1$

trades, especially when looking 10 seconds or more into the future.

5 Relation to inventory management

Order anticipation typically comes up in the context of speculation. But being able to anticipate order flow is also useful for HFT market makers managing their inventory risk. Consider a market making HFT who acquires a short position in a stock by selling to liquidity demanding non-HFT buyers. If the HFT subsequently realizes these non-HFT purchases will continue and push prices up, then the HFT can avoid a loss by exiting the short position quickly with a marketable buy order. In this scenario, the HFT buys ahead of non-HFTs to get out of a losing position.

Inventory management and speculation have different implications for how HFTs' marketable buying is related to existing positions. Marketable trades that reverse existing positions are consistent with inventory management.¹⁵ Speculation is consistent with the other leg of a marketable purchase occurring beforehand or afterwards, though a common expectation is to see a non-marketable sale afterward.

The NASDAQ data does not contain position information. But HFT positions can be estimated by cumulating their net buying across all NASDAQ trades. A downside to this method is that relying on data from a single exchange means it misses position changes caused by trades on other venues (Menkveld 2013, Reiss and Werner 1998). In particular, if an HFT purchases a share of GE on NASDAQ and sells the share on the NYSE, they will appear to be long one share of GE in the NASDAQ data when they actually have no position. Nevertheless, the estimates are informative.

Figure 3 Panel A plots cumulative HFT net buying for the first and tenth net marketable buying decile portfolios. The portfolios are the same as in Figure 1, except the plotted variable is net buying instead of net marketable buying.

If HFTs' marketable purchases in the sort period are unwinding short inventory positions accumulated over the prior 30 seconds, then the solid line indicating cumulative net buying for decile ten (stocks HFTs buy aggressively in second 0) should be falling in the pre-sort period. However, net buying for decile 10 is slightly increasing in the 30 seconds before the sort. Similarly, inventory management implies the dotted line showing cumulative net buying for decile one (stocks HFTs sell

¹⁵For research on inventory management by NYSE specialists, see Hasbrouck and Sofianos (1993), Hendershott and Seasholes (2008), and Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010).

aggressively) should be rising in the pre-sort period. Instead, it is slightly declining. Looking at net buying after the sort, we also see no evidence of position reversal there. Therefore, in expectation the other leg of the sort-period trades occurs on NASDAQ more than 30 seconds before or after the sort or on another venue.

This test shows no evidence HFT trades in the sort period dispose of inventory positions previously acquired on NASDAQ. It is possible they are disposing of inventory acquired on another exchange. An HFT might sell to liquidity demanding non-HFTs on the NYSE and then, realizing non-HFT buying will continue and push the price up, turn around and buy the shares back on NASDAQ. However, this scenario still implies the HFT correctly predicts the non-HFT's order flow, trades ahead of it to increase their profits, and is taking liquidity the non-HFT would have otherwise accessed. It is still anticipatory trading.

That said, the evidence is more consistent with a speculative trading interpretation. Marketable trades are too large a share of HFT volume to be primarily inventory management. They account for over half of HFT dollar volume in the sample and, since they pay the bid-ask spread and take fee, market makers want to minimize their use (Harris 2002). Consequently, it is probable a lot of the variation in marketable trades is from non-market making strategies. Speculation is also more consistent with HFT trades predicting non-HFT trades more strongly early in the day (Table 7). Market makers like to end the day with minimal inventory positions. Therefore, if the results are caused by HFTs exiting inventory positions, the effect should be stronger (not weaker) at the end of the day. Stronger results in small-cap stocks (Table 2) is also less consistent with inventory management, because these are the stocks in which HFT market making is least prevalent (Yao and Ye 2018).

While an implied holding period longer than a few seconds contradicts some notions of high-frequency trading, it is perhaps important to distinguish between holding periods of market making and speculative strategies. To this point, the HFT Hudson River Trading told the *Wall Street Journal* it trades 5% of U.S. volume using both market making and speculative strategies, with 60% of their strategies taking liquidity (Hope 2014). Their average holding period was about five minutes and positions held overnight commonly accounted for 25% of firm capital. Thus, holding periods of more than a few seconds are not unusual. Hudson River explicitly states in the article they "... don't try to race ahead of an institution's order ..." Nonetheless, for some subset of HFTs,

the evidence in this paper shows those HFTs' marketable trades do just that.

6 Conclusion

This study examines whether HFTs anticipate and then trade ahead of other investors' order flow. I find that if HFTs on net buy a stock with marketable orders during a particular second, then non-HFTs will subsequently buy this same stock with marketable orders. This forecasted non-HFT net marketable buying continues for at least five minutes after the sort second. I explore several explanations for these findings, including them being driven by HFTs reacting faster to news, non-HFTs chasing return trends, and trader misclassification. However, the findings are best explained by HFTs learning which stocks non-HFTs will be buying or selling and then trading ahead of that order flow and its associated price impact.

Consistent with this anticipatory trading hypothesis, the effects are stronger at times when non-HFTs may be impatient, such as at the market open, on high volume days, and in stocks with wide bid-ask spreads. I also find evidence consistent with HFTs varying in their skill at predicting non-HFT order flow, and that trades from HFTs who are the most skilled at predicting order flow predict larger price changes. These findings provide evidence supporting the existence of an anticipatory trading channel through which HFTs increase non-HFT trading costs.

A primary reason for interest in research on HFTs is to understand their effect on other market participants. Likely benefits from HFTs acting as hyper-efficient market makers include lower bid-ask spreads and reduced return reversals. It is harder to get a handle on the costs HFTs impose on others. This study takes a step in that direction by providing evidence indicating HFTs trade ahead of other investors' order flow. Future research could provide more detail on which characteristics of past trades and orders allow HFTs to predict the stocks non-HFTs will buy and sell. Additionally, research that specifically focuses on HFT trading around news could determine whether HFTs are simply faster at reacting to news or if they are also better at interpreting it. Research into these questions would improve our understanding of the ways HFTs acquire information and, in doing so, inform evaluations of the welfare costs in addition to the benefits of HFT participation in the price discovery process.

References

- Ait-Sahalia, Yacine, and Mehmet Saglam, 2016, High frequency market making, *Available at SSRN 2331613*.
- Baldauf, Markus, and Joshua Mollner, 2020, High-frequency trading and market performance, *Journal of Finance* forthcoming.
- Barbon, Andrea, Marco Di Maggio, Francesco Franzoni, and Augustin Landier, 2019, Brokers and order flow leakage: Evidence from fire sales, Discussion Paper, 6.
- Baron, Matthew, Jonathan Brogaard, Björn Hagströmer, and Andrei Kirilenko, 2019, Risk and return in high-frequency trading, *Journal of Financial and Quantitative Analysis* 54, 993–1024.
- Battalio, Robert, Shane A Corwin, and Robert Jennings, 2016, Can brokers have it all? on the relation between make-take fees and limit order execution quality, *The Journal of Finance*.
- Battalio, Robert H., and Tim Loughran, 2008, Does payment for order flow to your broker help or hurt you?, *Journal of Business Ethics* 80, 37–44.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, *Journal of Financial Economics* 116, 292–313.
- Boehmer, Ekkehart, Kingsley Fong, and Julie Wu, 2015, International evidence on algorithmic trading, *Available at SSRN 2022034*.
- Boehmer, Ekkehart, Dan Li, and Gideon Saar, 2018, The competitive landscape of high-frequency trading firms, *The Review of Financial Studies* 31, 2227–2276.
- Breckenfelder, Johannes, 2013, Competition between high-frequency traders and market quality, *Available at SSRN 2264858*.
- Brogaard, Jonathan, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Sokolov, 2017, High-frequency trading and extreme price movements, *Journal of Financial Economics* forthcoming.
- Brogaard, Jonathan, and Corey Garriott, 2019, High-frequency trading competition, *Journal of Financial and Quantitative Analysis* 54, 1469–1497.
- Brogaard, Jonathan, Bjorn Hagstromer, Lars Norden, and Ryan Riordan, 2016, Trading fast and slow: Colocation and liquidity, *Review of Financial Studies* forthcoming.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High-frequency trading and price discovery, *Review of Financial Studies* 27, 2267–2306.
- , 2019, Price discovery without trading: Evidence from limit orders, *The Journal of Finance* 74, 1621–1658.
- Brunnermeier, Markus, and Lasse Pedersen, 2005, Predatory trading, *Journal of Finance* 60, 1825–1863.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *The Quarterly Journal of Economics* 130, 1547–1621.
- Chao, Yong, Chen Yao, and Mao Ye, 2016, What drives price dispersion and market fragmentation across u.s. stock exchanges?, *Working Paper*.
- Chordia, Tarun, T Clifton Green, and Badrinath Kottimukkalur, 2018, Rent seeking by low-latency

