High-Frequency Trading around Large Institutional Orders[1]

Vincent van Kervel and Albert J. Menkveld

June 18, 2015

[1]Vincent van Kervel, VU University Amsterdam, FEWEB, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands, +31 20 598 2446, 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 Björn Hagströmer, Charles-Albert Lehalle, and Yazid Sharaiha for their extremely useful comments. We thank Sailendra Prasanna Mishra for excellent research assistence. 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 that HFTs go "with the wind" and trade in the same direction. We find that HFTs lean against the order in the first hour, but turn around and trade with the order in the case of multi-hour executions. This pattern could explain why institutional trading cost is 39% lower when HFTs lean against the order (by one standard deviation), but are 64% higher when they go with it.

1 Introduction

Migration to electronic trading created a new type of market participant: high-frequency traders (HFTs). The 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 00s. They first arrived in equity markets. Their market participation, in percentage terms, is typically a couple of deciles.

HFT triggered lots of academic study, in particular after NASDAQ released data that labeled HFTs in trades and quotes. The evidence is by and large favorable for HFT emphasizing reduced bid-ask spreads and increased price efficiency. The evidence is mixed however on how HFTs relate to "excess" volatility, e.g., flash crashes. Jones (2014) surveys the early HFT literature.[1]

Relatively unexplored is how HFT affects trading by an important group of end-users of securities markets: institutional investors. Retail investors benefit from a smaller bid-ask spread as there generally is enough depth at the best quote to execute their order. Institutional investors, however, need to "work their order" by splitting it up into smaller pieces and feeding them to the market sequentially. They care about "implementation shortfall," i.e., the average price at which the entire order executed relative to the price at which it started. In other words, how far did they push the price away from them? Institutional investors care about *cumulative* price impact rather than the half-spread paid on a single child order execution. Some have expressed concern that trading cost has increased, and relate it to HFT presence.[2]

The time trend seems to support institutional-investor concern. Figure 1 plots trading cost in U.S. equity markets from 2005 to 2010, which is the period in which HFT participation grew from

[1] Several empirical studies find that HFT activity reduces bid-ask spreads (Hendershott, Jones, and Menkveld, 2011; Hasbrouck and Saar, 2013; Menkveld, 2013; Brogaard, Hagströmer, Nordèn, and Riordan, 2015; van Kervel, 2015) and improves price efficiency (Brogaard, Hendershott, and Riordan, 2014; Boehmer, Fong, and Wu, 2014). The effect of HFT 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).

[2] See, for example, "Institutional Investors Air HFT Concerns" *Financial Times*, September 12, 2011, "Wealth Fund Cautions Against Costs Exacted by High-Speed Trading" *New York Times*, October 20, 2013 and "Berkshire's Munger: High-Frequency Tradings' Basically Evil," *Berkshire Munger*, May 3, 2013.

Figure 1: Time trend HFT participation and institutional trading cost

This figure plots the overall time trend in HFT equity market participation, institutional trading cost, and effective spread. The latter is measured as "implementation shortfall." This is the average execution price on a large order, expressed relative to the price at the start of order execution (multiplied by -1 for sell orders). Institutional trading cost and effective spread were taken from Anand, Irvine, Puckett, and Venkataraman (2013, Figure 1) who based it on Abel-Noser data. The HFT data was taken from "High Frequency Trading: Evolution and the Future," a report published by Capgemini in 2012.

