

 Open access • Journal Article • DOI:10.1111/JOFI.12759

High-Frequency Trading around Large Institutional Orders — [Source link](#)

Vincent van Kervel, Albert J. Menkveld

Institutions: Tinbergen Institute

Published on: 01 Jun 2019 - Journal of Finance (Wiley-Blackwell)

Topics: High-frequency trading and Market liquidity

Related papers:

- [Does Algorithmic Trading Improve Liquidity](#)
- [High-Frequency Trading and Price Discovery](#)
- [High-Frequency Market Making to Large Institutional Trades](#)
- [Continuous Auctions and Insider Trading](#)
- [The Flash Crash: High-Frequency Trading in an Electronic Market](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/high-frequency-trading-around-large-institutional-orders-2yozkz7fv1>

TI 2017-092/IV
Tinbergen Institute Discussion Paper



High-Frequency Trading around Large Institutional Orders

*Vincent van Kervel*¹

*Albert J. Menkveld*²

¹ Pontificia Universidad Católica de Chile

¹ VU University Amsterdam; Tinbergen Institute, The Netherlands

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

High-Frequency Trading around Large Institutional Orders¹

Vincent van Kervel and Albert J. Menkveld

September 26, 2017

¹Vincent van Kervel, Pontificia Universidad Católica de Chile, Avenida Vicuña Mackenna 4860, Macul 7820436, Santiago, Chile, vincentvankervel@gmail.com and Tinbergen Institute. Albert J. Menkveld, VU University Amsterdam, FEWEB, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands, +31 20 598 6130, albertjmenkveld@gmail.com and Tinbergen Institute. We thank Andrew Ainsworth, Jonathan Brogaard, Björn Hagströmer, Terrence Hendershott, Ron Kaniel, Robert Korajczyk, Charles-Albert Lehalle, Anna Obizhaeva, Neil Pearson, Marc Ponsen, Talis Putnins, Yazid Sharaiha, George Sofianos, René Wells, Avi Wohl, and Shihao Yu for their very useful comments. We are also grateful for comments received during conferences at Alan Turing Institute, NBIM, and Tel Aviv University and seminars at AFM, CFTC, Université Paris Dauphine, Goldman Sachs, SEC, University of Sydney, UNSW, and UTS. We thank Bernard Hosman and Sailendra Prasanna Mishra for excellent research assistance. Menkveld gratefully acknowledges NWO for a VIDI grant. The authors further acknowledge support from APG, DNB, NBIM, Swedbank Robur, and QuantValley/FdR Quantitative Management Initiative.

High-Frequency Trading around Large Institutional Orders

Abstract

Liquidity suppliers lean against the wind. We analyze whether high-frequency traders (HFTs) lean against large institutional orders that execute through a series of child orders. The alternative is HFTs trading “with the wind,” that is, in the same direction. We find that HFTs initially lean against these orders but eventually change direction and take position in the same direction for the most informed institutional orders. Our empirical findings are consistent with investors trading strategically on their information. When deciding trade intensity, they seem to trade off higher speculative profit against higher risk of detection by HFTs and being preyed on.

1 Introduction

Migration to electronic trading created a new type of market participant: high-frequency traders (HFTs). The U.S. Securities and Exchange Commission (SEC) characterized this type as “professional traders acting in proprietary capacity” who use “extraordinarily high-speed and sophisticated computer programs for generating, routing, and executing orders” and end the trading day “in as close to a flat position as possible.” HFTs entered securities markets in the late 2000s, appearing first in equity markets. Their market participation, in percentage terms, is typically a few deciles (SEC, 2010; ESMA, 2014).

High-frequency trading has triggered a great deal of academic study, particularly after Nasdaq released data that labeled HFTs in trades and quotes. The evidence is by and large favorable for HFTs emphasizing reduced bid-ask spreads and increased price efficiency. The evidence is mixed, however, on how HFTs relate to “excess” volatility, such as in flash crashes. Biais and Foucault (2014), O’Hara (2015a), and Menkveld (2016) are examples of survey articles on the young and rapidly growing HFT literature.¹

Relatively unexplored is how HFT relates to trading by an important group of end users of securities markets: institutional investors. Retail investors benefit from a smaller bid-ask spread, since there generally is enough depth at the best quote to execute their order. Institutional investors, however, need to “work their order” by splitting a parent order up into small child orders that are fed to the market sequentially. They care about “implementation shortfall,” that is, the average price at which the entire order executed relative to the price that prevailed when it started executing. In other words, how far did they push the price away from them while executing? Institutional investors care about *cumulative* price impact rather than the half-spread paid on, for example, a single market order. Some have expressed concern that trading costs have increased and attribute

¹Several empirical studies find that HFT activity reduces bid-ask spreads (Hendershott, Jones, and Menkveld, 2011; Hasbrouck and Saar, 2013; Menkveld, 2013; Brogaard, Hagströmer, Nordén, and Riordan, 2015; van Kervel, 2015) and improves price efficiency (Boehmer, Fong, and Wu, 2014; Brogaard, Hendershott, and Riordan, 2014). The effect of high-frequency trading activity on short-term volatility and crashes is mixed: Some studies document a negative correlation (Hasbrouck and Saar, 2013; Chaboud, Chiquoine, Hjalmarsson, and Vega, 2014; Hasbrouck, 2015) whereas others document a positive correlation (Gao and Mizrach, 2013; Ye, Yao, and Gai, 2013; Boehmer, Fong, and Wu, 2014; Kirilenko, Kyle, Samadi, and Tuzun, 2014).

this to HFT presence.²

Objective. Our objective is to document how HFTs trade while an institutional order executes through a series of child trades, and to provide an economic narrative inspired by theoretical studies. Three of them speak directly to our empirical setting. HFTs could be liquidity suppliers who take the other side of the order. They would, for example, be selling in the period when a buy order executes (Grossman and Miller, 1988). Or, they could prey on such buy order by also buying initially to sell when the order is almost finished executing, essentially riding the (transitory) price-pressure wave (Brunnermeier and Pedersen, 2005). Finally, they might need time to learn that the buy order is executing, to trade along with it (i.e., also buy) later (Yang and Zhu, 2015).³ Taking such long position relatively late however only makes money if the order is information-motivated (as opposed to an uninformed order seeking liquidity). These three studies have distinctively different predictions for how HFTs trade as intermediaries while an institutional order executes, and the way that the price responds (see Section 2).

The predicted patterns are taken to the data. The empirical analysis is based on a sample that combines proprietary data on order executions by institutional investors with public HFT trade data. The sample runs from January 1, 2011 through March 31, 2013 and covers trading in Swedish index stocks. The order execution data were provided by four large institutional investors (APG, DNB, NBIM, and Swedbank Robur) and include 801,341 child trades. We construct “parent orders,” which we refer to as meta orders, by stringing together child trades by a single institution in a single stock, possibly spanning multiple days (for details see Section 3.2 on p. 11). The final sample consists of 5,136 such meta orders that, on average, contain 156 child trades and last a bit longer than four hours. Not surprisingly, we find that these orders are directional in the sense that all child trades in a single order are either almost exclusively buys or almost exclusively sells. Finally, institutional orders are large on average: \$2.2 million, or 3.88% when expressed relative

²See, for example, “Institutional Investors Air HFT Concerns” *Financial Times*, September 12, 2011; “Wealth Fund Cautions against Costs Exacted by High-Speed Trading” *The New York Times*, October 20, 2013; and “Berkshire’s Munger: High-Frequency Tradings’ Basically Evil,” *Berkshire Munger*, May 3, 2013.

³Boulatov, Bernhardt, and Larionov (2016) propose a model that also generates back-running but considers the exogenous price impact function of Almgren and Chriss (1999); however, they go beyond two periods and analyze Nash strategies in a continuous-time setting.

to average daily volume.

An important benefit of this particular sample is that HFTs had to reveal their trades on NASDAQ OMX,⁴ which at the time was the dominant market with a market share of two-thirds for the local index stocks.⁵ We select Europe's largest high-frequency trading firms according to *Financial News*: Citadel, Flow Traders, Getco, IAT, IMC, Knight, Optiver, Spire, Susquehanna, and Virtu.⁶ Collectively, their participation rate in trades is almost a third in our sample.

Findings. The empirical analysis yields three main findings. First, the pattern of HFT trading fits none of the three theoretical studies perfectly. We find that HFTs (as a group) appear to lean against the wind when an order starts executing but if it executes more than seven hours (possibly multiple days), they seem to reverse course and trade with the wind. They not only close out their position, but they actually *build* a speculative position in the direction of the order. For buy orders, for example, HFTs sell initially and after seven hours start buying, but lots more than they sold initially. This first result is stated somewhat cautiously, since we are aware that HFTs could have entered offsetting positions in alternative markets or in highly correlated securities. We nevertheless consider this concern unlikely because perfectly correlated securities are hard to find for single stocks, and NASDAQ OMX is by far the largest equity exchange for Swedish stocks.

A more pressing concern is that this first result could be entirely mechanical. Institutional investors *and* HFTs might both respond to certain market conditions, which implies that any correlation in their trading pattern is simply driven by a third factor. To rule out such explanation, we create a placebo sample in which none of the institutional investors were active yet market conditions were similar (in terms of volume, idiosyncratic return, market return, volatility, and (market) order imbalance). We find HFT against-wind trading also in the placebo sample, but not the with-wind trading. The with-wind trading therefore seems to be robustly related to the presence of an institutional order. This finding motivates the remainder of the empirical analysis and the overall message of the paper.

⁴This changed on March 23, 2014, when NASDAQ OMX changed to voluntary reports. Many HFT firms opted to go under the radar and not report their trades. See "Changes to Post Trade Counterparty Visibility in NASDAQ OMX Nordic Blue Chip Shares," *GlobeNewswire*, February 6, 2014.

⁵These numbers are from Fidessa, a trade reporting company (<http://tinyurl.com/ozo8ytm>).

⁶See "Europe's Top 10 High-Frequency Kingmakers," *Financial News*, October 3, 2011.

Late HFT with-wind trading is consistent with the third theoretical study — HFTs need time to learn that an institutional order is executing (Yang and Zhu, 2015). But, as alluded to earlier, taking such speculative position late only makes sense if the orders that one preys on are informed-motivated. This is where we turn next.

The second main finding is that institutional orders appear mostly information-motivated, in particular the ones with long-lasting executions that HFTs eventually trade along with. A couple of findings support this claim. First, on days when institutions trade, they trade typically only one or a few stock names. On the spectrum from stock picking based on informational signals to uninformed trading of the market portfolio, institutions seems to be closer to the stock-picking end. Analysis of the informational content of their orders based on the extent to which prices rebound after their final child trade, supports the conjectured informed trading. More specifically, the information component dominates the transitory component in their order's overall price impact. Finally, in the cross-section of orders, this information component is larger for orders that HFTs trade with-wind on. This association is economically large as the average permanent price impact is 15 basis points instead of 6 after a one standard deviation increase in with-wind flow.

The third main finding is a set of results that frames the observed with-wind trading as strategic interplay between investors and HFTs. The overriding conjecture is that investors optimally trade on a private signal in *full* awareness of HFTs preying on the footprint they leave in the market. More specifically, when deciding how intensively to trade on their signal an investor faces the following trade-off. Trading more intensively raises his speculative profit, but it also raises the risk of being detected by HFTs and preyed on through their with-wind trading. The following findings generate empirical support for this conjecture.

- Investors seem to trade cautiously on their information as this information is revealed only slowly in prices. Half of their private signal is revealed after three hours, on average.
- The ability of HFTs to eventually detect information-motivated executions is demonstrated by them trading most strongly with-wind on the long-lasting orders that are most informed.
- Investors seem aware of actual detection as they decelerate trading after HFTs turn with-wind on their order. Their participation in volume, for example, drops by about 50%.

- Detection is costly to investors, and profitable to HFTs. A one standard deviation with-wind trading by HFTs associates with \$4,480 higher order-execution cost and with \$850 higher HFT gross profit.

Contribution relative to contemporary papers. Our paper contributes to three contemporary papers on trading by institutional investors and HFTs. Korajczyk and Murphy (2016) study how HFTs trade around large orders for a Canadian sample. They assign HFT labels based on trader behavior and infer large orders from aggregate flow by broker-client account. They document against-wind flow initially and with-wind flow later in the course of large order executions. They further find that the implementation shortfall paid on large orders is higher for “stressful trades,” that is, trades for which HFT liquidity provision is predicted to be lower.

Our study differs from that of Korajczyk and Murphy (2016) in three ways. First, we identify end-user net flow, as opposed to client flow per broker. End users often use multiple brokers to execute their orders (Linnainmaa and Saar, 2012). Second, we observe HFT names in the sample and therefore do not need to rely on inference based on behavior. Such inference is necessarily imperfect since speed, for example, is one important HFT characteristic that is hard to observe in trade and quote data. Third, they focus on HFT liquidity provision, whereas our focus is on with-wind trading by HFTs which is more robust in our sample. We develop an economic perspective and support it with further analysis (e.g., show that institutional orders are information-motivated).