- traders: Evidence from trading on macroeconomic announcements, *The Review of Financial Studies* 31, 4650–4687.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2002, Order imbalance, liquidity, and market returns, *Journal of Financial Economics* 65, 111–130.
- Clark-Joseph, Adam, 2013, Exploratory trading, *Working Paper*.
- Comerton-Forde, Carole, Terrence Hendershott, Charles M. Jones, Pamela C. Moulton, and Mark S. Seasholes, 2010, Time variation in liquidity: The role of market-maker inventories and revenues, *The Journal of Finance* 65, 295–331.
- Comerton-Forde, Carole, and Tālis J Putniņš, 2015, Dark trading and price discovery, *Journal of Financial Economics* 118, 70–92.
- Conrad, Jennifer, and Sunil Wahal, 2019, The term structure of liquidity provision, Discussion paper, .
- , and Jin Xiang, 2015, High-frequency quoting, trading, and the efficiency of prices, *Journal of Financial Economics* 116, 271–291.
- Di Maggio, Marco, Francesco Franzoni, Amir Kermani, and Carlo Sommovilla, 2019, The relevance of broker networks for information diffusion in the stock market, Discussion Paper, 2.
- Egginton, Jared F, Bonnie F Van Ness, and Robert A Van Ness, 2016, Quote stuffing, *Financial Management* 45, 583–608.
- E*Trade, 2009, Etrade q3 2009 rule 606 filing, www.etrade.com.
- Fama, Eugene F, and James D MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of political economy* 81, 607–636.
- Foucault, Thierry, Johan Hombert, and Ioanid Rosu, 2016, News trading and speed, *The Journal of Finance* 71, 335–382.
- Foucault, Thierry, and Albert J. Menkveld, 2008, Competition for order flow and smart order routing systems, *Journal of Finance* 63, 119–158.
- Gervais, Simon, Ron Kaniel, and Dan H. Mingelgrin, 2001, The high-volume return premium, *Journal of Finance* 56, 877–919.
- GETCO, 2010, Our firm: Getco, http://www.getcollc.com/index.php/getco/tertiary/our_firm/ Accessed: 6/2/2010.
- Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71–100.
- Griffith, Todd G, Bonnie F Van Ness, and Robert A Van Ness, 2017, High-frequency trading patterns around short-term volatility spikes, *The Journal of Trading* 12, 48–68.
- Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *The American Economic Review* pp. 393–408.
- Gurun, Umit G, and Alexander W Butler, 2012, Don't believe the hype: Local media slant, local advertising, and firm value, *The Journal of Finance* 67, 561–598.
- Harris, Larry, 2002, *Trading and exchanges: Market microstructure for practitioners* (Oxford University Press, USA).

- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *The Journal of Finance* 46, 179–207.
- , 2010, The best bid and offer: A short note on programs and practices, *Working Paper*.
- , 2018, High-frequency quoting: Short-term volatility in bids and offers, *Journal of Financial and Quantitative Analysis* 53, 613–641.
- , and Gideon Saar, 2013, Low-latency trading, *Journal of Financial Markets* 16, 646–679.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical analysis of nyse specialists, *Journal of Finance* 48, 1565–1593.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2011, Does algorithmic trading improve liquidity?, *Journal of Finance* 66, 1–33.
- Hendershott, Terrence, and Ryan Riordan, 2013, Algorithmic trading and the market for liquidity, *Journal of Financial and Quantitative Analysis* pp. 1001–1024.
- Hendershott, Terrence, and Mark S. Seasholes, 2008, Market maker inventories and stock prices, *American Economic Review* 97, 210–214.
- Hope, Bradley, 2014, Inside hudson river, <https://blogs.wsj.com/moneybeat/2014/10/15/inside-hudson-river-the-firm-that-does-5-of-all-stock-trading>.
- ITG, 2013, Dark liquidity aggregation state of play, <http://www.itg.com/2012/11/19/dark-liquidity-aggregation-state-of-play/print/> Accessed: 2/4/2013.
- Jovanovic, Boyan, and Albert J. Menkveld, 2015, Middlemen in limit order markets, *Working paper*.
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun, 2017, The flash crash: high frequency trading in an electronic market, *Journal of Finance*.
- Korajczyk, Robert A, and Dermot Murphy, 2019, High-frequency market making to large institutional trades, *The Review of Financial Studies* 32, 1034–1067.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1385.
- Lewis, Michael, 2014, *Flash boys: a wall street revolt* (WW Norton & Company).
- Madrigal, Vicente, 1996, Non-fundamental speculation, *The Journal of Finance* 51, 553–578.
- Menkveld, Albert J., 2013, High frequency trading and the new-market makers, *Journal of Financial Markets*.
- Menkveld, Albert J, and Marius A Zoican, 2017, Need for speed? exchange latency and liquidity, *The Review of Financial Studies* 30, 1188–1228.
- NASDAQ 10K, 2009, Our firm: Getco, http://www.getcollc.com/index.php/getco/tertiary/our_firm/ Accessed: 2/22/2017.
- Newey, Whitney K., and Kenneth D. West, 1994, Automatic lag selection in covariance matrix estimation, *Review of Economic Studies* 61, 631–653.
- O’Hara, Maureen, Gideon Saar, and Zhuo Zhong, 2019, Relative tick size and the trading environment, *The Review of Asset Pricing Studies* 9, 47–90.
- O’Hara, Maureen, Chen Yao, and Mao Ye, 2014, What’s not there: Odd lots and market data, *Journal of Finance* 69, 2199–2236.
- O’Hara, Maureen, and Mao Ye, 2011, Is market fragmentation harming market quality?, *Journal*

- of *Financial Economics* 100, 459–474.
- Phua, J Kenny, T Mandy Tham, and Chishen Wei, 2018, Do equity analysts learn from their colleagues?, *Available at SSRN 3300721*.
- Reiss, Peter C., and Ingrid M. Werner, 1998, Does risk sharing motivate interdealer trading?, *The Journal of Finance* 53, 1657–1703.
- SEC, 2010, <https://www.sec.gov/rules/sro/finra/2010/34-62885.pdf> Accessed: 3/20/2017.
- Stiglitz, Joseph E, 2014, Tapping the brakes: Are less active markets safer and better for the economy?, in *Federal Reserve Bank of Atlanta 2014 Financial Markets Conference*.
- Stoll, Hans R., 1978, The supply of dealer services in securities markets, *The Journal of Finance* 33, 1133–1151.
- Tabb, Larry, 2009, High frequency trading: What is it & should I be worried?, <http://www.world-exchanges.org/files/file/LarryTABB-HFT-part1.ppt> Accessed: 8/9/2012.
- Tetlock, Paul C., Maytal Saar-Tsechansky, and Sofus Macskassy, 2008, More than words: Quantifying language to measure firms’ fundamentals, *Journal of Finance* 63, 1437–1467.
- Tradebot, 2010, About us, <http://www.tradebotsystems.com> Accessed: 6/2/2010.
- Traders Magazine, 2013, New goldman sachs algo takes aim at small caps, <http://www.tradersmagazine.com/news/goldman-sachs-algorithm-small-109890-1.html> Accessed: 2/5/2013.
- Van Kervel, Vincent, and Albert J Menkveld, 2019, High-frequency trading around large institutional orders, *The Journal of Finance* 74, 1091–1137.
- Vayanos, Dimitri, 2001, Strategic trading in a dynamic noisy market, *The Journal of Finance* 56, 131–171.
- Von Beschwitz, Bastian, Donald B Keim, and Massimo Massa, 2015, First to ‘read’ the news: News analytics and institutional trading, .
- Weller, Brian M, 2018, Does algorithmic trading reduce information acquisition?, *The Review of Financial Studies* 31, 2184–2226.
- Yang, Liyan, and Haoxiang Zhu, 2019, Back-running: Seeking and hiding fundamental information in order flows, *Review of Financial Studies*.
- Yao, Chen, and Mao Ye, 2018, Why trading speed matters: A tale of queue rationing under price controls, *The Review of Financial Studies* 31, 2157–2183.
- Ye, Mao, Chen Yao, and Jiading Gai, 2013, The externalities of high-frequency trading, *Available at SSRN 2066839*.

Figure 1: Cumulative HFT vs. non-HFT Net Marketable Buying.