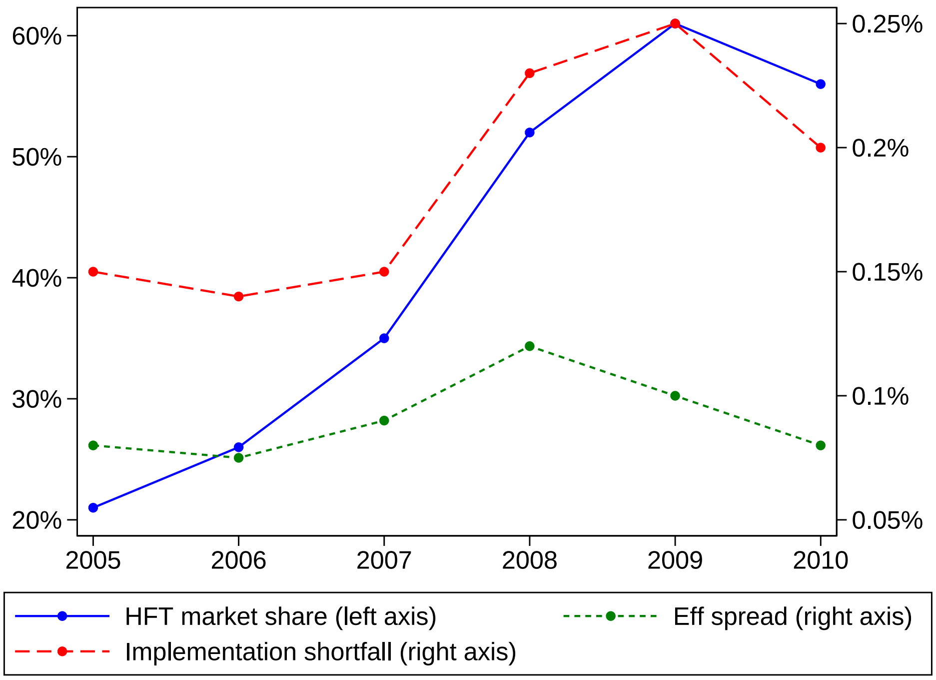

about 21% to 56%. The implementation shortfall for institutional orders grew from 15 basis points to 20 basis points, an increase of 33%. If one takes the effective half spread as a proxy for retail investor cost, then one finds that their cost did not change; it was four basis points at the start and at the end of this period. The 2008-2009 peak in trading cost coincides with the financial crisis and is arguably due to elevated volatility. We should not overinterpret these time trends, but they do seem to warrant further study.

This paper's objective is to relate the HFT trading pattern to implementation shortfall in a direct fashion. We conjecture that HFT net flow (i.e., buy minus sell volume) over the lifetime of an institutional order correlates with the order's implementation shortfall. In particular, HFTs reduce cost when they lean "against the wind," i.e., they trade in a direction opposite to the institutional order. They increase cost when they "go with the wind."

The empirical analysis is based on a sample that combines proprietary institutional-investor execution data with publicly available HFT trade data (no inference needed). The sample runs from January 1, 2011 until March 31, 2013 and pertains to trading in Swedish stocks. The execution data was provided by four large institutional investors (APG, DNB, NBIM, and Swedbank Robur). It consists of 801,341 child order executions. We construct daily "meta orders" by grouping into a single order all the child trades of an institution in a particular stock on a particular day. For brevity, we will refer to these meta orders as institutional orders. The final sample contains 5,910 orders that, on average, contain 135 child order executions. Not surprisingly, we find that these orders are directional, i.e., an institution's child order executions on a particular "stock-day" are either almost exclusively buys or almost exclusively sells. Finally, institutional orders are large on average: $1.940 million or 4.0% when expressed as a percentage of average daily volume.

An important benefit of this particular sample is that HFTs had to reveal their trades at NASDAQ-OMX[3] which was the dominant market with a two-thirds market share.[4] We select the Europe's largest HFT firms according to Financial News: Citadel, Flow Traders, Getco, IAT, IMC, Knight, Optiver, Spire, Susquehanna, and Virtu.[5] Collectively, their participation rate in trades is almost a third in our sample.

The empirical analysis yields two main findings. First, HFTs (as a group) lean against the wind in the first hour of an institutional-order execution, but go with the wind for multi-hour executions. The with-wind flow is so strong that HFT net flow over the lifetime of the order is eventually positive for long-lasting institutional buy orders and negative for sell orders. HFTs therefore seem to be actively taking positions as opposed to simply mean-reverting "inventory." We stated this result cautiously as we are aware that HFTs could have entered offsetting positions in alternative markets, or highly correlated securities. We consider it somewhat unlikely as perfectly correlated

[3]This changed in 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.

[4]These numbers are taken from Fidessa, a trade reporting company (http://tinyurl.com/ozo8ytm).

[5]See "Europe's Top 10 High-Frequency Kingmakers," Financial News, Oct 3, 2011.

securities are hard to find for stocks, and NASDAQ-OMX is by far the largest equity exchange for Swedish stocks in our sample.