The second contemporary paper is that of Tong (2015), who relates the average implementation shortfall to high-frequency trading intensity for a U.S. equity sample. The author averages across all institutional investors who participate in the Abel/Noser dataset and documents that high shortfall days coincide with days of high HFT intensity, both for HFT “market making” and HFT “directional trades.” The benefit of our dataset is that it has *intraday* timestamps (as opposed to daily timestamps) and identifies HFTs by name (as opposed to an exchange-labeled category). This allows us to study *exactly* how HFTs trade during the lifetime of an order.⁷

The third related paper is that of Hirschey (2014), who documents that this second’s HFT

⁷Tong (2015) identifies HFT market making by mean reversion in cumulative net flow, whereas we identify it as leaning against an order. Note that predatory trading on institutional orders also implies mean reversion in net position (Brunnermeier and Pedersen, 2005). For this reason, we prefer to identify market making as leaning against an order.

aggressive flow predicts non-HFT aggressive flow in the next 30 seconds. We add to the author's findings in two ways. First, we focus on institutional investor flow, which is a subset of non-HFT flow. The latter also contains flow by retail investors and other intermediaries, such as proprietary trading by sell-side banks. Second, our results suggest that HFTs respond to investor flow at low frequencies. We believe that HFTs should be thought of more broadly than as machines engaged in sub-millisecond speed races. In fact, we suspect that their superior information technology enables them to generate better signals from "big data" assembled over multiple hours and trade on it profitably.

Some words of caution. Our study faces several limitations and we therefore caution against over-interpretation of the results. First, we do not claim causality. All the results are either unconditional or conditional (regression) associations. We nevertheless believe that the interaction of a set of sophisticated intermediaries with large institutional investors should be of general interest. Moreover, the empirical patterns being more supportive of one theory (i.e., Yang and Zhu, 2015) over others helps us understand more deeply the functioning of securities markets.

Second, HFTs seem to trade with-wind on institutional orders, but they might not be the only ones. The data identify HFTs, but not other important intermediaries such as the proprietary trading desks of sell-side banks or hedge funds. In the data, the order flow of such intermediaries is grouped together with end-user flow.

Third, with-wind trading by intermediaries most likely plagued investors long before the rise of HFTs (Harris, 1997). Intermediaries have existed for centuries. One could even argue that back-running was easier in human-intermediated markets because broker-dealers observed their client flow directly, since it had to pass through their hands. They were not allowed to trade on this information but enforcement was extremely difficult (traders could tip each other off). Our contribution is simply to document with-wind trading on institutional orders without claiming that it is somehow unique to HFTs operating in electronic markets. Indeed, there could very well be *less* of it now. Finally, in a broader sense, cream skimming could also be interpreted as with-wind trading. For example, a New York Stock Exchange specialist would only improve the price if he suspected the order originated from an uninformed investor, thereby raising the price impact for

informed investors in equilibrium (Rock, 1990). Another example is dealers paying for retail order flow (Easley, Kiefer, and O’Hara, 1996).

Fourth, with-wind trading is not necessarily costly to investors as a group. If the HFT sector is competitive, HFTs could willingly incur a (net) loss when initially trading against the wind. Doing so, they might discover large long-lasting orders on which they can then profitably back-run. This implies a transfer from large long-lasting orders to small institutional orders or to retail orders.

2 Price and Trade Patterns, as Predicted by Theory

In this section, we summarize several priors for trading by HFTs around large institutional orders. Figure 1 depicts the expected price pattern and HFT cumulative net flow in the lifetime of an institutional buy order and shortly thereafter. The three panels correspond to market making (Grossman and Miller, 1988), predatory trading (Brunnermeier and Pedersen, 2005), and back-running (Yang and Zhu, 2015). These predictions are briefly discussed in the next three paragraphs.

Grossman and Miller (1988) predict that market-making HFTs sell to an institutional buy order and close out their position in the long run. The price rises when the institutional investor buys in order to compensate risk-averse market makers for supplying “immediacy.” Market makers will offload their position at fundamental values *eventually* but incur a price risk on their inventory in the short run. Selling to the institutional buyer at a (temporarily) elevated price compensates them for such risk. The identifying features are therefore that HFTs sell in the lifetime of the institutional buy order and the price impact is transitory (see Panel (a) of Figure 1).

Brunnermeier and Pedersen (2005) model strategic trading by predators who are aware that an institutional investor is forced to trade for liquidity reasons. Their findings imply that predatory HFTs first trade along with an institutional buy order and thus add to price pressure, but eventually turn around and sell to the institutional order. Since HFTs enter positions at a price below the price at which they exit, they earn a positive profit in expectation. The institutional investor suffers, since price pressure is stronger than what it otherwise would have been. Predatory trading is front-running in a general sense. The key features are as follows: HFTs trade along with the institutional order initially and unwind their position at elevated prices and the price impact is transitory (see


Panel (b) of Figure 1).

Yang and Zhu (2015) propose back-running as intermediaries learning about the execution of a privately informed order. If HFTs could sniff out those orders, they would trade along with them. This is profitable for them, since these informed orders continue to execute until the privately known fundamental value is reached. If HFTs can unwind their position after completion of the institutional order, they earn positive expected profit. Institutional investors suffer as HFTs effectively take part of their “informational rent.” That is, prices are pushed to fundamental values more quickly when HFTs join the trade. The key features are as follows: Buying by HFTs is delayed due to an initial learning period and the price impact is permanent (see Panel (c) of Figure 1).


Figure 1: Price and trade patterns around institutional buy orders predicted by theory

This figure schematically depicts three possible price and trade patterns around institutional buy orders, as predicted by theory. Each panel summarizes the predicted price pattern in the lifetime of an institutional order and shortly thereafter (top graph). It further shows how intermediaries (i.e., HFTs) trade by plotting their cumulative net flow in this period (bottom graph). The three panels correspond to market making (Grossman and Miller, 1988), predatory trading (Brunnermeier and Pedersen, 2005), and back-running (Yang and Zhu, 2015), respectively.

(a) Market making (Grossman and Miller, 1988)




(b) Predatory trading (Brunnermeier and Pedersen, 2005)



Continued on next page.

(Figure 1 continued)

(c) Back-running (Yang and Zhu, 2015)



3 Trade Environment and Data

This section describes the NASDAQ OMX trading environment and presents the public and proprietary datasets. The public dataset contains trades with exchange member identities that are used to identify HFT trades. The proprietary dataset contains the child order execution records of four large institutional investors. These investors cannot be identified in the public data, since they are not exchange members. They use brokers to route their orders to the exchange. Both datasets pertain to trading in the 30 Swedish index stocks from January 1, 2011, through March 31, 2013.

3.1 Trading on NASDAQ OMX

NASDAQ OMX runs mostly like a standard limit-order market to trade their Swedish stocks. The most notable idiosyncratic feature is *ex post* trade transparency on who traded. Trade records that are revealed in real time contain the usual fields, that is, a time stamp in milliseconds, a transaction price, and a transaction quantity. However, at the end of each trading day, NASDAQ OMX also

reveals who was trading with whom for each transaction. This identification is carried out at the exchange member level and therefore does not reveal end users. HFTs, banks, and brokers are exchange members, but not institutional investors who trade through banks or brokers. A total of 89 exchange members were active in our sample.

At the time of our sample, NASDAQ OMX faced competition from other trading venues such as multilateral trading facilities and dark pools. Its market share for exchange-traded volume was 65%. The most active rival exchange, Chi-X, had a 20% market share. The remaining 15% was shared by four other exchanges. An important caveat of this study is that we do not observe the trades by HFTs on these alternative markets. We revisit this issue in the robustness analysis presented in Section 7.

3.2 Public and Proprietary Data

Public data. Two sets of public data are used in this study: equity transactions with member identification and index future returns. Both are obtained from the Thomson Reuters Tick History. As mentioned in the introduction, we use member identities to identify the aggregate net flow of the 10 largest HFTs.

Proprietary data. The proprietary data consist of the child order transactions of four large institutional investors highly active in Swedish index stocks. The data contain detailed NASDAQ OMX execution data, that is, a timestamp (to the second), price, and quantity.

As parent-order information is missing, the best we can do is infer them from child order executions. We string these child trades together into “meta orders” and create observations by stock-period-institution. This is done as follows. First, for each stock and each day, all child order executions by a single institution are aggregated into stock-day-institution orders. Then, following Korajczyk and Murphy (2016), these orders are further strung together across days if there is a child trade in the last half hour of one day followed by a child trade in the first half hour in the next day. It turns out that 577 out of the 5,136 orders span multiple days. We refer to institutional meta orders as institutional orders in the remainder of the manuscript for brevity.

Two filters are applied to establish the sample used in all the analyses. First, institutional orders with low directionality are excluded, since the focus is on orders that built a position in the course of the day (as opposed to intra-day trading strategies). Directionality is based on net flow and defined as the absolute value of the difference between the buy and sell volumes, divided by total volume (all in shares).⁸ Orders with directionality below 0.90 are removed from the sample. This filter removes 11.5% of orders. We find that 95% of the remaining orders consist of either purely buys or purely sells. Second, HFT net flow is winsorized at the 1% and 99% levels. This takes care of extreme outliers in this variable.

3.3 Summary Statistics

Table 1 presents summary statistics of the 5,136 stock-period-institution observations. It reports statistics for the entire sample as well as for institutional buy and sell orders separately.

Institutional orders are large and long-lasting. The average order size is \$2.2 million or 3.88% of average daily volume. It consists of 156 child trades that last 4.26 hours, on average. The fraction of the orders executed through market orders (as opposed to limit orders) is 40.6%, which we consider close to the 34.5% documented for ITG Posit data in the U.S. (O'Hara, 2015b). The average implementation shortfall is \$4,294 or 8.27 basis points when expressed relative to order size. Institutional sell orders are somewhat larger on average and execute in a shorter period of time.

In the period it executes, the institutional order is a sizable part of the market. Its net volume for buy or sell orders is \$2.2 million which compares to an overall volume of \$61.2 million in this period (i.e., it is 3.6%). Its net number of trades is 156 out of 5,507 in the market (2.8%). The directionality is at least 0.99 on average. This level is much above the 0.90 threshold that was used to filter the data (see Section 3.2).

While the institutional order executes, HFTs are highly active. They trade \$13.1 million which is 21% when compared to overall volume. This volume is executed through, on average, 1,514 net

⁸It is inspired by the imbalance measure of Chordia and Subrahmanyam (2004). The precise definition of directionality is $|sells - buys| / (sells + buys)$, where buys and sells are the buy and sell share volume, respectively.

Table 1: Summary statistics

This table presents summary statistics for all order executions in Swedish index stocks from January 1, 2011 through March 31, 2013. An observation is defined as the period (consisting potentially of multiple days) during which an institution executed a (meta) order for a particular stock. Observations are therefore indexed by stock-period-institution (SPI). The table presents SPI means and standard errors (in brackets) of various trade variables for the active institution (Inst), the group of HFTs (HFTs), and the entire market (All). The leftmost column describes the variable with the unit of measurement in brackets. Most are self-explanatory, except for the following: directionality is defined as $|\text{buys}-\text{sells}|/(\text{buys}+\text{sells})$ (note: all buys and sells are based on whether the institution bought or sold, *not* based on which side was the market order), ADV is average daily volume based on the full sample, implementation shortfall is defined as what was paid to execute a buy order minus what one would have paid had one been able to buy all at the midquote price that prevailed just before the order started executing (sell-order definition is analogous; relative implementation shortfall is obtained by dividing by total order size), and HFTs' gross trading revenue pertains to their trading in the lifetime of the institutional order where any nonzero position at the end time is valued at the end-of-day price. Swedish Kronas are converted to U.S. dollars at the sample average exchange rate.