This figure plots cumulative standardized net marketable buying for stocks sorted into portfolios by HFTs' net marketable buying at the one-second horizon. The focus is whether the sort forecasts changes in cumulative non-HFT imbalances after the sort period. The time t cumulative imbalance is the time $t-1$ cumulative imbalance plus the time t imbalance. The left y-axis is for HFT net marketable buying, HFT_{NMB} , and the right y-axis is for non-HFT net marketable buying, $non-HFT_{NMB}$. Stocks are sorted into deciles each second based on HFT net marketable buying. Decile breakpoints are calculated from non-zero observations during the prior trading day. Stocks in decile ten are marked as those HFTs are buying and those in decile one are marked as those HFTs are selling. Portfolio means are calculated by first averaging across all observations for a day and then averaging across all days in the sample. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Both variables in the figure are divided by their respective standard deviation among all sample stocks that day.

Figure 2: Bid-Ask Midpoint Returns

This figure plots NBBO midpoint returns in basis points for stocks sorted into portfolios by HFTs' net marketable buying. Stocks are sorted into deciles each second based on HFT net marketable buying. Decile breakpoints are calculated from non-zero observations during the prior trading day. Stocks in decile ten are marked as those HFTs are buying and those in decile one are marked as those HFTs are selling. Portfolio means are calculated by first averaging across all observations for a day and then averaging across all days in the sample. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks.

Figure 3: HFT Net Buying around the HFT Net Marketable Buying Sorts

Cumulative standardized HFT net buying (HFT_{NB}) around sorts on HFT net marketable buying (HFT_{NMB}). The time t cumulative imbalance is the time $t-1$ cumulative imbalance plus the time t imbalance. Panel A plots aggregate HFT_{NB} for stocks sorted on aggregate HFT_{NMB} . Panel B plots HFT_{NB} of three HFT groups (*High*, *Mid*, and *Low*) for stocks sorted by the *High* group's HFT_{NMB} . HFTs are sorted into groups by the coefficient on the first lag of the individual HFT's HFT_{NMB} in a regression of *non- HFT_{NMB}* on ten lags of the HFT's HFT_{NMB} , *non- HFT_{NMB}* , and returns estimated during the prior month. Table 5 provides more detail on the sorting procedure. Stocks are sorted into deciles each second based on HFT net marketable buying. Decile breakpoints are calculated from non-zero observations during the prior trading day. Stocks in decile ten are marked as those HFTs are buying and those in decile one are marked as those HFTs are selling. Portfolio means are calculated by first averaging across all observations for a day and then averaging across all days in the sample. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Net buying is shares bought minus shares sold. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. In both panels, all HFT_{NB} measures are divided by the standard deviation of aggregate HFT_{NB} among all sample stocks that day to ease interpretation.

Figure 4: Non-HFT Price Impact with and without competition from HFTs

Impulse response functions for returns estimate the price impact of a 100 share shock to non-HFT net marketable buying ($non-HFT_{NMB}$) following Hasbrouck (1991). The first estimate uses the full VAR. The second is an estimate for a counterfactual world with no HFT competition. It sets HFT_{NMB} in every second to zero and measures price impact if $non-HFT_{NMB}$ each second after the shock is the same as in the first scenario. The solid line is the estimated response, and dashed lines indicate 95% confidence intervals. Estimates come from the VAR:

$$\begin{aligned}
 R_t &= \alpha_1 + \sum_{i=0}^{10} \gamma_{1,i} HFT_{NMB,t-i} + \sum_{i=0}^{10} \beta_{1,i} non-HFT_{NMB,t-i} + \sum_{i=0}^{10} \omega_{1,i} C_{t-i} + \sum_{i=1}^{10} \lambda_{1,i} R_{t-i} + \epsilon_{1,t} \\
 HFT_{NMB,t} &= \alpha_2 + \sum_{i=1}^{10} \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=0}^{10} \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \omega_{2,i} C_{t-i} + \sum_{i=1}^{10} \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \\
 non-HFT_{NMB,t} &= \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \omega_{3,i} C_{t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t} \\
 C_t &= \begin{cases} HFT_{NMB,t} \times non-HFT_{NMB,t}, & \text{if } HFT_{NMB,t} > 0 \text{ and } non-HFT_{NMB,t} > 0 \\ HFT_{NMB,t} \times non-HFT_{NMB,t} \times -1, & \text{if } HFT_{NMB,t} < 0 \text{ and } non-HFT_{NMB,t} < 0 \\ 0, & \text{otherwise} \end{cases}
 \end{aligned}$$

where R_t is NBBO midpoint return. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method. Following estimation, the standardized share shocks and return response are converted to share and basis point units for ease of interpretation.

Figure 5: Persistence of Differences in Individual HFTs' Prediction Ability

Given an HFT is among the top third of HFTs in terms of how their trades predict non-HFT order flow, the figure plots the probability they will still be in the top third one, two, and three months later. The dotted line at 33.3% is what would be expected in months 1–3 if there were no persistence in which HFTs' trades are most strongly correlated with future non-HFT trades. The four lines indicate different sorting methods using the 10-lag regression discussed in Table 5. HFT_{NB} is HFT net buying. HFT_{NMB} is the HFT's net marketable buying. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Net buying is shares bought minus shares sold. The probabilities are calculated for each sort month and then averaged across all sort months in the sample. HFTs that leave the sample after the sort period are assigned to the lowest correlation group.

Figure 6: Does HFT net marketable buying predict itself or non-HFT net marketable buying better?

Comparison of HFT coefficients ($\gamma_{2,i}$ and $\gamma_{3,i}$ below) from regressions where HFT and non-HFT trading are the dependent variables. The solid line depicts the estimate and dotted lines show 95% confidence intervals. To make the comparison in the same units, the standardized coefficients are converted to show the share change in the dependent variable for a 100 share change in the HFT variable. The conversion uses the mean standard deviation of the variables in Table 1. Y-axis is cut off above 0.5 to zoom in on higher lags. The regressions are equations 5 and 6 from Table 3:

$$\begin{aligned}
 HFT_{NMB,t} &= \alpha_2 + \sum_{i=1}^J \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=1}^J \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^J \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \\
 non-HFT_{NMB,t} &= \alpha_3 + \sum_{i=1}^J \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^J \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^J \lambda_{3,i} R_{t-i} + \epsilon_{3,t}
 \end{aligned}$$

where R_t is NBBO midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, $non-HFT_{NMB,t}$ is non-HFT net marketable buying, and J is the number of lags. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method.

Table 1
Summary Statistics

The sample consists of 96 stocks chosen by randomly selecting six NASDAQ-listed and six NYSE-listed stocks from each of the top eight NYSE size deciles of CRSP common stocks passing liquidity filters described in Section 2. Size portfolio breakpoints are computed among NYSE-listed stocks. Size portfolios for year t are formed on December 31st of year $t-1$. Deciles one through five are small-cap, six through eight are mid-cap, and nine through ten are large-cap. Panel A summarizes stock characteristics calculated from the pooled time-series of all stock-day observations. Market capitalization, price, and dollar volume are end of day values. NQ is the share of total dollar volume that traded on NASDAQ, HFT is the fraction of NASDAQ dollar volume traded by HFTs, and N is stocks per day in the sample. Panel B summarizes intra-day returns and net buying measures by reporting the mean and median daily standard deviation of these variables across all stock days. Each day, for every stock, the following are calculated: the standard deviation of NBBO midpoint returns ($BAM\ Ret$), HFTs' net buying (HFT_{NB}), HFTs' net marketable buying (HFT_{NMB}), and non-HFTs' net marketable buying ($non-HFT_{NMB}$). Net buying is shares bought minus shares sold. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. For this table only, imbalance measures are expressed in shares.