Second, the implementation shortfall on institutional orders correlates significantly with HFT net flow, controlling for standard covariates. The average shortfall is 7.4 basis points. It is reduced by 2.9 basis points for a one standard deviation against-wind flow, a reduction of 39%. A one standard deviation with-wind flow increases shortfall by 4.7 basis points, an increase of 64%. The magnitude is larger when implementation shortfall is measured in dollars. Here, a one standard deviation with-wind flow increases it by $2,965, an increase of 104% relative to the sample average of $2,860. We further note that the implementation shortfall in our sample is the same order of magnitude as what is reported in Anand, Irvine, Puckett, and Venkataraman (2013, Table 1). They document an order-size weighted shortfall of 25 basis points. The equivalent number for our sample is 14.7 basis points.

Two additional findings are worth mentioning. First, HFTs do not seem to detect large, long-lasting institutional orders right from start as they lean against the wind initially also for these orders. Second, HFT gross trading revenue is $364 on average for the stock-days where at least one of the institutions is executing an order. We find that, when controlling for standard covariates, a one standard deviation against-wind flow raises their profit by $701. A one standard deviation with-wind flow raises it by $1,164. The latter is 100%*(1,164/2,962)=39% of the institution's increase in implementation shortfall for a one-standard deviation with-wind flow.

Our paper contributes to three contemporary papers on trading by institutional investors and HFT. Korajczyk and Murphy (2014) study how HFTs trade around "large orders" for a Canadian sample. They assign HFT labels based on traders' 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 "effective spread" paid on large orders is higher for "stressful trades." These are trades for which HFT liquidity provision is predicted to be lower.

Our study differs from Korajczyk and Murphy (2014) in four 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 as speed, for example, is one important HFT characteristic that is hard to observe in trade and quote data. Third, Korajczyk and Murphy (2014) apply a minimum threshold of 80% passive trades[6] in their HFT identification scheme. This likely removes lots of HFT speculators from the sample, and therefore retains mostly market-making types.[7] Finally, we measure end-user cost by implementation shortfall. Institutional investors prefer it over a spread measure for two main reasons. First, it measures *cumulative* price impact. Second, they might "earn the spread" along the way as some child order executions might be the result of their price quote being taken by someone else (i.e., an institution trades passively in such case). For example, Menkveld and Yueshen (2015) document how the large seller who reportedly contributed to Flash Crash executed half his trades passively.

The second contemporary paper is Tong (2015) who relates the average implementation shortfall to "HFT intensity" for a U.S. equity sample. She averages across all institutional investors who participate in the Abel-Noser dataset. She 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. Contrary to Tong (2015) we find that "HFT market-making" lowers institutional cost.[8]

The third related paper is Hirschey (2014) who documents that current-second HFT aggressive flow predicts non-HFT aggressive flow in the next 30 seconds. We add to his 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 other intermediaries, e.g., sell-side banks trading for own account. Second, our results suggest that HFT responds to investor flow at low frequencies. This will be a concern to

[6]A trader trades passively if his limit order gets taken by another trader. He trades aggressively when he takes another trader's limit order.

[7]Hagströmer and Nordèn (2013, Figure 1) show that only half a dozen out of 29 HFTs would pass the 80% threshold in their NASDAQ-OMX sample.

[8]Tong (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 a mean-reversion in net position (Brunnermeier and Pedersen, 2005). It is for this reason that we prefer to identify "market-making" as leaning against an order.

large investors who have to spread their flow over multiple hours.

2 Background 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 child order execution records of four large institutional investors. These investors cannot be identified in the public data as 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.

2.1 Trading at NASDAQ-OMX

NASDAQ-OMX runs mostly like a standard limit-order market to trade their Swedish stocks. The most notable idiosyncratic feature is that there is *ex-post* trade transparency on who traded. Trade records that are revealed in real-time contain the usual fields, i.e., a time stamp in milliseconds, a transaction price, and a transaction quantity. But, at the end of each trading day, NASDAQ-OMX also reveals who was selling to who for each transaction. This identification is done at the level of exchange member and therefore does not reveal end-users. HFTs, banks, and brokers are exchange members, but not institutional investors who trade through banks or brokers. 89 exchange members were active in our sample.

At the time of our sample, NASDAQ-OMX faced competition from other regular exchanges and multilateral trading facilities (e.g., 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 five other exchanges.