	All inst orders			Inst buy orders			Inst sell orders		
	Inst	HFTs	All	Inst	HFTs	All	Inst	HFTs	All
Avg order size (\$100,000)	22 (1)			19 (1)			27 (1)		
Avg order size relative to ADV (%)	3.88 (0.11)			3.41 (0.13)			4.63 (0.18)		
Avg nr child trades	156 (5)			145 (6)			173 (7)		
Avg duration (hours)	4.26 (0.08)			4.66 (0.10)			3.63 (0.11)		
Avg fraction trades through market orders (%)	40.6 (30.4)			40.2 (30.0)			41.3 (30.9)		
Avg imp shortfall (\$)	4,294 (375)			4,366 (496)			4,181 (566)		
Avg imp shortfall (bps)	8.27 (0.74)			8.21 (1.08)			8.37 (0.86)		
Avg HFT gross trading revenue (\$)	285 (82)			261 (123)			323 (85)		
Avg volume (\$100,000)	22 (1)	131 (3)	612 (14)	19 (1)	145 (4)	675 (19)	27 (1)	108 (4)	508 (20)
Avg directionality (based on \$100,000)	1.00 (0.00)	0.08 (0.00)	0.00 (0.00)	1.00 (0.00)	0.07 (0.00)	0.00 (0.00)	1.00 (0.00)	0.08 (0.00)	0.00 (0.00)
Avg volume (#trades)	156 (5)	1,514 (32)	5,507 (116)	145 (6)	1,718 (45)	6,304 (160)	177 (8)	1,178 (43)	4,194 (153)
Avg directionality (based on #trades)	0.99 (0.00)	0.08 (0.00)	0.00 (0.00)	1.00 (0.00)	0.07 (0.00)	0.00 (0.00)	0.99 (0.00)	0.08 (0.00)	0.00 (0.00)
Number of observations	5,136			3,151			1,985		

trades. Note that this suggests that the average trade size is \$8,652 (\$13.1 million / 1,514) which is about half of what the same ratio yields for institutional child order size: \$14,100. Directionality in HFT trades is only 8%. The relatively small trade size and low directionality are in line with a well known HFT preference for keeping inventory at a minimum (SEC, 2010).

Inspection of start- and end-times of the institutional orders shows the following daily pattern (see Figure 9 in the online appendix). These times are quite evenly dispersed throughout the day with two exceptions. End-times tend to concentrate at the close with 31% of the orders ending in the last fifteen minutes of the trading day. Start- and end-times peak between 15:45 and 16:00, which is right after U.S. equity markets open. These observations led us to flag these orders in later panel regressions by adding dummies.

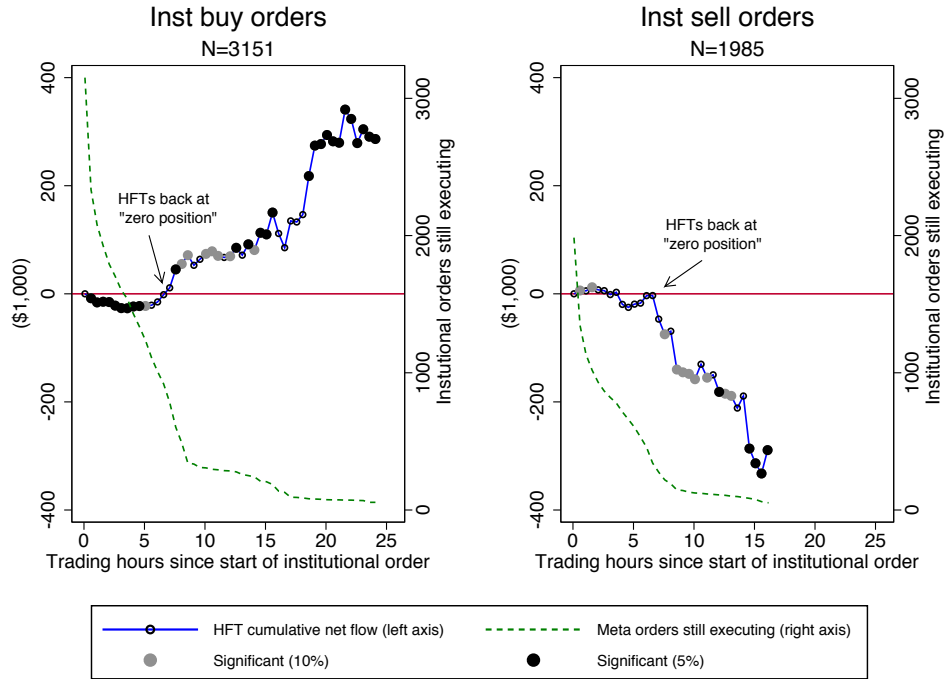
The online-appendix contains an additional set of summary statistics which show that HFTs mean-revert their position substantially within the day, but also across days. Table 7 in this appendix again reports directionality (i.e., $|B-S|/(B+S)$), but this time for stock-days, stock-weeks, and stock-months. It is 0.084, 0.048, and 0.028, respectively. These numbers suggest that HFTs are close to flat at the end of the day, but not entirely. Part of the position they accumulate within a day, they smooth out across weeks or even months. Not mean-reverting entirely to zero is perhaps not surprising, as HFT might hedge their position with correlated assets and obtain a flat position *in terms of risk*. The table further shows that the directionality of institutional investors, on the other hand, is much larger across days, weeks and months: 0.99, 0.91, and 0.83, respectively.

4 HFT Net Flow during the Execution of Institutional Orders

In this section, we investigate whether HFTs “lean against the wind” or “go with the wind” while an institutional order executes through a series of child orders. We also set up a placebo sample to study whether HFT behavior is associated with the presence of the institutional order, or simply relates to the market conditions that prevailed at the time the institution implemented its order.

Figure 2: HFT net flow in the lifetime of an institutional order

This figure plots average HFT net flow from the start of an institutional order to various time points in the lifetime of the order. These time points are multiples of 30 minutes. At each point the average is taken across all institutional orders that are still executing at that point (i.e., the final child trade has not arrived yet). The size of the sample at each time point is indicated by the dashed green line (right y-axis). Statistical significance is tested based on the t -value of the mean across stock-institution fixed effects (same as the overall mean), with residuals clustered at the stock-period level.



4.1 HFT Net Flow in the Lifetime of an Institutional Order

In this subsection, we plot how HFT net flow develops in the lifetime of an institutional order. We track HFT net flow cumulatively, starting from the first institutional child trade. Snapshots are taken every 30 minutes. At the time of the snapshot, HFT net flow is averaged over all institutional orders that are still executing (i.e., their final child trade has not arrived yet). This implies that, for the first half hour, for example, the average is taken across short- and long-lasting orders. We consider this a reasonable approach, since HFTs are unlikely to know the order duration *ex ante* (a view supported by the data, as shown later in Table 3). If the average net flow is negative, then HFTs lean against the order (against wind). If it is positive, then they trade along with it (with wind). An equivalent analysis is done for institutional sell orders.

Figure 2 plots the results for buy orders (left panel) and sell orders (right panel). The plots show that HFTs lean against buy orders in the first seven hours of execution, since they go short for approximately \$20,000. This result is statistically significant only in the first five hours, as indicated by the solid dots. The right panel shows that HFTs lean against sell orders only in the first two hours. This result is statistically weaker (only significant at a 10% significance level).

Strikingly, HFTs turn around and go with the order if it lasts more than seven hours for buys and more than two hours for sells. The result turns statistically significant only if the order lasts longer than seven hours. HFT cumulative net flow steadily grows until it plateaus at around \$300,000 with-wind. It reaches this level after about 20 hours for institutional buys and after 15 hours for institutional sells.

Overall it seems that with-wind trading is slightly stronger for sell orders since initial against-wind trading is weaker and the large with-wind position is reached earlier. We suspect this is due to the higher execution intensity of sell orders. Indeed, Table 1 reveals that sell orders are 38% larger and about an hour shorter, on average. This finding might not be all that surprising as investors might rationally choose to specialize (Van Nieuwerburgh and Veldkamp, 2009). As a result, they are better informed about stocks they own (and consider selling) than stocks they do not own (but consider buying).⁹

The dashed green line indicates the number of institutional orders still executing at each point in time (right y-axis). For buy orders, we observe significant with-wind flow for the 500 orders that are still executing after seven hours. This is about a sixth of all buy orders. For sell orders, there is significant with-wind flow for 200 orders, which is about a tenth of all sell orders.

Two additional results were added to the online appendix and their findings are worth mentioning here. First, Figure 10 shows that in the lifetime of an order the investor trades into a position linearly for both buy and sell orders. Sell orders however execute slightly more intensively, since their slope is about 50% larger in magnitude. This is not surprising since, as noted earlier, sell orders are larger and execute over a shorter horizon. Second, Table 8 shows that treating HFTs

⁹In Section 5.2 we show that larger orders tend to be more informed and there is no *additional* effect from the “meta order is buy” dummy. It therefore is the larger size of sell orders that makes them more informed and therefore exhibit stronger with-wind trading.

as a group seems like a reasonable approach. If, in the aggregate, HFTs trade against-wind in the lifetime of an order, each HFT trades against wind. We find a similar result for with-wind trading. Having said that, we do want to mention that there is some heterogeneity across HFTs in terms of how intensive they trade against-wind or with-wind. Finally, this finding of relative homogeneity for these HFTs is somewhat remarkable given that the HFT literature discusses a variety of HFT strategies (see, e.g., Menkveld, 2016, for a literature survey).

4.2 HFT Net Flow Relative to Placebo Days

A placebo sample is created to identify whether the HFT flow pattern is related to the institutional order or to market conditions that prevailed at the time of the order.

The matching procedure. The placebo sample is constructed by matching each stock-period when one of the institutional investors was active to a similar stock-period for the same stock but when none of the four institutional investors were active. We proceed as follows. The placebo stock-period is selected based on matching five trade variables across two periods:

1. From market open until the first child trade of the order in the “treated” sample and
2. During the lifetime of the order, that is, from the first to the last child trade.

The period from market open until the first child trade is added to account for potential endogenous timing by the institution with regards to when to start executing (Hendershott, Jones, and Menkveld, 2013). The five trade variables used are volume, market return, idiosyncratic return¹⁰, realized volatility (based on one-minute midquote returns), and order imbalance (i.e., standard initiator-signed imbalance or otherwise referred to as market-order imbalance). The imbalance measure should pick up HFTs building or managing positions *in general*, and if the large institutional orders would just be run-of-the-mill position building then any effect would disappear when comparing the treated and untreated samples.

¹⁰We used a beta estimate from Reuters that is based on monthly data for a five-year rolling window (<http://www.reuters.com/finance/markets/indices>).

A “nearest-neighbor” matching procedure is used to identify the placebo days. We follow Davies and Kim (2009, p. 183) with one modification: The distance is measured in standard deviation units as opposed to percentages.¹¹

Note that the placebo analysis controls for a momentum-based explanation of HFT behavior. If institutional orders have a price impact and if HFTs trade on momentum, then a with-wind pattern occurs naturally. However, if the with-wind HFT pattern is present in the original sample *minus* the placebo sample (i.e., in the differential), then this alternative explanation becomes less likely. Both samples are constructed to have the same price pattern in terms of market return and idiosyncratic return.

Table 2 presents trade statistics for both the (treated) institutional order sample and the placebo sample. It shows that the match seems reasonable in terms of distance. Its value is small and evenly distributed across all matched variables.

Results of the matched sample analysis. Figure 3 plots HFT net flow for the treated and the placebo sample. The treated sample line is the same as in Figure 2. The only difference is that the dots now denote significance of HFT net flow tested against placebo HFT net flow, as opposed to testing against zero. In other words, it pertains to a test on whether the *differential* between the two HFT net flows is zero.

In the placebo sample, HFT net flow mostly leans against the wind. It is negative for buy orders throughout and positive for sell orders in the first eight hours. When compared to the placebo sample, treated buy orders show significantly more against-wind trading in the first five hours and more with-wind trading after eight hours. For sell orders results are weaker, and we only

¹¹The relative distance measure of Davies and Kim (2009) is inappropriate when matching on returns because the distance gets large when returns approach zero (division by close-to-zero). We therefore choose to express distance in standard deviation units. Specifically, for the ten matching variables x^k and for treated observation i , we select the matching observation j as follows:

$$\operatorname{argmin}_{j \in J} \sum_{k=1}^{10} \left(\frac{|x_i^k - x_j^k|}{\sqrt{1/J \sum_{l=1}^J |x_i^k - x_l^k|^2}} \right),$$

where j runs over all candidate placebo periods.

Figure 3: HFT net flow in placebo sample

This figure plots average HFT net flow for the placebo sample (and the main sample for reference). It mirrors Figure 2 but this time includes a placebo sample that consists of trading periods that match the “treated” sample in terms of trading conditions (volume, index return, idiosyncratic return, volatility, and market order imbalance) but this time does not contain trades by any of the four institutional investors. Statistical tests pertain to the differential across treated and placebo stock-periods. These are done based on the t -value of the mean across stock-institution fixed effects (same as the overall differential mean), with residuals clustered at the stock-period level.




Table 2: Quality of the match with placebo periods

This table presents the mean and standard deviation of the five variables that were used to construct a placebo sample. The main sample consists of the period when an institution executed an order. Each such period is matched with the same period elsewhere in the sample but this time without an order by any of the four institutional investors. The match is done by a nearest-neighbor algorithm. It uses the average of the five trade variables computed for two time intervals: (i) from market open until the start of order execution (i.e., first child trade) and (ii) from the start of execution until the end of it. The five trade variables are dollar volume, index return, the stock's idiosyncratic return, realized volatility (based on one-minute midquote returns), and Nasdaq market order imbalance (buyer minus seller initiated volume). The distance reported in the table is the average distance across all matching variables. Distance is expressed in standard deviation units.