Panel A: Daily Stock Characteristics

		Mkt Cap Mil. \$	Price \$	Volume Mil. \$	NQ %	HFT %	N
All Stocks	mean	5,302	26.38	58.1	27.2	27.6	93.2
	median	1,301	22.05	12.6	25.0	27.7	93.0
	std dev	12,909	19.96	111.4	13.5	13.7	1.6
	min	22	0.91	0.1	0.7	0.0	89.0
	max	125,331	166.82	2,153.1	80.7	78.4	96.0
Small-cap	mean	367	16.52	4.0	26.3	16.7	33.4
	median	293	14.77	1.9	22.5	14.8	33.0
Mid-cap	mean	1,900	26.10	34.2	27.1	28.9	35.8
	median	1,565	25.04	15.5	26.0	29.2	36.0
Large-cap	mean	17,252	40.55	169.3	28.5	40.7	24.0
	median	9,413	31.37	120.2	26.9	40.9	24.0

Panel B: Standard Deviation of Intra-day Variables

		$BAM\ Ret$ %	HFT_{NB} shares	HFT_{NMB} shares	$non-HFT_{NMB}$ shares
All Stocks	mean	0.080	83	80	125
	median	0.027	28	26	42
Small-cap	mean	0.052	13	11	27
	median	0.031	9	7	18
Mid-cap	mean	0.085	89	81	138
	median	0.022	35	34	49
Large-cap	mean	0.114	175	178	247
	median	0.023	116	123	138

Table 2
Stocks Sorted by HFT Net Marketable Buying

This table shows standardized non-HFT net marketable buying and NBBO midpoint returns for stocks sorted on HFTs' net marketable buying at the one-second horizon. Stocks are sorted into deciles at time t based on HFT net-marketable buying. Decile breakpoints are calculated from non-zero observations during the prior trading day. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Non-HFT net marketable buying and returns for portfolios are averaged across all observations for a day, and the mean of the daily time series is reported in the table. Parentheses indicate t -statistics for the time-series means, with autocorrelation corrections following the optimal automatic lag selection method of Newey and West (1994). The reported imbalance is divided by its standard deviation among all stocks that day to ease interpretation.

Panel A: Non-HFT Net Marketable Buying

Decile	Seconds					
	$[t-30, t-1]$	$t-1$	t	$t+1$	$[t+1, t+30]$	$[t+1, t+300]$
All Stocks						
10 (HFT Buying)	0.26 (4.39)	0.10 (22.01)	0.58 (24.92)	0.10 (26.17)	0.65 (14.35)	1.15 (2.41)
9	0.12 (10.35)	0.04 (25.24)	0.22 (24.86)	0.03 (21.91)	0.26 (23.47)	0.50 (7.28)
2	-0.11 (-10.69)	-0.04 (-24.92)	-0.22 (-26.79)	-0.03 (-29.18)	-0.26 (-25.85)	-0.62 (-8.59)
1 (HFT Selling)	-0.29 (-4.52)	-0.10 (-20.24)	-0.58 (-32.28)	-0.09 (-19.66)	-0.69 (-10.15)	-1.89 (-4.26)
Small-cap						
10 (HFT Buying)	1.09 (5.49)	0.25 (18.14)	1.09 (20.60)	0.21 (13.91)	1.39 (9.68)	2.42 (1.54)
1 (HFT Selling)	-1.04 (-4.68)	-0.26 (-11.71)	-1.07 (-27.77)	-0.19 (-10.16)	-1.27 (-5.40)	-3.27 (-2.27)

Panel B: NBBO midpoint returns in basis points

Decile	Seconds					
	$[t-30, t-1]$	$t-1$	t	$t+1$	$[t+1, t+30]$	$[t+1, t+300]$
All Stocks						
10 (HFT Buying)	4.43 (25.26)	4.51 (25.49)	0.92 (15.91)	0.55 (17.98)	1.23 (11.78)	0.57 (1.16)
9	3.24 (29.04)	3.20 (24.28)	0.65 (9.41)	0.47 (14.53)	0.70 (14.03)	0.55 (3.63)
2	-3.04 (-30.62)	-3.11 (-27.98)	-0.51 (-8.95)	-0.34 (-11.07)	-0.51 (-7.42)	0.05 (0.30)
1 (HFT Selling)	-4.30 (-26.30)	-4.47 (-27.66)	-0.81 (-14.28)	-0.46 (-14.58)	-1.07 (-13.48)	-0.58 (-1.42)
Small-cap						
10 (HFT Buying)	7.98 (19.77)	6.63 (19.64)	1.31 (13.75)	0.61 (11.48)	2.54 (8.41)	1.36 (0.86)
1 (HFT Selling)	-7.59 (-15.25)	-6.55 (-21.68)	-1.19 (-15.07)	-0.57 (-13.08)	-2.64 (-14.37)	-3.09 (-2.58)

Table 3

Intra-day VAR Estimates for Individual Stock-day Observations

For each stock-day observation, the following vector autoregressions (VARs) with ten lags are estimated:

$$R_t = \alpha_1 + \sum_{i=1}^{10} \gamma_{1,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{1,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{1,i} R_{t-i} + \epsilon_{1,t} \quad (4)$$

$$HFT_{NMB,t} = \alpha_2 + \sum_{i=1}^{10} \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \quad (5)$$

$$non-HFT_{NMB,t} = \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t} \quad (6)$$

where R_t is the NBBO midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, and $non-HFT_{NMB,t}$ is non-HFT net marketable buying. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method.

lag	γ (HFT)		β (non-HFT)		λ (R)	
	μ	t -stat	μ	t -stat	μ	t -stat
$y = non-HFT_t$						
1	0.0007	2.21	0.0754	37.32	0.9181	13.68
2	0.0021	15.47	0.0246	39.42	-0.0145	-4.28
3	0.0020	13.12	0.0184	38.72	-0.0008	-0.53
4	0.0018	12.62	0.0153	33.37	-0.0028	-1.95
5	0.0015	13.52	0.0145	38.66	-0.0032	-2.41
6	0.0017	14.00	0.0111	25.98	0.0047	4.68
7	0.0015	13.97	0.0093	33.33	0.0016	1.54
8	0.0016	15.41	0.0084	28.33	0.0005	0.73
9	0.0016	13.59	0.0089	30.21	-0.0002	-0.22
10	0.0016	13.48	0.0117	30.19	-0.0025	-2.59
$y = HFT_t$						
1	0.0243	25.55	0.0028	2.64	1.9956	10.19
2	0.0088	24.58	0.0007	1.49	0.0290	4.79
3	0.0056	19.07	-0.0001	-0.21	-0.0047	-2.16
4	0.0039	17.41	-0.0008	-2.32	-0.0124	-6.52
5	0.0042	15.42	-0.0013	-3.34	-0.0131	-7.84
6	0.0026	10.34	-0.0011	-3.11	-0.0125	-7.26
7	0.0022	9.79	-0.0014	-3.98	-0.0117	-7.40
8	0.0014	7.30	-0.0007	-2.35	-0.0127	-9.49
9	0.0015	7.26	-0.0012	-3.34	-0.0148	-8.26
10	0.0007	3.28	-0.0019	-5.56	-0.0121	-7.83
$y = R_t$						
1	0.0185	10.92	0.0269	14.82	-0.1551	-55.30
2	0.0115	11.82	0.0155	12.64	-0.1310	-63.87
3	0.0085	11.93	0.0119	9.85	-0.1070	-68.61
4	0.0067	9.60	0.0102	11.30	-0.0906	-66.86
5	0.0056	9.94	0.0078	10.86	-0.0749	-71.81
6	0.0042	10.60	0.0066	11.69	-0.0645	-72.24
7	0.0035	7.67	0.0037	6.65	-0.0522	-70.93
8	0.0023	10.21	0.0041	10.01	-0.0420	-77.02
9	0.0017	5.68	0.0018	3.24	-0.0308	-77.50
10	-0.0001	-0.38	-0.0003	-0.73	-0.0188	-67.19

Table 4

Example Patterns Predicting non-HFT Net Marketable Buying and Returns

This table reports two patterns that predict non-HFT net marketable buying ($non-HFT_{NMB}$), HFT net marketable buying (HFT_{NMB}), and NBBO midpoint returns in basis points (R). Stocks are sorted in second t into *Buy* and *Sell* portfolios based on values for second $t-1$ and accumulated values from seconds $t-5$ to $t-1$:

	Pattern 1		Pattern 2	
	Buy	Sell	Buy	Sell
$\sum_{l=1}^5 non-HFT_{NMB,t-l}$	> 2	< 2	$> non-HFT_{NMB,t-1}$	$< non-HFT_{NMB,t-1}$
$non-HFT_{NMB,t-1}$	> 0	< 0	> 1	< -1
$\prod_{l=1}^5 (1 + R_{t-l}) - 1$	≥ 0	≤ 0	≥ 0	≤ 0
R_{t-1}	$= 0$	$= 0$	$= 0$	$= 0$

Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Portfolio values are averaged across all observations for a day, and the mean of the daily time series is reported in the table. Parentheses indicate t -statistics for the time-series means, with autocorrelation corrections following the optimal automatic lag selection method of Newey and West (1994).

		Seconds						
		$[t-30, t-1]$	$[t-5, t-1]$	$t-1$	t	$t+1$	$[t+1, t+30]$	$[t+1, t+300]$
Panel A: Pattern 1								
$non-HFT_{NMB}$	Buy	9.10 (85.75)	7.00 (99.22)	5.04 (97.38)	0.25 (23.65)	0.27 (35.90)	2.69 (36.66)	7.43 (10.67)
	Sell	-9.08 (-72.84)	-6.97 (-96.20)	-5.03 (-100.03)	-0.23 (-29.66)	-0.26 (-22.62)	-2.74 (-21.34)	-8.16 (-13.29)
	Diff	18.19 (99.98)	13.97 (106.00)	10.07 (103.37)	0.48 (29.76)	0.53 (32.48)	5.43 (34.46)	15.60 (33.61)
<i>Return</i> (bps)	Buy	7.01 (24.04)	5.89 (26.56)	0.00 -	0.36 (17.44)	0.14 (5.95)	1.68 (13.80)	2.96 (7.21)
	Sell	-7.11 (-23.32)	-5.88 (-28.47)	0.00 -	-0.26 (-8.64)	-0.15 (-8.61)	-1.64 (-14.51)	-2.61 (-8.58)
	Diff	14.12 (23.95)	11.77 (27.76)	0.00 -	0.61 (16.62)	0.29 (9.32)	3.32 (15.71)	5.58 (14.79)
HFT_{NMB}	Buy	2.73 (20.43)	1.59 (24.51)	0.73 (36.62)	0.08 (7.05)	0.08 (13.61)	0.47 (6.04)	0.96 (3.27)
	Sell	-2.48 (-38.73)	-1.51 (-36.83)	-0.70 (-38.43)	-0.07 (-16.49)	-0.07 (-17.42)	-0.34 (-9.11)	-0.35 (-1.45)
	Diff	5.21 (35.97)	3.10 (31.65)	1.43 (39.86)	0.16 (10.34)	0.15 (20.62)	0.81 (8.89)	1.31 (6.94)
Panel B: Pattern 2								
$non-HFT_{NMB}$	Buy	12.04 (55.99)	8.61 (61.08)	3.83 (86.23)	0.45 (24.09)	0.50 (29.94)	4.11 (32.72)	10.05 (9.65)
	Sell	-12.11 (-45.01)	-8.60 (-55.81)	-3.84 (-75.04)	-0.44 (-21.21)	-0.49 (-18.35)	-4.33 (-16.77)	-11.44 (-10.43)
	Diff	24.15 (59.62)	17.21 (61.58)	7.67 (86.93)	0.89 (24.81)	0.99 (25.30)	8.44 (27.72)	21.50 (28.14)
<i>Return</i> (bps)	Buy	8.73 (25.81)	6.65 (28.85)	0.00 -	0.42 (12.23)	0.17 (4.22)	1.72 (11.48)	2.26 (3.75)
	Sell	-8.88 (-27.70)	-6.71 (-31.83)	0.00 -	-0.26 (-6.15)	-0.20 (-6.42)	-1.63 (-13.31)	-2.02 (-4.07)
	Diff	17.61 (26.70)	13.37 (30.80)	0.00 -	0.68 (13.19)	0.37 (7.06)	3.35 (15.54)	4.28 (12.55)
HFT_{NMB}	Buy	4.20 (20.31)	2.32 (21.55)	0.72 (24.32)	0.12 (6.48)	0.11 (12.72)	0.57 (4.54)	1.08 (2.68)
	Sell	-3.72 (-29.84)	-2.20 (-27.07)	-0.69 (-30.61)	-0.11 (-10.18)	-0.09 (-8.85)	-0.27 (-4.62)	-0.17 (-0.48)
	Diff	7.92 (32.49)	4.53 (25.89)	1.41 (28.53)	0.23 (8.54)	0.20 (14.10)	0.85 (5.80)	1.25 (3.50)

Table 5
Examining Differences in Individual HFTs' Prediction Ability

This table tests whether trades from some HFTs consistently predict future buying and selling pressure better than others. The test examines whether the HFTs who predict buying and selling pressure the best one month continue to do so the next month. Each day, the following regression is run for each HFT:

$$non-HFT_{d,s,t} = \alpha_{d,i} + \sum_{l=1}^J \gamma_{d,i,l} HFT_{d,s,i,t-l} + \sum_{l=1}^J \beta_{d,i,l} non-HFT_{d,s,t-l} + \sum_{i=1}^J \lambda_{d,i,l} R_{d,s,t-l} + \epsilon_{d,i,s,t}, \quad (7)$$

where d indexes days, s indexes stocks, t indexes seconds, J is either 10 or 30 lags, and i indexes HFTs. $non-HFT$ is non-HFT net marketable buying, HFT is either the HFT's net buying (HFT_{NB}) or net marketable buying (HFT_{NMB}), and R is the stock's NBBO midpoint return. Net buying is shares bought minus shares sold. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. The individual HFTs' net marketable buying and net buying measures are divided by the standard deviation of aggregate HFT net marketable buying and net buying, respectively. The left set of columns in the table below report results from regressions where HFT is HFTs' net marketable buying, and the right set report results where HFT is their net buying. For regression (7), I require there to be more than 100 non-zero net marketable buying observations to ensure relatively precise coefficient estimates. Then for each month, among HFTs for whom regression (7) could be estimated at least 15 days during the current and following month, HFTs are split into three groups based on their γ_i coefficients. There are three groupings: 1) HFTs are split based on their $\overline{\gamma_{d,i,1}}$ for that month in a 10-lag regression, 2), the split is based on their $\overline{\sum_{l=1}^{10} \gamma_{d,i,l}}$ in the 10-lag regression, and 3) the split is based on their $\overline{\sum_{l=1}^{30} \gamma_{d,i,l}}$ in the 30-lag regression. The cross-sectional average $\overline{\gamma_{d,i,1}}$, $\overline{\sum_{l=1}^{10} \gamma_{d,i,l}}$, and $\overline{\sum_{l=1}^{30} \gamma_{d,i,l}}$ coefficients are then calculated the following month (i.e., the post-sort month). For each group's monthly time series, the table reports the time-series mean, p-value for the test that the mean is different from zero using standard errors corrected for autocorrelation following the Newey and West (1994) optimal lag selection method, time-series median, number of months out of 11 that the value is greater than zero, and p-value from a non-parametric Wilcoxon rank sum test that the time-series median is equal to zero. I report p-values for the t -test because the small number of months means standard rules of thumb (e.g., t -stat > 1.96) for determining statistical significance do not apply. High-frequency traders go in and out of the sample, so months are weighted by the number of HFTs in that month's group.

	<i>HFT</i> = Net Mkt. Buying					<i>HFT</i> = Net Buying				
	μ_{t+1}	p-value	median	#>0 of 11	Wilcoxon p-value	μ_{t+1}	p-value	median	#>0 of 11	Wilcoxon p-value
$\overline{\gamma_{d,i,1}}$										
H in month _{<i>t</i>}	0.024	< 0.01	0.024	11	< 0.01	0.031	< 0.01	0.029	11	< 0.01
M in month _{<i>t</i>}	0.008	< 0.01	0.008	11	< 0.01	0.010	< 0.01	0.010	11	< 0.01
L in month _{<i>t</i>}	-0.005	0.12	-0.001	4	0.24	-0.003	0.01	-0.002	2	0.04
High-Low	0.029	< 0.01	0.028	11	< 0.01	0.034	< 0.01	0.032	11	< 0.01
$\overline{\sum_{l=1}^{10} \gamma_{d,i,l}}$										
H in month _{<i>t</i>}	0.095	< 0.01	0.095	11	< 0.01	0.107	< 0.01	0.089	11	< 0.01
M in month _{<i>t</i>}	0.033	< 0.01	0.030	11	< 0.01	0.033	< 0.01	0.033	11	< 0.01
L in month _{<i>t</i>}	-0.006	0.66	0.010	7	0.90	-0.001	0.80	0.000	6	0.90
High-Low	0.102	< 0.01	0.094	11	< 0.01	0.108	< 0.01	0.089	11	< 0.01
$\overline{\sum_{l=1}^{30} \gamma_{d,i,l}}$										
H in month _{<i>t</i>} 1	0.136	0.04	0.134	10	< 0.01	0.124	< 0.01	0.118	11	< 0.01
M in month _{<i>t</i>} 1	0.053	< 0.01	0.042	11	< 0.01	0.048	< 0.01	0.044	10	< 0.01
L in month _{<i>t</i>} 1	0.023	0.02	0.028	9	0.05	0.014	0.11	0.010	10	0.02
High-Low1	0.114	0.02	0.129	9	< 0.01	0.110	< 0.01	0.114	11	< 0.01