	Sample stock-periods					Placebo stock-periods					Distance
	Volume	r^{Index}	r^{Idio}	Volatility	Imb	Volume	r^{Index}	r^{Idio}	Volatility	Imb	
<i>Panel A: Institutional buy orders</i>											
Mean (open-start)	14.5	0.0	1.0	5.1	0.6	13.4	-0.2	1.6	4.2	0.6	4.2
Std (open-start)	21.3	79.4	79.7	13.5	2.9	18.6	67.6	64.9	6.9	2.4	
Mean (start-end)	29.5	3.0	7.8	6.1	0.5	0.3	1.9	4.4	4.1	1.2	3.2
Std (start-end)	45.6	98.7	91.1	45.0	3.6	0.4	70.1	75.6	26.6	3.9	
<i>Panel B: Institutional sell orders</i>											
Mean (open-start)	13.5	1.7	-0.5	3.2	0.5	13.0	1.7	-0.1	3.0	0.5	3.6
Std (open-start)	20.2	60.7	75.9	9.6	3.2	17.5	52.2	61.6	6.3	2.1	
Mean (start-end)	21.2	-1.5	-10.5	3.7	0.3	0.2	-0.7	-5.7	3.1	0.7	3.7
Std (start-end)	38.3	53.6	65.2	29.6	3.0	0.3	48.8	53.0	19.2	2.9	

see significantly more with-wind trading between five and seven hours. For sells, there seems to be a with-wind pattern in the placebo sample as well, but this applies to a small sample of orders that are still executing. When pooling buys and sells into pseudo-buy orders,¹² Figure 11 in the online appendix shows that the treated sample has significantly more against-wind flow in the first three hours, and more with-wind flow after seven hours. Taken together, these findings suggest that some of the against-wind pattern is due to market conditions, whereas the with-wind pattern seems truly associated with the presence of the institutional order.

One interpretation of these placebo results is that HFTs use a market-making strategy in normal market conditions. In the treated sample, institutional orders have a price impact (see Table 2), that is, the average idiosyncratic return is positive for buy orders and negative for sell orders. Given that the placebo sample is matched on this variable, it seems HFTs trade against price changes. They sell when prices go up and buy when prices go down, at least initially. The wedge between the two lines suggests that HFTs switch from market making to speculation when they detect a persistent,

¹²Buy and sell orders can be meaningfully pooled into pseudo-buy orders by multiplying HFT cumulative net flow by minus one for sell orders.

directional long-lasting order. They mostly stick to market marking in the placebo sample (except for institutional sell orders that last longer than eight hours in which case the placebo pattern becomes erratic as the number of orders still executing is less than 100, see Figure 2).

It is important to stress that, while there are no order executions by the four identified institutional investors in the placebo sample, other institutional investors could have executed similar orders. We are not overly worried, though, since this would bias against us finding anything. In other words, this observation implies not only that the results that we *do* find are really there, but also that they underestimate the true strength of the effect.

4.3 Do HFTs Detect Large Long-Lasting Orders Early?

The empirical patterns we documented thus far were motivated by the conjecture that HFTs are unable to detect large long-lasting orders early. The HFT cumulative net flow for example one hour into the execution of orders pool across *all* orders that are still executing. In this subsection we analyze whether this is a reasonable conjecture by splitting the sample based on order size or order duration. If the subsample of large, long-lasting orders also exhibits against-wind trading initially, then this supports our conjecture.

The results in Table 3 appear largely supportive. First, if the order sample is split into small and large orders then HFTs trade against-wind initially for both, but only significantly so for buy orders. The sell-order insignificance is not surprising given that there was no such significance in the full sample either. (We report only the odd hours to save space; the even ones show a similar pattern.) These observations hold irrespective of whether the order size is measured relative to a stock's average daily volume (ADV) or measured simply as the order's dollar value. One could even argue that the initial against-wind pattern is stronger for large orders and more significant.

Second, the initial against-wind pattern also holds for both subsamples when they are based on order duration as opposed to size. Long-lasting orders are defined as those that last longer than four hours (half a trading day). The result is again significant only for buy orders.

Table 3: HFT net flow in the lifetime of an institutional order, by subsample

This table documents HFT cumulative net flow in the lifetime of institutional orders where the sample is split based on order size or order duration. For completeness it also contains results for the full sample which are depicted in Figure 2. We only report odd hours for brevity but the complete set of results is available upon request. Statistical significance is established based on the t -value of the mean across stock-institution fixed effects (same as the overall mean), with residuals clustered at the stock-period level. The superscripts ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Hours since first child order execution							
	1	3	5	7	9	11	13	15
<i>Panel A: HFT net flow mean in the lifetime of institutional buy orders</i>								
All	-15.9***	-26.5***	-22.5*	11.1	52.8	70.2*	71.5	110.0**
t-stat	(4.2)	(3.4)	(1.8)	(0.6)	(1.3)	(1.7)	(1.6)	(2.0)
N	2,079	1,603	1,236	785	335	293	259	206
Size ≤ median %ADV	-13.8**	-33.4***	-23.3	24.5	145.0*	115.3	36.2	
t-stat	(2.3)	(2.7)	(1.1)	(0.9)	(1.7)	(1.3)	(0.4)	
N	846	634	468	264	63	60	43	
Size > median %ADV	-17.4***	-21.9**	-22.0	4.4	31.4	58.5	78.6	135.7**
t-stat	(3.5)	(2.3)	(1.4)	(0.2)	(0.7)	(1.2)	(1.6)	(2.2)
N	1,233	969	768	521	272	233	216	186
Size ≤ \$1,000,000	-11.3**	-27.5***	-25.5	0.8	51.2	48.7	-9.1	-97.9
t-stat	(2.4)	(2.9)	(1.6)	(0.0)	(0.9)	(0.8)	(0.2)	(1.3)
N	1,185	893	665	387	106	97	79	47
Size > \$1,000,000	-21.9***	-25.2**	-18.9	21.1	53.5	80.8	106.9*	171.5***
t-stat	(3.5)	(2.0)	(1.0)	(0.8)	(1.0)	(1.5)	(1.9)	(2.6)
N	894	710	571	398	229	196	180	159
Duration < 4 hours	-18.0***	-23.6						
t-stat	(2.5)	(0.9)						
N	640	164						
Duration > 4 hours	-15.0***	-26.8***	-22.5*	11.1	52.8	70.2*	71.5	110.0**
t-stat	(3.3)	(3.3)	(1.8)	(0.6)	(1.3)	(1.7)	(1.6)	(2.0)
N	1,439	1,439	1,236	785	335	293	259	206
<i>Panel B: HFT net flow mean in the lifetime of institutional sell orders</i>								
All	5.5	-1.4	-19.1	-46.8	-145.2*	-155.1*	-188.8*	-313.5***
t-stat	(1.0)	(0.1)	(1.0)	(1.5)	(1.9)	(1.7)	(1.7)	(2.5)
N	1,132	820	603	276	142	117	101	77
Size ≤ median %ADV	7.6	-13.7	15.3	49.1				
t-stat	(0.7)	(0.6)	(0.4)	(0.8)				
N	294	197	126	40				
Size > median %ADV	4.8	2.5	-28.2	-63.0*	-184.2**	-191.8**	-245.3**	-358.5***
t-stat	(0.7)	(0.2)	(1.3)	(1.8)	(2.4)	(2.2)	(2.2)	(2.9)
N	838	623	477	236	132	109	93	72
Size ≤ \$1,000,000	2.5	-10.7	-9.1	24.3				
t-stat	(0.3)	(0.6)	(0.3)	(0.5)				
N	426	278	181	59				
Size > \$1,000,000	7.3	3.4	-23.4	-66.1*	-187.0**	-196.5**	-251.9**	-373.6***
t-stat	(1.0)	(0.2)	(1.0)	(1.8)	(2.3)	(2.1)	(2.1)	(2.9)
N	706	542	422	217	127	104	88	69
Duration < 4 hours	14.0	45.7						
t-stat	(1.4)	(1.5)						
N	409	97						
Duration > 4 hours	0.7	-7.7	-19.1	-46.8	-145.2*	-155.1*	-188.8*	-313.5***
t-stat	(0.1)	(0.7)	(1.0)	(1.5)	(1.9)	(1.7)	(1.7)	(2.5)
N	723	723	603	276	142	117	101	77

5 The Nature of Trading by Institutional Investors

The patterns of how HFTs trade around institutional orders, documented in the previous section, do not simply unequivocally support one of three theoretical patterns summarized in Figure 1: market making, predatory trading, or back-running. Rather, the HFT against-wind pattern is consistent with market making, but is the with-wind pattern a sign of predatory trading or back-running? To answer this question we exploit an important distinguishing feature between the two: Predatory trading hypothesizes that an uninformed investor trades to satisfy a pressing liquidity need, whereas back-running assumes that the investor trades to capitalize on a private (informational) signal. In this section we employ a standard microstructure technique to assess whether at least some institutional orders are information-motivated. And, if so, are these the ones that HFT back-run on?


5.1 Plain-Vanilla Trading of the Market Portfolio or Stock Picking?

But, before turning to a full blown price-impact analysis, some exploratory analysis is useful. If these institutions are uninformed investors trading for liquidity reasons, one expects them to trade the market portfolio. If instead, they trade on private signals then one would expect them to *not* trade the market portfolio (as the market is more cheaply traded through index futures), but small sets of stocks, or even a single stock. This broad distinction suggests computing how many stock (names) an institution trades on days that it trades. This is where we turn next.

Figure 4 plots the histogram of how many names an institution trades when it trades. Trading the market portfolio would entail trading all the 30 constituent names. The histogram reveals that none of the institutions ever traded the entire market portfolio — far from it: They most often traded a single stock only. On average, they traded 5.6 names, with a standard deviation of 5.3. We now turn to a price-impact analysis to study whether such trading could be interpreted as “stock picking” based on information.

Figure 4: How many names did an investor trade on days that he traded?

This figure depicts the empirical distribution of the number of (stock) names that an institutional investor traded on the days that he traded.



5.2 Information in Institutional Orders and Its Correlation with HFT Net Flow

To measure the informational content of institutional orders we extrapolate from the standard microstructure “spread decomposition” approach (Glosten, 1987; Hendershott, Jones, and Menkveld, 2011). This decomposition is used to identify the information in a *single* trade. This is done by simply waiting long enough for the transitory component in the trade’s price impact to disappear, which is often taken to be five minutes. More specifically, the spread decomposition measures the “effective spread” of a (market) buy trade as its immediate price impact: trade price minus the midquote price that prevailed just before the execution, the pre-trade midquote. This trade’s (permanent) “adverse-selection” component is the midquote five minutes later minus the pre-trade midquote. The residual (transitory) component is referred to as the trade’s “realized spread.” Sell trades are pooled with buy trades by multiplying prices by minus one and thus make them “pseudo buys.” Estimates of the effective spread and its two components relies on averaging them across a

large sample of trades.

An equivalent price-impact analysis to assess the informational content of institutional orders that comprise *multiple* trades proceeds as follows. The overall price impact (OPI) for buy orders is defined as the transaction price of the final child trade minus the midquote that prevailed just before execution of the first child trade, the pre-order midquote. The decomposition into transitory and permanent price impact (TPI and PPI, respectively) is then done by waiting 24 hours after the last child executed and use the midquote that prevails then.¹³ Sell orders are pooled with buy orders after converting them to pseudo buys. In all subsequent price-impact analysis we focus on idiosyncratic returns and therefore subtract off from prices the stock's beta times the market return. The reason is that we expect investors to trade on firm-specific information because, if they trade, they trade either a single stock or only a few stocks (see Figure 4).

Figure 5 illustrates the information content for institutional orders sorted by terciles of HFT cumulative net flow, calculated from order start to end. The figure shows that the information content in institutional orders is large on average, and larger for those where HFT trade with-wind. The overall price impact is 8.5 basis points on average, with 8.1 basis points for the tercile of against-wind orders and 11.4 basis points for the tercile of with-wind orders. The permanent price impact is larger than the transitory impact, 5.7 vs. 2.8, respectively. Interestingly, this wedge is larger for with-wind orders, suggesting that, in the cross-section of orders, HFTs trade with-wind on the more informed orders. This finding however might be driven by alternative factors, for example order size, which is why we turn to regression analysis next to remove the effect of various covariates.¹⁴

Table 4 presents the results of price-impact regressions and shows that the positive correlation between permanent price impact and HFT with-wind trading is robust. The main explanatory variables of interest in these regressions are with-wind and against-wind HFT cumulative net flow.

¹³A wait longer than 24 hours could improve accuracy, but at the cost of reduced statistical power, since the fundamental value continues to evolve. We believe a 24-hour wait strikes a good balance, since longer waits imply qualitatively similar but statistically insignificant results.