Table 6

Differences among HFTs in How Strongly Their Marketable Trades Forecast Returns

This table examines whether trades from the HFTs whose trades are most strongly correlated with future non-HFT order flow also predict larger future NBBO midpoint returns. HFTs are split into three groups each month based on the correlation between their marketable trades and future non-HFT order flow. The split is based on the average $\gamma_{d,i,1}$ coefficient from Regression 7 in Table 5 estimated during the prior month. Table 5 provides more detail on the methodology for sorting HFTs. Once HFT groups are determined, all their trades are then aggregated among HFTs in the same group during the post-sort month. The table below compares coefficients on the top and bottom HFT groups' aggregate net marketable buying in regressions of the following form:

$$R_t = \alpha + \sum_{i=1}^{10} \gamma_i HFT_{NMB,t-i}^G + \sum_{i=1}^{10} \beta_i non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_i R_{t-i} + \epsilon_{t,R} \quad (8)$$

where R_t is NBBO midpoint return, $HFT_{NMB,t}^G$ is HFT net marketable buying for either group, and $non-HFT_{NMB,t}$ is non-HFT net marketable buying. All variables are measured at the one-second horizon. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Returns and non-HFT net marketable buying are divided by their one-second standard deviations, whereas the HFT groups' net marketable buying is divided by the one-second standard deviation of HFT_{NMB} aggregated among all HFTs. The values reported are the time-series mean and median of daily cross-sectional mean coefficients. For t -tests, the null hypothesis is that the mean of the daily time series equals zero. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method. The difference column group also reports the p-value from a non-parametric Wilcoxon rank sum test that the time-series medians are equal.

lag	High in month _{t-1} HFTs' γ_i			Low in month _{t-1} γ_i			Difference		
	mean	t -stat	median	mean	t -stat	median	mean	t -stat	rank sum p-value
1	0.0328	10.44	0.0259	0.0268	10.69	0.0235	0.0060	2.41	0.02
2	0.0210	10.80	0.0169	0.0160	11.53	0.0146	0.0051	3.39	0.01
3	0.0160	10.62	0.0133	0.0125	10.10	0.0111	0.0035	3.00	0.02
4	0.0118	11.16	0.0101	0.0086	8.73	0.0086	0.0033	3.43	0.02
5	0.0098	9.77	0.0082	0.0087	9.31	0.0079	0.0012	1.06	0.18
6	0.0080	8.16	0.0060	0.0064	5.93	0.0057	0.0016	1.68	0.28
7	0.0061	6.20	0.0046	0.0049	4.79	0.0046	0.0013	1.25	0.75
8	0.0038	6.43	0.0029	0.0023	2.37	0.0028	0.0015	1.37	0.68
9	0.0028	3.40	0.0016	0.0030	2.94	0.0032	-0.0002	-0.16	0.20
10	-0.0004	-0.57	-0.0006	-0.0004	-0.43	0.0009	0.0000	0.03	0.41
Σ 1-10	0.1119	10.69	0.0945	0.0887	10.59	0.0794	0.0231	3.78	0.00

Table 7

Conditioning on Times when non-HFTs are Hypothesized to be Impatient

This table reports coefficients on HFTs' net marketable buying from the VAR in Table 3 conditional on times when non-HFTs are impatient. Estimates are from the equation,

$$non-HFT_{NMB,t} = \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t},$$

where R_t is NBBO midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, and $non-HFT_{NMB,t}$ is non-HFT net marketable buying. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method. In Panel B, high volume and high absolute daily aggregate net marketable buying imbalance days are those among the top 10% of days the past month (i.e., rank relative to prior 19 trading days is 19 or 20). The daily average spread and relative spread are calculated by duration-weighting intra-day observations. For either spread measure, the high/low spread groups consist of stocks in the top/bottom third of the sample the prior day.

Panel A: Comparing γ estimates from the open and close to the middle of the day

lag							Difference with Mid-day			
	9:30–10:00am		10:00am–3:30pm		3:30–4:00pm		9:30–10:00am		3:30–4:00pm	
	mean	<i>t</i> -stat	mean	<i>t</i> -stat	mean	<i>t</i> -stat	mean	<i>t</i> -stat	mean	<i>t</i> -stat
1	0.0026	4.43	0.0003	0.90	−0.0050	−10.68	0.0023	4.48	−0.0053	−11.88
2	0.0033	7.96	0.0018	12.55	0.0004	1.30	0.0015	3.85	−0.0014	−3.95
3	0.0023	7.72	0.0018	13.39	0.0009	3.14	0.0005	1.69	−0.0009	−3.86
Σ 1-10	0.0207	14.14	0.0131	12.24	0.0065	5.21	0.0076	6.92	−0.0067	−4.68

Panel B: Comparing γ estimates on high volume/imbalance days to all other days

lag	Volume				Imbalance			
	High	Not High	Difference		High	Not High	Difference	
	mean	mean	mean	<i>t</i> -stat	mean	mean	mean	<i>t</i> -stat
1	0.0036	0.0004	0.0033	3.94	0.0018	0.0006	0.0013	2.13
2	0.0039	0.0018	0.0021	4.00	0.0031	0.0020	0.0013	2.88
3	0.0029	0.0019	0.0010	2.62	0.0024	0.0020	0.0005	0.91
Σ 1-10	0.0235	0.0149	0.0086	4.80	0.0192	0.0156	0.0042	2.10

Panel C: Comparing γ estimates from high spread stocks to low spread stocks

lag	Bid-Ask Spread				Relative Spread (Spread / Midpoint)			
	High	Low	Difference		High	Low	Difference	
	mean	mean	mean	<i>t</i> -stat	mean	mean	mean	<i>t</i> -stat
1	0.0050	−0.0019	0.0069	17.66	0.0040	−0.0003	0.0043	7.21
2	0.0031	0.0015	0.0015	3.92	0.0034	0.0019	0.0015	2.96
3	0.0025	0.0018	0.0007	2.27	0.0029	0.0019	0.0010	2.70
Σ 1-10	0.0215	0.0127	0.0087	8.55	0.0199	0.0153	0.0046	2.79

Table 8
VAR Estimates on Days without News

This table reports coefficients on HFTs' net marketable buying from the VAR equation in Table 3 after excluding either stock days with news or stock days within five days of news. The table shows results for lags 1–3, the sum of all 10 lags (Σ), and the 10-lag sum on days without news minus the 10-lag sum on all days ($\Sigma - \Sigma_{NE}$). News days are defined as days with a Factiva news article, an analyst earnings forecast revision, an analyst recommendation change, management earnings guidance, a form 8-K filing, or a large absolute value of the stock's market-adjusted return. Estimates are from the equation where the dependent variable is non-HFT net marketable buying,

$$non-HFT_{NMB,t} = \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t},$$

where R_t is NBBO midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, and $non-HFT_{NMB,t}$ is non-HFT net marketable buying. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method.

Exclude Stock-days with	Lag			Sum of lags 1-10		#Stocks-days
	1	2	3	Σ	$\Sigma - \Sigma_{NE}$	
None Excluded (NE)	0.0007 (2.21)	0.0021 (15.47)	0.0020 (13.12)	0.0161 (20.30)		23,071
Factiva News on						
Day t	0.0018 (4.64)	0.0020 (7.18)	0.0019 (11.74)	0.0155 (17.08)	-0.0006 (-0.92)	10,838
Days $[t-5, t+5]$	0.0027 (2.66)	0.0031 (3.09)	0.0024 (2.61)	0.0175 (7.81)	0.0014 (0.62)	2,227
IBES Forecast, Recommendation, or Guidance on						
Day t	0.0004 (1.26)	0.0018 (13.31)	0.0020 (10.69)	0.0152 (21.23)	-0.0009 (-3.34)	19,112
Days $[t-5, t+5]$	0.0004 (1.09)	0.0018 (5.99)	0.0015 (6.40)	0.0142 (15.79)	-0.0019 (-2.98)	9,298
Form 8-K Filed on						
Day t	0.0006 (1.92)	0.0020 (14.34)	0.0020 (12.77)	0.0158 (21.76)	-0.0003 (-2.00)	22,089
Days $[t-5, t+5]$	0.0007 (1.88)	0.0020 (13.57)	0.0019 (13.41)	0.0160 (20.27)	-0.0001 (-0.28)	17,302
 Market-adjusted Return on Day t 						
> 1%	-0.0001 (-0.17)	0.0019 (8.97)	0.0019 (11.42)	0.0151 (17.08)	-0.0010 (-1.95)	9,478
> 0.5%	-0.0001 (-0.17)	0.0020 (7.52)	0.0015 (8.19)	0.0141 (15.86)	-0.0020 (-2.38)	5,114
> 0.25%	0.0002 (0.37)	0.0022 (6.26)	0.0013 (6.34)	0.0126 (10.72)	-0.0035 (-3.29)	2,577
 Mkt. Adj. Return on Day t > 0.25% or Factiva, IBES, or 8-K news on						
Day t	0.0014 (2.78)	0.0022 (6.58)	0.0014 (5.29)	0.0137 (10.62)	-0.0024 (-2.10)	5,602
Days $[t-5, t+5]$	0.0033 (1.79)	0.0028 (3.90)	0.0011 (1.46)	0.0167 (4.27)	0.0007 (0.18)	835

A Internet Appendix

Table IA.1
Summary of CRSP Universe

This table summarizes 2009 stock-day observations for CRSP common stocks with dual-class stocks removed. The table summarizes market capitalization (*mv*), dollar volume (*dolvol*), and price (*prc*). Market value and dollar volume are in millions. The column *szp* denotes size deciles. Size portfolio breakpoints are computed among NYSE-listed stocks. Size portfolios for year t are formed on December 31st of year $t - 1$. Deciles one through five are small-cap, six through eight are mid-cap, and nine through ten are large-cap.

<i>szp</i>	<i>nstocks</i>	<i>avg</i>	<i>sd</i>	<i>min</i>	<i>q1</i>	<i>q2</i>	<i>q3</i>	<i>max</i>
mv								
1	690	18	23	0	7	12	19	436
2	753	58	56	1	30	44	64	996
3	589	139	100	4	81	114	165	1,525
4	575	305	216	13	188	262	360	5,695
5	414	576	308	24	389	509	672	3,526
6	312	977	447	31	698	883	1,126	4,194
7	275	1,677	728	101	1,193	1,538	1,981	6,643
8	212	2,871	1,092	373	2,164	2,672	3,341	10,955
9	213	5,645	2,450	523	3,926	5,208	6,876	33,010
10	202	35,668	43,505	2,375	13,512	19,839	34,675	415,274
dolvol								
1	690	0.15	1.99	0.00	0.00	0.01	0.04	324.23
2	753	0.45	2.88	0.00	0.01	0.05	0.19	231.69
3	589	1.27	5.30	0.00	0.12	0.34	0.89	590.10
4	575	3.34	13.92	0.00	0.59	1.36	2.97	1,533.59
5	414	7.41	17.36	0.00	1.85	3.63	7.49	1,704.35
6	312	14.23	26.80	0.00	4.14	7.70	15.22	1,674.99
7	275	26.72	34.07	0.06	9.90	17.49	31.94	2,655.60
8	212	44.17	57.97	0.07	17.25	30.10	54.75	7,143.57
9	213	90.52	120.99	0.52	40.24	65.87	106.64	7,129.60
10	202	374.36	556.13	7.96	128.37	218.11	381.70	19,972.16
prc								
1	690	2.13	3.24	0.01	0.54	1.24	2.63	78.00
2	753	4.99	5.10	0.01	1.65	3.51	6.71	62.00
3	589	7.87	7.92	0.05	3.10	5.96	10.06	329.79
4	575	12.59	10.99	0.10	5.74	9.71	16.33	155.72
5	414	17.14	14.09	0.06	8.32	14.09	22.36	195.98
6	312	20.60	12.78	0.26	11.30	17.94	26.66	103.86
7	275	29.12	71.06	0.25	14.26	22.42	32.79	1,549.00
8	212	38.63	58.21	0.75	17.76	26.94	39.49	731.00
9	213	32.17	22.92	0.35	18.77	27.70	39.77	306.58
10	202	45.08	43.05	1.02	24.63	37.39	52.40	622.87

Table IA.2
Sample Universe

This table summarizes 2009 stock-day observations for the set of stocks from which the sample is constructed. The stocks consist of CRSP common equities with dual-class stocks removed. Stocks are also excluded from the sample universe if they fall in the bottom two size deciles, if their price at the end of 2008 is less than \$5, or if average daily dollar volume in December 2008 is less than \$1 million dollars. The table summarizes market capitalization, *mv*, dollar volume, *dolvol*, and price, *prc*. Market value and dollar volume are in millions. The column *szp* denotes size deciles. Size portfolio breakpoints are computed among NYSE-listed stocks. Size portfolios for year t are formed on December 31st of year $t - 1$. Deciles one through five are small-cap, six through eight are mid-cap, and nine through ten are large-cap.

szp	nstocks	avg	sd	min	q1	q2	q3	max
mv								
3	44	175	90	22	114	160	216	674
4	322	319	162	24	215	289	381	2,151
5	347	554	244	28	392	506	655	2,499
6	293	956	421	31	692	875	1,105	4,194
7	261	1,656	695	169	1,189	1,529	1,961	5,977
8	205	2,858	1,057	373	2,169	2,672	3,329	10,955
9	209	5,568	2,216	523	3,915	5,176	6,835	17,693
10	201	35,774	43,588	2,375	13,525	19,917	34,872	415,274
dolvol								
3	44	1.88	3.86	0.00	0.60	1.15	2.16	229.51
4	322	3.36	7.18	0.00	1.02	1.88	3.56	888.89
5	347	6.59	13.04	0.00	1.98	3.64	7.07	1,704.35
6	293	13.31	22.07	0.00	4.11	7.51	14.46	1,283.42
7	261	25.71	31.28	0.06	9.82	17.16	30.99	2,655.60
8	205	43.30	46.58	0.07	17.20	29.88	54.12	1,969.46
9	209	85.76	84.23	0.52	39.87	65.17	105.01	2,761.70
10	201	375.33	557.28	7.96	128.60	218.68	382.83	19,972.16
prc								
3	44	12.66	11.78	0.91	6.84	10.60	15.58	329.79
4	322	14.82	10.17	0.73	7.95	12.12	18.88	111.85
5	347	18.17	11.53	0.53	9.78	15.62	23.60	120.33
6	293	21.51	12.55	0.26	12.62	18.69	27.31	103.86
7	261	30.37	72.71	1.04	15.38	23.34	33.45	1,549.00
8	205	39.33	58.76	1.02	18.28	27.27	39.79	731.00
9	209	32.65	22.84	1.03	19.25	27.98	40.08	306.58
10	201	45.28	43.07	1.02	24.85	37.53	52.49	622.87

Table IA.3
Summary of News Data

This table summarizes news data. News for a stock comes from the Factiva news archive. Panel A shows the distribution among sample stocks in the total number of articles and number of trading days with news. The left two columns in Panel B show the top 10 sources for time-stamped news, and the right three columns shows the number of stamped vs. all articles from three major business news publications.

Panel A: Distribution of Articles

Per Stock	Time-stamped Articles		All Articles	
	Articles	Days with Articles	Articles	Days with Articles
mean	70	29	762	130
sd	147	33	1611	78
0%	1	1	14	12
25%	12	8	128	60
50%	22	16	227	105
75%	83	40	884	207
100%	1106	158	11877	251

Panel B: Source Summary

Top Time-stamped News Sources		Stamp/No-stamp Breakdown for Major Sources		
	Articles		Stamped	All
Dow Jones News Service	871	Dow Jones	4413	6597
Associated Press Newswires	751	Reuters	1317	3234
MidnightTrader	604	Wall Street Journal	174	1556
PR Newswire (U.S.)	488			
Reuters News	365			
Regulatory News Service	341			
Business Wire	337			
MarketWatch	259			
Market News Publishing	234			
DJ em Portugu??s	219			

Table IA.4

Robustness of VAR to Price Feed Latency

This table tests whether potentially mismatched NASDAQ trade and NBBO quote timestamps affect the VAR results in Table 3 by re-running the VAR using NASDAQ BBO midpoint returns rather than NBBO midpoint returns. The NASDAQ trade and NASDAQ BBO timestamps are precisely aligned. Since calculating the NASDAQ BBO is computationally intensive, the VAR uses only a subset of the sample period, January 1st to March 4th of 2009. Panel A reports results with NASDAQ BBO midpoint returns, and Panel B reports results with NBBO midpoint returns over the same time period. For each stock-day observation, the following vector autoregressions (VARs) with ten lags are estimated:

$$R_t = \alpha_1 + \sum_{i=1}^{10} \gamma_{1,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{1,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{1,i} R_{t-i} + \epsilon_{1,t} \quad (9)$$

$$HFT_{NMB,t} = \alpha_2 + \sum_{i=1}^{10} \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \quad (10)$$

$$non-HFT_{NMB,t} = \alpha_3 + \sum_{i=1}^{10} \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^{10} \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^{10} \lambda_{3,i} R_{t-i} + \epsilon_{3,t} \quad (11)$$

where R_t is the bid-ask midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, and $non-HFT_{NMB,t}$ is non-HFT net marketable buying. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method.