¹⁴Note that the heterogeneity of orders, the against-wind ones seem liquidity-motivated and the with-wind ones information-motivated, seems to call for a next generation of multi-period optimal-execution models that consider such heterogeneity explicitly. An early example of such model is Choi, Larsen, and Seppi (2017).

Figure 5: Institutional order’s overall, transitory, and permanent price impact, by HFT net flow tercile

This figure plots the average overall price impact (OPI), decomposed into a transitory price impact (TPI) and a permanent price impact (PPI). It plots them separately for the full sample, and the against-wind and with-wind HFT net flow terciles. Sell order observations are converted to pseudo buys by multiplying where appropriate with minus one. They can then be pooled with buy orders. OPI is the stock return over the lifetime of the order. PPI is the return from the start of an order until 24 hours *after* the end of it. TPI is OPI minus PPI. All variables are based on idiosyncratic returns, meaning we subtract beta times the market return from the raw return. OPI, TPI, and PPI are computed by HFT cumulative net flow terciles. This net flow is computed over the lifetime of the order. The lowest tercile has strongly negative values (against-wind), the middle one is mostly zero (neutral), and the highest one has strongly positive values (with-wind).

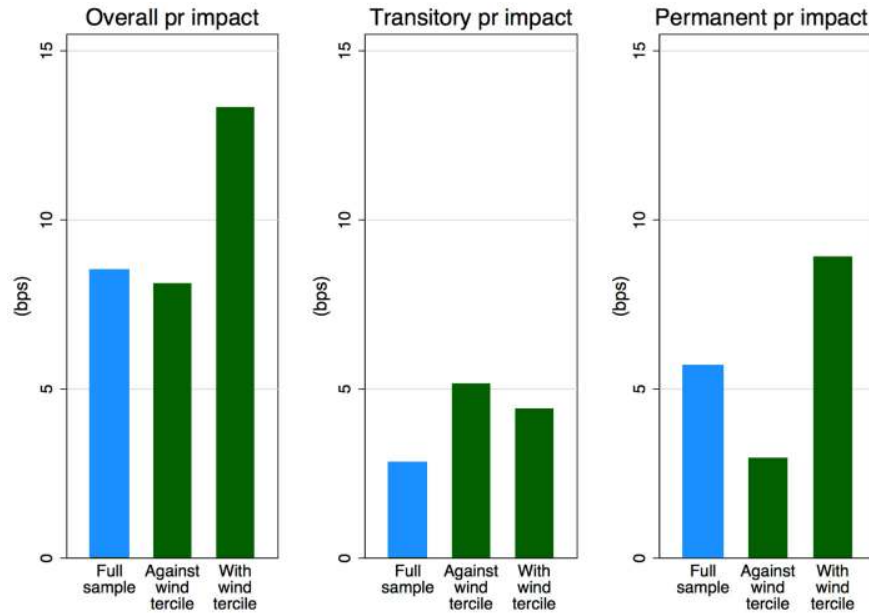


Table 4: Institutional order’s overall, transitory, and permanent price impact regressed on HFT net flow and control variables

This table presents the results of panel regressions where the dependent variable is an institutional order’s overall price impact (OPI), its transitory price impact (TPI), or its permanent price impact (PPI). Sell order observations are converted to pseudo buys by multiplying where appropriate with minus one. They can then be pooled with buy orders. OPI is the stock return over the lifetime of the order. PPI is the return from the start of an order until 24 hours *after* the end of it. TPI is OPI minus PPI. All variables are based on idiosyncratic returns, meaning we subtract beta times the market return from the raw return. The main explanatory variable is HFT net flow accumulated over the lifetime of the institutional order. It is expressed either relative to shares outstanding (%) or in dollar terms (\$). The values are then standardized and subsequently signed based on whether the direction is the same as the institutional order, in which case it receives a positive sign and is loaded into the with-wind variable (which is zero otherwise), or the opposite, in which case it receives a negative sign and is loaded into the against-wind variable (which again is zero otherwise). Order- and market-specific variables are added as controls. Variables names are mostly self-explanatory and should become clear with the following notes. Order size and ADV (based on the full sample) are both measured in shares. The dummies “Start at U.S. open” and “End at U.S. open” are equal to one if the order starts or ends, respectively, in the thirty-minute interval after the opening of the U.S. market. Turnover and volatility are measured from the start to the end of the order. Turnover is the stock volume divided by shares outstanding and volatility is realized volatility based on one-minute midquote returns. The market concentration measure HHI is the Herfindahl-Hirschman Index, defined as the sum of daily squared market shares across lit trading venues. All explanatory variables are standardized. Regressions include stock and institution fixed effects and standard errors are clustered by stock-date. *t*-values are in parentheses. Variable units are in brackets. Superscripts ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	OPI (bps)	TPI (bps)	PPI (bps)	OPI (\$)	TPI (\$)	PPI (\$)
<i>Order-specific variables</i>						
Against-wind HFT cum net flow (%)	-4.74*	1.78	-6.53			
	(-1.7)	(0.5)	(-1.5)			
With-wind HFT cum net flow (%)	6.50*	-2.78	9.28**			
	(1.9)	(-0.8)	(2.1)			
Against-wind HFT cum net flow (\$)				-0.01**	-0.00	-0.01
				(-2.1)	(-0.6)	(-0.8)
With-wind HFT cum net flow (\$)				0.01**	-0.00	0.02**
				(2.5)	(-0.7)	(2.3)
Order size relative to ADV (%)	11.33***	2.79	8.55***	0.02***	0.00	0.02***
	(5.8)	(1.3)	(3.0)	(5.9)	(1.2)	(3.1)
Order duration (hours)	0.99	2.40	-1.42	0.00	0.00	-0.00
	(0.3)	(0.8)	(-0.3)	(0.1)	(0.8)	(-0.4)
Trades through market orders (%)	1.70*	-0.19	1.89	0.00	-0.00	0.00
	(1.8)	(-0.1)	(1.0)	(1.4)	(-0.3)	(1.0)
Meta order is buy dummy	-1.33	-8.02	6.69	-0.00	-0.02	0.02
	(-0.5)	(-1.3)	(1.0)	(-0.1)	(-1.4)	(1.2)
Start at U.S. open dummy	1.01	1.05	-0.04	0.00	0.00	-0.00
	(0.4)	(0.2)	(-0.0)	(0.1)	(0.1)	(-0.0)
End at U.S. open dummy	-9.65***	-8.53	-1.12	-0.02**	-0.02	0.00
	(-2.7)	(-1.4)	(-0.2)	(-2.5)	(-1.6)	(0.1)
Closing auction dummy	-3.20	-5.03	1.83	-0.00	-0.01	0.01
	(-0.8)	(-1.0)	(0.3)	(-0.3)	(-1.5)	(1.0)
<i>Market-specific variables</i>						
Stock volatility (%)	1.13	-0.77	1.89	0.00	-0.00	0.00
	(0.6)	(-0.4)	(0.7)	(0.3)	(-0.9)	(0.8)
Turnover (%)	-1.26	-0.91	-0.35	-0.00	-0.00	-0.00
	(-0.3)	(-0.2)	(-0.1)	(-0.6)	(-0.5)	(-0.1)
Market concentration HHI (%)	-0.78	19.82	-20.60	-0.01	0.07	-0.08
	(-0.0)	(0.5)	(-0.5)	(-0.3)	(1.0)	(-1.0)
Observations	5,136	5,136	5,136	5,136	5,136	5,136
R-squared	0.030	0.002	0.008	0.030	0.002	0.008

The with-wind variable is defined as HFT cumulative net flow if (in the lifetime of the order) it carries the same sign as the institutional order, and zero otherwise. The against-wind variable is defined analogously, such that total HFT flow equals with-wind flow minus against-wind flow. Standard covariates are added as controls.

The OPI regressions show that HFT with-wind flow's coefficient is positive, both when these variables are expressed in relative terms or in dollar terms. Rerunning these regressions for OPI's two components separately reveals that PPI is driving this result. Stronger with-wind HFT flow is associated with more informed orders in the cross-section. This association is economically large as a one standard deviation increase in with-wind flow implies a 9.28 basis points higher PPI, which is high given that the overall PPI average is 5.69 basis points.

The regressions further show that the consistently significant covariate is order size. Its coefficient is significantly positive for both overall and permanent price impact, and is insignificant for transitory price impact. This is consistent with institutional investors using their private information strategically, trading larger quantities when having stronger signals. Note further that the with-wind flow's positive association with the information content of orders (PPI) is therefore not simply an artifact of with-wind activity for *larger* orders. The regression analysis shows that when size is controlled for, the positive association remains.

If the institutional order ends in the half hour after the opening of the U.S markets, the OPI is significantly lower. This is likely due to higher overall trading activity in this period and improved liquidity, as the TPI regression suggests that most of the reduction comes from the transitory component.

Taken together, the findings in this section are consistent with HFTs back-running on informed institutional orders. Not only are the orders in our sample largely informed orders, the with-wind activity is strongest for the most informed among them. We caution that this should be taken as evidence that there is no predatory trading surrounding orders *in general*. Predatory trading requires stressed orders (e.g., forced liquidations) and is therefore more likely to surround orders by highly leveraged institutions (e.g., hedge funds). Such institutions are not part of our sample, which consists of orders by two mutual funds, a sovereign wealth fund, and a pension fund.

6 Putting It All Together: An Economic Interpretation of With-Wind Trading

This final section develops an economic perspective that could rationalize the with-wind trading by HFTs. It first develops predictions based on economic reasoning and then adds further empirical evidence to verify whether these predictions hold up in the data.

6.1 The Interplay between Investors Trading on Information and HFTs

We conjecture that the with-wind trading pattern is driven by an interplay between two types of agents who both trade strategically:

- HFTs are continuously looking for signals on information-motivated trading by investors in order to prey on it.
- Investors are aware of this and try to hide their private signals, but face a trade-off: Trading at high intensity yields a large speculative position (and thus raises profit) but it also increases the risk of being detected and preyed on (which reduces profit).

This interplay leads us to articulate the following four predictions that will be taken to the data in the next four subsections.

Pred1: Investors trade intensively but cautiously on their private information and prices therefore are slow to reveal it.

Pred2: HFTs prey on such information-motivated trading, need time to detect it, but ultimately are able to identify among all long-lasting orders those that are information motivated.

Pred3: Investors are aware of HFT detection risk, yet trade intensively before detection and decelerate after detection.

Pred4: Investors do indeed experience larger implementation shortfall after being detected, and HFTs generate positive gross trading revenue.

6.2 Speed with Which Investor Information Is Revealed in Prices (Pred1)

The first prediction states that investors trade intensively but cautiously on their private information so that prices are slow to reveal it. To measure how successful investors are in this respect, we turn to “unbiasedness regressions.” They have been used before to study the speed of learning (Frenkel, 1977; Hodrick, 1987; Biais, Hillion, and Spatt, 1999). The general specification is:

$$v = [V - E(V|I_0)] = \alpha_t + \beta_t [P_t - E(V|I_0)] + \varepsilon_t, \quad (1)$$

where V is the fundamental value to be learned, $E(V|I_0)$ is its best estimate at the time learning starts, v therefore is the value yet to be discovered at time 0, and P_t is the market price at time t . Performing this regression for various t and plotting the alpha and beta estimates as a function of t yields insight in the speed of learning. Biais, Hillion, and Spatt (1999), for example, use it to study whether the series of indicative prices in the 90 minutes before the Paris Bourse opens reflect information or if they are just noise. In their analysis, v is the close-to-close return, and $[P_t - E(V|I_0)]$ is the close-to-indicative-pre-opening-price-at- t .

We will use these unbiasedness regressions to study how quickly prices reflect possible private information embedded in the institutional investors’ orders. In this application we proxy the dependent variable v with the permanent price impact, PPI. The explanatory variable $(P_t - E(V|I_0))$ is taken to be the midquote at time t minus the midquote that prevailed at the start of the meta-order.¹⁵ The analysis is done for idiosyncratic returns (consistent with the PPI analysis in Section 5.2) and pools sell orders with buy orders by making them pseudo buys (i.e., by changing signs appropriately).

Before turning to the empirical results, it is useful to develop the estimators for alpha and beta a bit further to help interpreting the results. It will turn out to be useful to decompose the regressor into a martingale and a pricing error component. Let

$$m_t = E(V|I_t) - E(V|I_0), \quad (2)$$

be the martingale component¹⁶ and rewrite the regressor as

¹⁵We assume that this midquote at time zero is the best estimate of V at that time: $P_0 = E(V|I_0)$.