Table IA.4 — continued

Panel A: NASDAQ BBO time-series average of mean daily coefficients

lag	γ (HFT)		β (non-HFT)		λ (R)	
	μ	t -stat	μ	t -stat	μ	t -stat
<i>y = non-HFT_t</i>						
1	0.0018	5.99	0.0515	26.96	0.1714	29.85
2	0.0025	17.08	0.0217	33.62	−0.0002	−0.30
3	0.0022	9.64	0.0159	39.73	−0.0002	−0.44
4	0.0022	8.78	0.0139	19.04	0.0004	1.62
5	0.0021	10.02	0.0127	31.66	−0.0002	−0.82
6	0.0020	8.63	0.0101	20.81	0.0006	1.59
7	0.0015	9.56	0.0073	27.62	0.0006	2.37
8	0.0014	7.84	0.0073	16.94	0.0006	3.01
9	0.0015	8.11	0.0075	16.57	0.0003	1.69
10	0.0017	9.67	0.0097	31.78	0.0000	−0.27
<i>y = HFT_t</i>						
1	0.0292	26.26	0.0059	6.92	0.3536	33.51
2	0.0069	5.20	−0.0005	−0.60	0.0148	18.34
3	0.0029	5.11	0.0002	0.43	0.0042	5.48
4	0.0032	9.88	−0.0004	−0.76	0.0020	4.05
5	0.0021	5.47	−0.0006	−0.90	0.0009	2.24
6	0.0011	4.72	−0.0009	−1.73	0.0006	1.24
7	0.0008	2.71	−0.0011	−1.94	0.0000	0.03
8	0.0001	0.19	−0.0010	−2.59	−0.0004	−1.15
9	−0.0001	−0.34	−0.0009	−1.20	−0.0005	−1.56
10	−0.0004	−1.26	−0.0005	−0.93	−0.0011	−4.00
<i>y = R_t</i>						
1	0.0222	18.79	0.0310	42.20	−0.0514	−26.91
2	0.0135	18.97	0.0177	29.16	−0.0359	−26.36
3	0.0088	21.64	0.0142	26.74	−0.0257	−28.86
4	0.0074	20.69	0.0109	18.32	−0.0209	−22.87
5	0.0049	17.89	0.0095	15.41	−0.0172	−27.38
6	0.0039	12.20	0.0067	12.52	−0.0145	−22.60
7	0.0028	12.24	0.0050	10.11	−0.0111	−24.67
8	0.0020	4.77	0.0051	13.58	−0.0089	−26.60
9	0.0017	3.90	0.0062	11.09	−0.0073	−17.33
10	0.0000	−0.10	0.0028	6.42	−0.0070	−10.37

Table IA.4 — continued

Panel B: NBBO time-series average of mean daily coefficients

lag	γ (HFT)		β (non-HFT)		λ (R)	
	μ	t -stat	μ	t -stat	μ	t -stat
<i>y = non-HFT_t</i>						
1	0.0013	4.58	0.0515	27.82	0.5567	20.80
2	0.0025	18.78	0.0218	30.73	0.0031	1.43
3	0.0022	9.84	0.0160	38.39	0.0000	−0.02
4	0.0020	7.45	0.0140	18.35	0.0007	0.59
5	0.0021	9.08	0.0129	31.71	−0.0003	−0.24
6	0.0020	6.91	0.0102	21.32	0.0022	1.76
7	0.0014	8.32	0.0073	28.05	0.0021	2.21
8	0.0013	6.79	0.0074	18.35	0.0025	2.53
9	0.0014	8.33	0.0074	16.87	0.0016	1.92
10	0.0017	8.79	0.0098	28.75	0.0008	1.16
<i>y = HFT_t</i>						
1	0.0306	26.21	0.0080	8.35	1.0500	32.56
2	0.0076	5.59	0.0008	0.94	0.0287	10.59
3	0.0030	5.97	0.0007	1.55	0.0041	1.80
4	0.0031	10.42	−0.0001	−0.21	0.0011	0.76
5	0.0021	5.36	−0.0001	−0.11	−0.0015	−1.14
6	0.0011	4.90	−0.0007	−1.44	−0.0036	−3.19
7	0.0006	2.51	−0.0009	−1.71	−0.0027	−3.98
8	−0.0002	−0.67	−0.0009	−2.29	−0.0039	−10.39
9	−0.0003	−0.69	−0.0009	−1.30	−0.0042	−2.77
10	−0.0005	−1.81	−0.0001	−0.15	−0.0044	−4.26
<i>y = R_t</i>						
1	0.0219	13.64	0.0263	20.90	−0.1260	−30.00
2	0.0138	12.19	0.0156	24.01	−0.1087	−40.03
3	0.0102	14.88	0.0140	10.84	−0.0888	−42.30
4	0.0081	13.95	0.0092	12.38	−0.0758	−36.43
5	0.0067	7.40	0.0079	10.38	−0.0597	−38.59
6	0.0055	8.85	0.0067	11.90	−0.0534	−47.89
7	0.0046	11.02	0.0042	3.64	−0.0423	−32.13
8	0.0026	5.82	0.0034	5.85	−0.0359	−28.80
9	0.0022	6.88	0.0029	3.66	−0.0265	−29.25
10	−0.0002	−0.44	0.0001	0.58	−0.0180	−26.73

Figure IA.1: HFTs' Share of NASDAQ Dollar Volume

This figure shows HFTs' share of dollar volume on the NASDAQ Stock Market. The calculation includes all stocks with CRSP share code 10 or 11 trading on NASDAQ, regardless of listing venue.

Figure IA.2: Liquidity Removing Trades as a Percent of HFT Dollar Volume

This figure shows liquidity removing trades as a percent of HFTs' dollar volume on the NASDAQ Stock Market. Liquidity removing trades are those in which the HFT initiates the trade with a marketable order, which is functionally equivalent to a market order. The calculation includes all stocks with CRSP share code 10 or 11 trading on NASDAQ, regardless of listing venue.

Figure IA.3: Market Share By Trading Venue

Market share is reported as percent of dollar volume. NASDAQ is the NASDAQ Stock Market, NYSE is the New York Stock Market, and TRF is the FINRA Trade Reporting Facility that includes trades that do not occur on a stock exchange (e.g., trades executed in dark pools or by off-exchange market making firms).

Figure IA.4: Response of non-HFT and HFT Net Marketable Buying to a 100 Share Shock to the Other Trader Type

This figure plots the impulse response function of non-HFT net marketable buying to a 100 share shock to HFT net marketable buying, and then the response of HFT net marketable buying to a 100 share shock to non-HFT net marketable buying. The solid line is the estimated response function, and dotted lines indicate 95% confidence intervals. Estimates come from the VAR in Table 3:

$$\begin{aligned}
 R_t &= \alpha_1 + \sum_{i=1}^J \gamma_{1,i} HFT_{NMB,t-i} + \sum_{i=1}^J \beta_{1,i} non-HFT_{NMB,t-i} + \sum_{i=1}^J \lambda_{1,i} R_{t-i} + \epsilon_{1,t} \\
 HFT_{NMB,t} &= \alpha_2 + \sum_{i=1}^J \gamma_{2,i} HFT_{NMB,t-i} + \sum_{i=1}^J \beta_{2,i} non-HFT_{NMB,t-i} + \sum_{i=1}^J \lambda_{2,i} R_{t-i} + \epsilon_{2,t} \\
 non-HFT_{NMB,t} &= \alpha_3 + \sum_{i=1}^J \gamma_{3,i} HFT_{NMB,t-i} + \sum_{i=1}^J \beta_{3,i} non-HFT_{NMB,t-i} + \sum_{i=1}^J \lambda_{3,i} R_{t-i} + \epsilon_{3,t}
 \end{aligned}$$

where R_t is NBBO midpoint return, $HFT_{NMB,t}$ is HFT net marketable buying, $non-HFT_{NMB,t}$ is non-HFT net marketable buying, and J is the number of lags. Variables are measured at the one-second horizon and divided by their standard deviation among all stocks that day. Net marketable buying is shares bought in buyer-initiated trades minus shares sold in seller-initiated trades. Imbalances are divided by 20-day trailing average daily volume to make them comparable across stocks. Equations are estimated for each stock every day. Coefficients are then averaged across all stocks for a day, with stocks weighted by the minimum number of non-zero observations among the three variables. The full-sample estimate is the mean of that daily coefficient time-series. Standard errors correct for autocorrelation using the Newey and West (1994) optimal lag selection method. The impulse response functions assume no contemporaneous effects among variables, so the ordering of the variables does not matter.