¹⁶It is a martingale by construction as takes the expectation of a random variable on expanding information sets.

$$P_t - E(V|I_0) = m_t + e_t \quad (\text{i.e., } \underline{\text{martingale term plus error term}}), \quad (3)$$

where the error term e_t is implicitly defined as the residual after taking out the martingale component. The estimators for alpha and beta can then be written as

$$\hat{\alpha}_t = v - \hat{\beta}_t (m_t + e_t) \quad \text{and} \quad \hat{\beta}_t = \frac{\text{cov}(v, m_t) + \text{cov}(v, e_t)}{\text{cov}(v, m_t) + \sigma_e^2 + 2\text{cov}(m_t, e_t)}. \quad (4)$$

These expressions lead to the following insights. First off, if the price at the end of the meta-order execution reflects its information perfectly (as in, e.g, Kyle, 1985), then $\hat{\beta}_T = 1$ and $\hat{\alpha}_T = 0$. Second, in the hypothetical case of no pricing errors, then once the order starts executing the beta estimate is one, and alpha reveals how much of the information v is in the price already. (Note this case is hypothetical because economic models need noise trading or private-value based trading for informed investors to participate in the game, otherwise one lands in the no-trade theorem (Milgrom and Stokey, 1982).) Third, if there are pricing errors and if they are orthogonal to the information process (driven by m_t and v), then beta is below one and alpha tends to reveal v at a slower rate. Finally, if pricing errors correlate with information, then beta might be above one if this correlation is negative, below one if it is positive, and alpha will tend to reveal the information more quickly or more slowly, respectively.

Figure 6 plots the alpha and beta estimates separately for the tercile of with-wind orders and the tercile of against-wind orders. The with-wind orders are largely information-motivated (see Section 5.2) and the plot should therefore reveal the speed of learning. The against-wind orders are largely uninformed (again Section 5.2) and therefore serve as a benchmark. We run the regressions with t varying from zero to four hours. To keep a balanced sample, all these regressions are done for the subsample of orders that last at least four hours.

The estimated alpha and beta patterns lead to a couple of observations. First, the alpha estimates for the with-wind orders start at about 16 basis points which is equal to the average PPI for these orders. The private signal is revealed slowly over time, and after about two hours one-third appears to be revealed. Alpha keeps declining gradually afterwards, but becomes statistically insignificant.

Half of the private information seems to be revealed after about three hours. For against-wind orders, the alpha estimates are insignificant arguably due to these orders being largely uninformed.

Second, the beta estimates are often significantly below one in the early hours but eventually, after three and half hours, they are statistically indistinguishable from one. The other salient feature is that they fall further below one for with-wind orders. These patterns suggest that pricing errors are relatively large in the early hours, and more so for with-wind orders. Their effect on the beta estimates is not undone by the negative correlation with information (see discussion following (4)). The stronger pattern for with-wind orders suggests that investors do seem to trade very carefully on their information and let pricing errors grow to obfuscate their signal. Or, another way of saying this is: For against-wind orders more of the long-term price impact is public information, which raises the covariance between m_t and v , thus bringing beta closer to one, all else equal.

Third, the empirical patterns for alpha and beta in the with-wind sample seem “reasonable” as they can be generated from an extended Kyle (1985) model. The extension needed is minimal as it simply adds a public information stream to the price process, but does not affect equilibrium strategies. This addition matches real-world price changes as these are unlikely to be *solely* driven by the private information of a monopolist informed agent.¹⁷ Technically this extension does nothing more than add a random walk process to prices. The alpha and beta pattern that such model can generate (obtained through simulations) appears remarkably close to the with-wind sample pattern (see Figure 12 in the online appendix).

In sum, these empirical results are consistent with the first prediction that investors trade stealthily, so that prices are slow to reveal it.

Figure 6: Unbiasedness regressions to measure speed of information revelation

This figure illustrates the speed of information revelation in the first four hours, for the subset of orders that lasted longer than four hours. It does so separately for the tercile of orders where HFTs traded against-wind and the tercile of orders where they traded with-wind. The unbiasedness regression that is estimated is:

$$PPI = \alpha_t + \beta_t (P_t - P_0) + \varepsilon_t,$$

where PPI is the long-term price impact of the order, t measures the time elapsed since the meta order started executing in increments of five minutes, $P_t - P_0$ is the midquote return from the start of order execution until time t . Both PPI and the midquote return are idiosyncratic (i.e., the market return component is removed). The analysis is done on the pooled set of orders where sell orders are converted to pseudo-buy orders. Alphas are tested in terms of being statistically different from zero. Betas are tested relative to one.

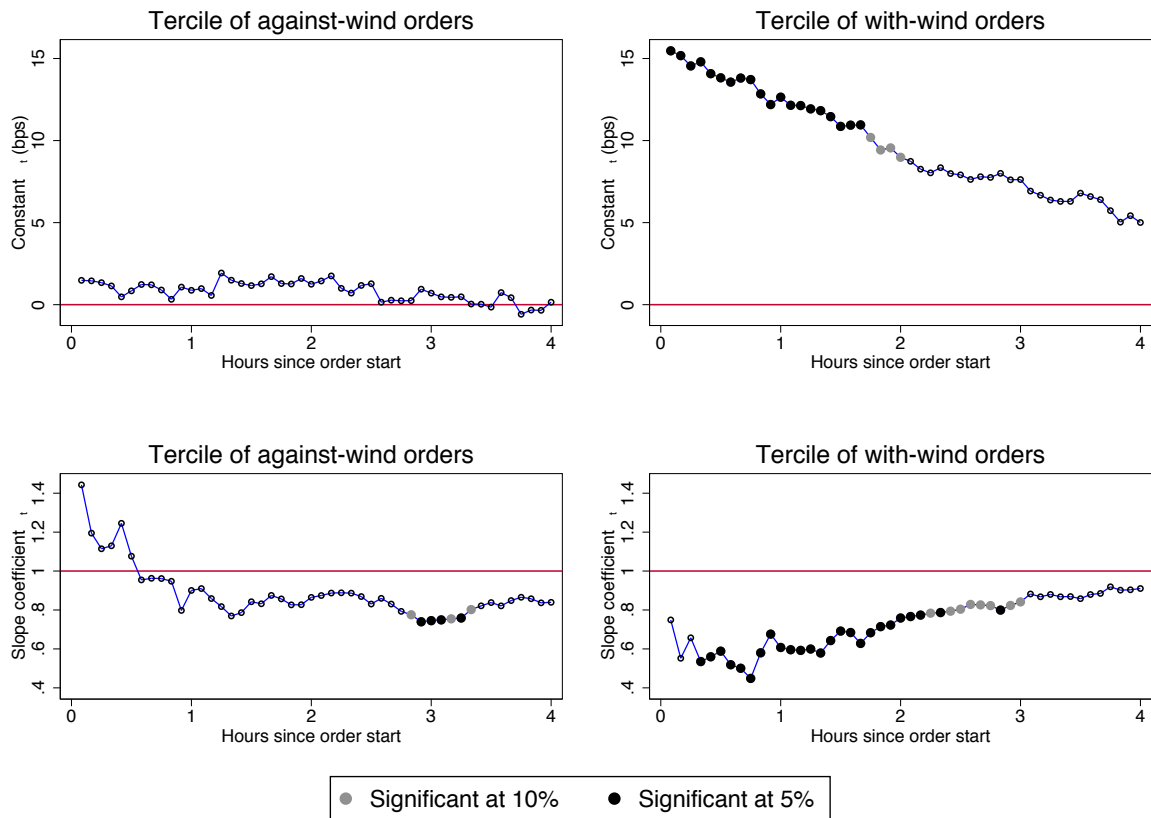



Figure 7: HFT net flow in the lifetime of an institutional order by information content

This figure plots how HFT cumulative net flow develops in the lifetime of institutional orders, where the latter are split in terciles based on their informational content as measured by PPI. It essentially is the 3D version of Figure 2 with informational content on the additional axis. Sell orders are pooled with buy orders by converting them to pseudo-buy orders. The panel below the graph reports coefficient estimates for the odd hours (to conserve space) to reveal what part of it is statistically significant. The analysis is done through the fifteenth hour of execution to retain at least 50 observations in each tercile.



	Hours since first child order execution							
	(1)	(3)	(5)	(7)	(9)	(11)	(13)	(15)
<i>Low PPI tercile</i>								
Mean	-14.5***	-27.7**	-35.8**	-8.1	42.8	36.1	48.9	98.5
t-stat	(2.7)	(2.5)	(2.0)	(0.3)	(0.8)	(0.6)	(0.7)	(1.3)
N	1091	856	643	377	176	153	133	109
<i>Middle PPI tercile</i>								
Mean	-13.9***	-1.8	17.9	42.9	112.0	90.9	59.8	178.4
t-stat	(2.5)	(0.2)	(0.9)	(1.5)	(1.4)	(1.1)	(0.6)	(1.3)
N	968	688	494	271	110	95	85	65
<i>High PPI tercile</i>								
Mean	-8.7	-18.5*	-2.9	31.7	96.6	151.5***	183.2***	224.4***
t-stat	(1.6)	(1.8)	(0.2)	(1.2)	(1.6)	(2.5)	(2.7)	(2.9)
N	1152	879	702	413	191	162	142	109

6.3 HFT With-Wind Trading and Information in Long-Lasting Orders (Pred2)

The Section 5.2 regressions showed that with-wind trading by HFTs correlates with how informed institutional orders are, and the unbiasedness regressions in Section 6.2 showed that this information is revealed in prices very slowly. In spite of such slow revelation, HFTs do seem able to detect these orders and trade along with them. But, how skilled are these HFTs? Is it simply the fact that these orders are long-lasting that makes HFTs trade along with them? Or, are only *some* long-lasting orders informed, and are HFTs able to screen these out to trade along with? The latter would generate more support for our overarching conjecture that investors willingly accept higher risk of being detected when trading high intensity on their information. This is what is stated in the second prediction and will be explored empirically in this subsection.

Figure 7 illustrates HFT cumulative net flow in the lifetime of institutional orders split into terciles, based on their informational content as measured by PPI. The figure is essentially a 3D version of the signature plot in Figure 2, with information on the additional axis. Sell orders are pooled with buy orders by converting them to pseudo-buy orders.

The figure illustrates that HFTs appear to trade with-wind only on the most informed of long-lasting orders. HFTs cumulative net flow reaches its highest level after about 15 hours for the tercile of high PPI (i.e., most informed) orders. The regression results that accompany this graph are reported below it and reveal that these levels are statistically significant. They turn significantly positive after ten hours of trading. The with-wind levels are substantially lower for the other terciles and, more importantly, they are statistically insignificant. Finally, note that for all three terciles HFTs start out leaning against the wind in the early hours. This against-wind trading is statistically significant, albeit marginally so for the high PPI tercile. Note that after between five and seven hours, all terciles show positive levels which is consistent with the turning point reported in Figure 2.¹⁸

¹⁷For example, all traders will observe market-index changes and update their estimate of an equity's value by adding beta times this change.

¹⁸We report regression coefficients only for odd hours to conserve space, but they are representative of all hours. The full regression results are reported in Table 9 in the online appendix.

The evidence presented here is consistent with the second prediction. There is cross-sectional heterogeneity in the long-lasting orders. Some seem clearly more information-motivated and HFTs only significantly trade with-wind on those ones.

6.4 Trading Before and After HFT Detection (Pred3)

As HFTs seem able to detect the informed orders, are investors aware of such detection risk? Do they decelerate trading after being detected? And, is their trading before detection consistent with trading off a higher speculative profit with heightened detection risk? Such trade-off is admittedly hard to identify in the data. Yet, if their pre-detection trading is somehow more intense as compared to their overall trading and they reduce this intensity after being detected, then such pattern is at least consistent with strategic trading in full awareness of HFT detection risk. Such pattern is predicted in Pred3 and verified empirically in this section.

To study trading before and after detection, the trading period of with-wind orders is split into two sub-periods: before and after detection by HFTs. We identify detection as the time, going backwards from the end of a with-wind order, that HFT cumulative net flow crossed zero. The pre- and post-detection periods are labeled pre-WW and post-WW, respectively. We characterize pre-WW and post-WW trading by generating three set of trade statistics. The first set focuses on how the institutional investor traded, the second one characterizes overall trading, and the final set zeroes in on trading by HFTs. These trade variables are also computed for the overall sample (i.e., including against-wind orders) to provide benchmark values.

The results are presented in Table 5 and lead to the following observations. First, the investor statistics suggest that he does trade intensively in the pre-WW period. For example, in that period he trades 6,530 shares per five minutes which is about 50% higher than the investor overall average. This higher volume rate is driven both by trading more often and by trading in larger sizes. Pre-WW he trades six instead of four times per five minutes. His trade size is 2,770 shares pre-WW against his overall average of 1,600 shares per trade. This intense trading however does not lead to an elevated participation in volume. His participation pre-WW is about 20% compared to 24%

Table 5: What prompts HFTs to change from against-wind to with-wind trading?

This table presents the results of an analysis on what could prompt detection, that is, HFTs changing from against-wind to with-wind trading in the lifetime of an institutional order. It reports the average of various trading variables pertaining to the institutional investor, the market, and HFTs. All variables are flow variables, expressed in units per five-minute interval. The averages are calculated for the lifetime of all institutional orders and for the tercile of orders where HFTs trade strongly with the wind. The execution period for this subset of orders is split into a pre-WW and a post-WW sub-period. The separating interval is the last five-minute interval in which the HFT cumulative net flow surpasses zero. The period following this interval, if available, is labeled post-WW; the preceding period is labeled pre-WW. The variable units are in brackets and are reported right after the variable names. Child trades are signed using bulk volume classification (Easley, de Prado, and O’Hara, 2016). Specifically, if the one-second return around a child trade is positive, a purchase is labeled market order and a sale is labeled limit order. For the imbalance variable, buy (sell) indicates the total market volume where the buyer (seller) initiated the transaction. Volatility and variance ratios are computed based on midquote returns. Superscripts ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Variable	Overall avg	Pre-WW	Post-WW	Diff	
Investor	Inst volume (1,000 shares per 5 min)	3.39	6.53	3.02	-3.51	***
	Number of child trades (per 5 min)	3.97	5.92	3.31	-2.61	***
	Size child trades (1,000 shares)	1.60	2.77	1.31	-1.46	***
	Investor trades through market orders(%)	40.41	41.48	38.96	-2.52	***
	Investor participation in share volume (%)	24.24	20.17	10.51	-9.66	***
	Std(5-min “Investor participation in share volume”)	8.86	14.25	8.84	-5.41	***
Market	Stock volume (1,000 shares per 5 min)	30.40	40.11	38.98	-1.13	
	Order imbalance, buy-sell /(buy+sell) (%)	6.49	14.35	7.91	-6.44	***
	Volatility (bps per 5 minutes)	12.83	14.18	15.20	1.02	***
	Variance ratio (1-min-ret var/ 60 times 1-sec-ret var)	0.87	0.88	0.86	-0.02	***
	Variance ratio (5-min-ret var/ 5 times 1-min-ret var)	0.83	0.84	0.83	-0.01	
HFTs	HFT trades through market orders (%)	39.04	41.03	45.51	4.48	***
	HFT participation in share volume (%)	22.64	26.02	25.66	-0.36	

overall.¹⁹ Finally, the more intense pre-WW trading is further evidenced by the extent to which trades were done by market orders (as opposed to passive orders, i.e., limit orders): 41% instead of 40%.²⁰ Incidentally, this small variation could explain why this attribute turns insignificant in explaining PPI in Table 4.

Second, investors do seem to decelerate strongly after HFTs turn to trade along with them. Comparing their pre- and post-WW trading, one notices that he trades significantly less intensive on all of his trade measures. His volume declines by 54%, trade frequency by 44%, trade size by 53%, his volume participation by 48%, and the fraction of market orders by 5%.

Third, we measure to what extent the investor actively managed execution by benchmarking against the commonly used volume-weighted-average-price (VWAP) algorithm. The perfect implementation of VWAP keeps the volume participation rate flat throughout the execution period. In other words, the standard deviation of this rate sampled, say, every five minutes is expected to be zero. In the overall sample, this standard deviation is about 9% which we do not consider excessive given a mean participation rate of 24% (i.e., the investor seems to often resort to a VWAP strategy). Interestingly, pre-WW this standard deviation increases to 14% to then drop significantly back to 9% post-WW. We consider this apparent strong deviation from VWAP in the pre-WW period as consistent with the investor optimally trading on a private signal (as is evidenced by the slow revelation result of Section 6.2).

Fourth, HFTs trade more aggressively after the detection point. Their relative number of trades through market orders increases significantly from 41% to 46%. Their participation in overall volume does not change significantly when comparing post-WW with pre-WW. They seem to resort to market orders to quickly change their position from against-wind to with-wind.


These results provide some empirical support consistent with the third prediction. Investors seem aware of detection risk as they decelerate their trading strongly after HFTs turn to with-trading. They seem to willingly take on this risk, probably to raise their speculative profit, as they

¹⁹Note that participation is far above order size divided by ADV as reported in Table 1. This wedge is mostly driven by short-duration orders where volume participation can be high, yet size expressed in ADV units is low.

²⁰This finding of the investor trading at above-average intensity (at least for the first few hours, i.e., pre-detection) on with-wind orders, combined with the observation that these orders are the most informed ones (Section 5), suggests that order size associates positively with PPI. Figure 13 in the online appendix relates the two directly and shows that this is indeed the case.

Figure 8: Implementation shortfall and HFT gross trading revenue by HFT net flow terciles

This figure plots an institution's implementation shortfall and the HFTs' gross trading revenue averages for the full sample, and the against-wind and with-wind HFT net flow terciles. Buy and sell meta orders are pooled after converting sell-order observations to pseudo buys by multiplying with -1 where needed. All HFT net flow observations are then sorted and put into equal-size bins. The tercile with the lowest values consists of strong against-wind HFT net flows, the middle tercile contains small HFT net flows in terms of size (we label this tercile neutral), and the tercile with the highest values contains strong with-wind HFT net flows.



do trade more intensively prior to being detected. Pre-detection, they trade more often and in larger size.

6.5 Does With-Wind Trading Correlate with Investor Cost and HFT Profit?

(Pred4)

The final prediction states that detection of the investor's informed orders by HFTs is costly to them, and profitable to HFTs. This completes the characterization of the investor-HFT interplay as the investor trading off costly detection with profitable speculation. To study this prediction empirically, we compute the (conditional) correlations between HFT with-wind trading and both the implementation shortfall on investor orders and HFT gross trading revenue.

Before turning to regressions to compute conditional correlations, we first plot the bivariate

relationship for both variables with the type of HFT trading. Figure 8 stratifies the sample by HFT cumulative net flow tercile (similar to Figure 5) and plots the full-sample average and averages for the against-wind and the with-wind tercile. It shows that implementation shortfall is four basis points below average for against-wind flow and nine basis points above it for with-wind flow.²¹ The strongly positive association of with-wind trading and transaction cost is consistent with the conjectured trade-off. We caution that this effect is likely to be overstated whereas the effect of against-wind is understated because of endogeneity. These biases result from reverse relationships. For example, against-wind trading causally reduces price pressures in inventory models, but in the data this effect might be dampened by reverse causality. The reason is that large price pressures make against-wind trading more profitable, which may generate a *positive* correlation between the two. For with-wind trading it works the other way. If the price impact is large then it is more likely to be detected by HFTs, and this reinforces the causal effect from with-wind trading to implementation shortfall.

The right-hand side plots of Figure 8 show that HFTs gross trading revenue is positive overall, and for both the against-wind and the with-wind terciles individually. HFTs seem to generate most revenue when trading with-wind: \$700 per order, which is about twice the overall average.

Table 6 complements Figure 8 by presenting *conditional* correlations. It again relates implementation shortfall and HFT gross trading revenue with the type of trading by HFTs, but this time adds control variables. The latter are either variables that further characterize execution of the meta order (e.g., duration of the order) or market characteristics in the lifetime of the order (e.g., volatility). All variables are standardized and therefore results become conditional correlations. The variables of interest — implementation shortfall, HFT gross trading revenue, and HFT against/with-wind trading — are expressed either in dollars or in relative terms for robustness.

The results lead to the following additional insights. First, also conditionally with-wind trading is positively correlated with implementation shortfall. The effect is highly statistically significant (i.e., beyond a 1% level). The same holds for the conditional correlation between with-wind trading

²¹We further note that the implementation shortfall in our sample is the same order of magnitude as that reported in Anand, Irvine, Puckett, and Venkataraman (2013, Table 1), who document an order size-weighted shortfall of 25 basis points. The equivalent number for our sample is 21 basis points.

Table 6: Implementation shortfall and HFT revenue regressed on against/with-wind HFT flow and controls

This table presents the panel regression results where implementation shortfall (IS) and HFT gross trading revenue are the dependent variables. The main explanatory variable is HFT net flow accumulated over the lifetime of the institutional order. It is expressed either relative to shares outstanding (%) or in dollar terms (\$). The values are then standardized and subsequently signed based on whether the direction is the same as the institutional order, in which case it receives a positive sign and is loaded into the with-wind variable (which is zero otherwise), or the opposite, in which case it receives a negative sign and is loaded into the against-wind variable (which again is zero otherwise). Order- and market-specific variables are added as controls. Variables names are mostly self-explanatory and should become clear with the following notes. Order size and ADV (based on the full sample) are both measured in shares. The dummies “Start at U.S. open” and “End at U.S. open” are equal to one if the order starts or ends, respectively, in the thirty-minute interval after the opening of the U.S. market. Turnover and volatility are measured from the start to the end of the order. Turnover is the stock volume divided by shares outstanding and volatility is realized volatility based on one-minute midquote returns. The market concentration measure HHI is the Herfindahl-Hirschman Index, defined as the sum of daily squared market shares across lit trading venues. All explanatory variables are standardized. Regressions include stock and institution fixed effects and standard errors are clustered by stock-date. *t*-values are in parentheses. Variable units are in brackets. Superscripts ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Imp shortfall (bps)	Imp shortfall (\$1,000)	HFT profit ratio (bps)	HFT gr tr rev (\$1,000)
<i>Order-specific variables</i>				
Against-wind HFT cum net flow (%)	-2.72 (-1.4)		0.38 (0.6)	
With-wind HFT cum net flow (%)	13.02*** (6.3)		1.55*** (2.7)	
Against-wind HFT cum net flow (\$)		-1.28 (-1.3)		0.33 (1.0)
With-wind HFT cum net flow (\$)		4.46*** (3.5)		0.86*** (2.6)
Order size relative to ADV (%)	4.63*** (3.4)	8.50*** (6.9)	-0.19 (-0.5)	-0.00 (-0.0)
Order duration (hours)	-0.75 (-0.3)	-0.42 (-0.2)	0.16 (0.3)	-0.21 (-0.8)
Trades through market orders (%)	2.17*** (3.3)	0.87*** (3.3)	-0.17 (-0.9)	-0.02 (-0.3)
Meta order is buy dummy	3.54* (1.9)	1.90 (1.6)	-0.28 (-0.5)	-0.16 (-0.9)
Start at U.S. open dummy	-3.25** (-2.3)	-0.31 (-0.5)	0.42 (0.6)	0.05 (0.3)
End at U.S. open dummy	-5.93*** (-2.6)	-2.26* (-1.8)	0.57 (1.0)	0.09 (0.5)
Closing auction dummy	-1.82 (-0.7)	-2.97** (-2.2)	-0.03 (-0.0)	-0.45 (-1.4)
<i>Market-specific variables</i>				
Stock volatility (%)	2.63** (2.5)	1.19** (2.0)	0.08 (0.2)	-0.06 (-0.4)
Turnover (%)	4.25 (1.6)	2.63* (1.7)	-0.90 (-1.4)	0.73** (2.3)
Market concentration HHI (%)	-7.94 (-0.4)	2.84 (0.3)	13.10** (2.4)	2.68 (1.2)
Observations	5,136	5,136	5,136	5,136
R-squared	0.062	0.058	0.007	0.005

and HFT gross trading profit. Taken together, these results provide more robust evidence on a positive association between with-wind trading the two variable of interest: transaction cost for investors and gross trading profit for HFTs. This is perhaps the strongest support for the fourth prediction.

Second, the against-wind trading does not significantly correlate with against-wind trading in either the implementation-shortfall or the HFT-revenue regressions. The signs are consistent with Figure 8, but there simply is no statistical significance.

Third, the control variables that are statistically significant in the implementation-shortfall regressions mostly confirm our priors. Transaction costs correlate positively with the relative size of the order, with the extent to which trades were done through market orders (as opposed to limit orders), and with market volatility. These signs are all consistent with the microstructure literature. The surprising results are that implementation shortfall is larger for buy orders and for high-turnover periods. These results are statistically the weakest ones, though, with a confidence level of only 90%.

Fourth, the two variables that turn statistically significant in the HFT-revenue regressions are also intuitive. HFTs earn a higher gross profit when turnover is higher (so that they can layoff inventory easier) and their gross profit margin is higher when the market is more concentrated on the incumbent venue (to probably make up for the higher fees).²²

7 Robustness

This section presents the robustness analysis. The results are added to the online appendix.

HFT trading around institutional orders in volume clock time (instead of calendar clock time). An alternative way to run the clock to study how HFTs trade around institutional orders is to use volume instead of calendar time. Volume time increments the clock by, for example, an hour when market volume has incremented by 1/8.5 of a stock's average daily volume (a trading

²²Nasdaq trading fees were 0.50 basis points for a market order in this period, compared to its main competitor entrant market Chi-X that charged 0.28 basis points.

day in Stockholm lasts 8.5 hours). We then track HFT cumulative net flow on this new time grid. To avoid issues of concatenating trades across days, we do this analysis only for the first day of an order's execution period.²³ Figure 14 in the online appendix shows that the HFT trading pattern in volume time echoes the pattern in calendar time of the signature plot of Figure 2: against-wind initially and with-wind eventually for long-lasting orders.

The with-wind pattern is somewhat weaker than in the calendar time analysis and this might be due to the following. There is with-wind trading on some orders that are executed in periods of low market volume. These orders by definition are short-lasting in *volume* time. They also have a high participation rate, which makes hiding more difficult, and explains why one sees lots of with-wind flow. However, such orders leave the sample quickly in a volume time analysis. During periods of low volume, the investor may rationally choose to take on more detection risk if his clock ticks in natural time. For example, his information advantage dissipates in calendar time when the trading day finishes (and overnight news arrives that makes his private signal redundant), or when a public news announcement comes closer.

Is with-wind trading an end-of-day phenomenon (i.e., HFTs clearing out their inventory)?

To study whether with-wind trading is an end-of-day-phenomenon, we do the following additional analysis. First, we plot the dispersion of order start- and end-times of the with-wind subsample in Figure 15 in the online appendix. It shows that these times are dispersed throughout the day in a way similar to the full sample (Figure 9). Further analysis reveals that the HFT trading pattern is insensitive to whether the order ends in the final trading hour or not. Figure 16 in the online appendix redoes the paper's signature plot after splitting the sample based on whether an order stopped executing before the final trading hour. The general pattern of against-wind first and eventually with-wind for long-lasting orders emerges for both subsamples. These results suggest that our main result is not an end-of-day-phenomenon.

²³That is, some order might stretch across multiple days according to a calendar time concatenation, but not according to volume time concatenation. To compare apples to apples, we decided to do this analysis only based on the first day of each meta order.

NASDAQ OMX market share. The NASDAQ OMX market share of the exchange-traded volume was 65% at the time of our sample. The HFT cumulative net flow analysis is based on NASDAQ OMX only because this was the sole exchange that reported the trader's identity for each trade (at the end of the trading day). If HFTs trade randomly across exchanges, then the NASDAQ OMX subsample is representative and all the results go through. If not, NASDAQ OMX is likely to be the exchange whose majority of trades execute simply because it is the exchange with the largest market share. This line of thought inspired us to conduct the following robustness check: Select only orders on those stock-days when the NASDAQ OMX market share was beyond some threshold and rerun the main analysis. The largest market share we can select while still having sufficient observations for statistical inference is 70%, which reduces the sample from 5,136 orders to 863 orders. We can then confirm statistical significance of our main result in Figure 17: HFTs initially trade against-wind but eventually trade with-wind for long-lasting orders.²⁴

Unfortunately we cannot go beyond the 70% cutoff level and incomplete HFT data therefore remains a concern, but we are not overly worried for the following reasons. First, if the Nasdaq HFT position were perfectly offset by opposite positions in non-Nasdaq markets, then the HFT Nasdaq position would be economically meaningless, or just “noise”, and this would work against us when trying to establish empirical patterns. The fact that we find statistically significant patterns therefore suggests that such position is not pure noise. Second, Table 7 in the online appendix shows that HFT directionality (on Nasdaq) approaches zero when it is sampled at lower frequencies. HFTs seem to mean-revert at least some of their Nasdaq position across days which they would not need to do if they had perfectly offsetting positions in non-Nasdaq markets. Taken together, these arguments suggest that our HFT cumulative net flow variable is potentially noisy but carries signal as well.

HFTs build position *before* the first child trade. Although we find that HFTs collectively trade against-wind after an investor's first child trade, one could argue that they might have traded with-

²⁴We also redid the signature plot by stratifying the sample into tercile based on market concentration. Figure 18 in the online appendix again shows that the pattern is robust and, not surprisingly, stronger for the high concentration subsample (as HFT cumulative net flow is measured more precisely there since it is only based on Nasdaq data, i.e., on data of the incumbent market).

wind *ahead* of the execution of this first child trade. If the investor polled interest in, for example, dark pools or submitted a non-marketable limit order on an exchange, then this signal could have been picked up by HFTs. It could have prompted them to trade with-wind immediately after receiving such signal. In this case, the against-wind flow after the first child trade could just be due to the unwinding of a predatory with-wind position. We do not observe investor orders (only their trades), so we cannot directly test this hypothesis. Instead, we start tracking HFT cumulative net flow one hour before an investor's first child trade. Figure 19 shows that this flow is small and statistically insignificant. We therefore consider this alternative interpretation highly unlikely (since HFTs must have picked up a signal, if any, at least an hour before an investor's first trade).

8 Conclusion

Our most robust finding is that HFTs trade along with institutional orders, with the wind, if such orders last at least a few hours. Additional analysis suggests that in these cases investors' trading is information-motivated and HFTs try to detect such trading in order to join it. Such HFT with-wind trading is costly to investors and profitable to HFTs.

HFTs appear to need considerable time to detect these informed orders (several hours) which appears at odds with their typical characterization as ultra-high speed traders. We believe HFTs should be more broadly characterized as high-tech traders. Technology enables them to trade at sub-microsecond speed, but also endows them with an ability to parse "big data" to detect informed trading across several hours (such period becomes big data if one includes all message traffic for the security at hand and all related securities).

HFTs trading with-wind on institutional orders does not necessarily improve market quality. One could argue that prices become more efficient in the short run. HFT trading in the same direction as informed investors makes prices reveal private information more quickly. The worrisome side effect is that, in the long run, prices could become less efficient. Institutional investors could discontinue costly analyst research, since informational rents have to be shared with others in the trading process. Research might no longer be privately profitable. This could become socially costly if informational externalities are large (i.e., information benefits the allocation of capital

across entrepreneurs).

We believe the market structure debate should re-center on end-user costs. Data are hard to come by, but it should be in the interest of end users — retail and institutional investors — to make their trade data available. Alternatively, regulators could demand the reporting of more granular data, much in the spirit of what U.S. regulators did after the 1987 crash. Exchanges were required to identify retail orders in the consolidated equity audit trail data (CAUD). For each trade they completed, brokers had to report whether it was a principal or an agency trade and, for agency trades, whether it was for a retail investor or for an institutional investor. This would enable more analysis to inform future debates on market quality. The recent SEC initiative to amend rule 613 and create a consolidated audit tape (for regulatory use only) seems like a step in the right direction.²⁵

²⁵See <https://www.sec.gov/divisions/marketreg/rule613-info.htm>.

References

- Almgren, R., and N. Chriss, 1999, “Value Under Liquidation,” *Risk*, pp. 61–63.
- Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2013, “Institutional Trading and Stock Resiliency: Evidence From the 2007-2009 Financial Crisis,” *Journal of Financial Economics*, 108(3), 773–797.
- Biais, B., and T. Foucault, 2014, “HFT and Market Quality,” *Bankers, Markets & Investors*, (128), 5–19.
- Biais, B., P. Hillion, and C. Spatt, 1999, “Price Discovery and Learning during the Preopening Period in the Paris Bourse,” *Journal of Political Economy*, 107, 1218–1248.
- Boehmer, E., K. Fong, and J. Wu, 2014, “International Evidence on Algorithmic Trading,” *Working Paper EDHEC Business School*.
- Boulatov, A., D. Bernhardt, and I. Larionov, 2016, “Predatory and Defensive Trading in a Dynamic Model of Optimal Execution by Multiple Traders,” Manuscript, Higher School of Economics (Moscow).
- Brogaard, J., B. Hagströmer, L. Nordèn, and R. Riordan, 2015, “Trading Fast and Slow: Colocation and Market Quality,” *Forthcoming Review of Financial Studies*.
- Brogaard, J., T. Hendershott, and R. Riordan, 2014, “High Frequency Trading and Price Discovery,” *Review of Financial Studies*, 27(8), 2267–2306.
- Brunnermeier, M. K., and L. H. Pedersen, 2005, “Predatory Trading,” *Journal of Finance*, 60(4), 1825–1863.
- Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2014, “Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market,” *Journal of Finance*, 69(5), 1540–6261.
- Choi, J. H., K. Larsen, and D. Seppi, 2017, “Information and Trading Targets in a Dynamic Market Equilibrium,” *Working Paper Carnegie Mellon University*.
- Chordia, T., and A. Subrahmanyam, 2004, “Order Imbalance and Individual Stock Returns: Theory and Evidence,” *Journal of Financial Economics*, 72(1), 485–518.
- Davies, R. J., and S. S. Kim, 2009, “Using Matched Samples to Test for Differences in Trade Execution Costs,” *Journal of Financial Markets*, 12(2), 173–202.
- Easley, D., M. L. de Prado, and M. O’Hara, 2016, “Discerning information from trade data,” *Journal of Financial Economics*, 120(2), 269 – 285.
- Easley, D., N. M. Kiefer, and M. O’Hara, 1996, “Cream-Skimming or Profit-Sharing? The Curious Role of Purchased Order Flow,” *Journal of Finance*, 51, 811–833.

- ESMA, 2014, “High-Frequency Trading Activity in EU Equity Markets,” European Securities and Markets Authority Economic Report 2014-1.
- Frenkel, J. A., 1977, “The Forward Exchange Rate, Expectations, and the Demand for Money: The German Hyperinflation,” *American Economic Review*, 67, 653–670.
- Gao, C., and B. Mizrahi, 2013, “Market Quality Breakdowns in Equities,” *Working Paper Rutgers University*.
- Glosten, L., 1987, “Components of the Bid Ask Spread and the Statistical Properties of Transaction Prices,” *Journal of Finance*, 42, 1293–1307.
- Grossman, S. J., and M. H. Miller, 1988, “Liquidity and Market Structure,” *The Journal of Finance*, 43(3), 617–633.
- Harris, L. E., 1997, “Order Exposure and Parasitic Traders,” Manuscript, University of Southern California.
- Hasbrouck, J., 2015, “High Frequency Quoting: Short-Term Volatility in Bids and Offers,” *Working Paper NYU*.
- Hasbrouck, J., and G. Saar, 2013, “Low-latency Trading,” *Journal of Financial Markets*, 16(4), 646–679.
- Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, “Does Algorithmic Trading Improve Liquidity?,” *Journal of Finance*, 66(1), 1–33.
- , 2013, “Implementation Shortfall with Transitory Price Effects,” in *High Frequency Trading: New Realities for Traders, Markets and Regulators*, ed. by D. Easley, M. Lopez de Prado, and M. O’Hara. Risk Books, London, chap. 9.
- Hirschey, N. H., 2014, “Do High-Frequency Traders Anticipate Buying and Selling Pressure?,” *Working Paper University of Texas at Austin*.
- Hodrick, R. J., 1987, *The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets*. Harwood Academic Publishers, Chur, Switzerland.
- Kirilenko, A. A., A. S. Kyle, M. Samadi, and T. Tuzun, 2014, “The Flash Crash: The Impact of High Frequency Trading on an Electronic Market,” *Working Paper CFTC*.
- Korajczyk, R., and D. Murphy, 2016, “High Frequency Market Making to Large Institutional Trades,” *Working Paper Northwestern University*.
- Kyle, A., 1985, “Continuous Auctions and Insider Trading,” *Econometrica*, 53(6), 1315–1335.
- Linnainmaa, J. T., and G. Saar, 2012, “Lack of Anonymity and the Inference from Order Flow,” *Review of Financial Studies*, 25(5), 1414–1456.

- Menkveld, A. J., 2013, “High Frequency Trading and the New Market Makers,” *Journal of Financial Markets*, 16(4), 712–740.
- , 2016, “The Economics of High-Frequency Trading: Taking Stock,” *Annual Review of Financial Economics*, 8, 1–24.
- Milgrom, P., and N. Stokey, 1982, “Information, Trade and Common Knowledge,” *Journal of Economic Theory*, 26, 17–27.
- O’Hara, M., 2015a, “High Frequency Market Microstructure,” *Journal of Financial Economics*, 116, 257–270.
- , 2015b, “High frequency market microstructure,” *Journal of Financial Economics*, 116(2), 257–270.
- Rock, K., 1990, “The Specialist’s Order Book and Price Anomalies,” Manuscript, Harvard University.
- SEC, 2010, “Concept Release on Equity Market Structure,” Release No. 34-61358; File No. S7-02-10.
- Tong, L., 2015, “A Blessing or a Curse? The Impact of High Frequency Trading on Institutional Investors,” *Working Paper University of Iowa*.
- van Kervel, V., 2015, “Competition for Order Flow with Fast and Slow Traders,” *Review of Financial Studies*, 28(7), 2094–2127.
- Van Nieuwerburgh, S., and L. Veldkamp, 2009, “Information Immobility and the Home Bias Puzzle,” *The Journal of Finance*, 64(3), 1187–1215.
- Yang, L., and H. Zhu, 2015, “Back-Running: Seeking and Hiding Fundamental Information in Order Flows,” *Working Paper MIT*.
- Ye, M., C. Yao, and J. Gai, 2013, “The Externalities of High Frequency Trading,” *Working Paper University of Illinois at Urbana-Champaign*.