2.2 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 Thompson Reuters Tick History.

As mentioned in the introduction, we use member identities to identify the aggregate net flow of the ten largest HFTs.

Proprietary data. The proprietary data consists of child order transactions of four large institutional investors who were highly active in Swedish index stocks. The data contain detailed NASDAQ-OMX execution data, i.e., a millisecond time stamp, price, and quantity.

The child order transactions are aggregated to a stock-day-institution level. For each stock and each day, all child order executions by a single institution are aggregated into an institutional "meta order." The rationale for constructing meta orders is that an execution desk at the institution gets orders from different portfolio managers and will internally match buy and sell orders. They will therefore worry about obtaining best execution on *net* flow at the institutional level. We further define meta orders at the daily frequency, as opposed to lower frequencies, as we are interested in the trading behavior of HFTs. The latter are known to make intraday round-trips and prefer to "go home flat." We will refer to institutional meta orders as institutional orders in the remainder of the manuscript for brevity.

Two filters were applied to establish the sample that was used in all analysis. First, institutional orders with low directionality are excluded as 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 buy and sell volume, divided by total volume (all in shares).[9] Orders with directionality below 0.90 are removed from the sample. This filter removes 11.5% of the orders. We find that 95% of the remaining orders consists of either purely buys or purely sells. Second, HFT net flow is winsorized at the 1% and 99% level. This takes care of some extreme outliers in this variable.

Table 1: Summary statistics

This table presents the mean and standard deviation (in brackets) of various trade variables for the active institution (Inst), the group of HFTs (HFTs), and the entire market (All). The sample consists of days where an institution executed an order in a Swedish index stock between January 1, 2011 and March 31, 2013. The left-most column contains a brief description of the variable with, in brackets, the unit of measurement. Most are self-explanatory except the following: net flow is buys minus sells, directionality is defined as |buys-sells|/(buys+sells), ADV is average daily volume based on the full sample, implementation shortfall is defined as what was paid for executing a buy order minus what would have been paid if the price on all child transactions were equal to the midquote price at the time the order started executing (equivalent definition for a sell order; the relative implementation shortfall is obtained by dividing by total order size), and HFT gross trading revenue pertains to their trading in the lifetime of an institutional order where any nonzero position at the end time is valued at end-of-day price. Weighted averages are obtained by weighting with order size. Swedish Kronas are converted to US 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 volume (10,000 shares)	14	170	767	13	179	807	17	155	698
	(39)	(184)	(877)	(36)	(200)	(974)	(43)	(152)	(679)
Avg net flow (10,000 shares)	1.62	0.09	0.00	12.70	0.10	0.00	-17.09	0.08	0.00
	(41)	(22)	(0)	(35)	(22)	(0)	(43)	(21)	(0)
Avg directionality (based on shares)	1.00	0.08	0.00	1.00	0.07	0.00	1.00	0.08	0.00
	(0.01)	(0.07)	(0.00)	(0.01)	(0.07)	(0.00)	(0.01)	(0.08)	(0.00)
Avg volume (#Trades)	139	2,910	10,216	130	3,055	10,810	153	2,665	9,212
	(248)	(2,009)	(6,271)	(235)	(2,178)	(6,690)	(267)	(1,656)	(5,342)
Avg net flow (#Trades)	25	-8	0	129	-5	0	-151	-12	0
	(281)	(307)	(0)	(235)	(299)	(0)	(265)	(321)	(0)
Avg directionality (based on #Trades)	0.99	0.08	0.00	1.00	0.07	0.00	0.99	0.08	0.00
	(0.01)	(0.07)	(0.00)	(0.01)	(0.07)	(0.00)	(0.01)	(0.07)	(0.00)
Avg volume ($100,000)	20	251	1126	18	258	1164	25	239	1063
	(47)	(194)	(1,034)	(45)	(209)	(1,166)	(51)	(164)	(757)
Avg net flow ($100,000)	1.94	0.76	0.00	17.58	0.85	0.00	-24.47	0.61	0.00
	(51)	(35)	(0)	(44)	(36)	(0)	(51)	(33)	(0)
Avg directionality (based on $100,000)	1.00	0.08	0.00	1.00	0.07	0.00	1.00	0.08	0.00
	(0.01)	(0.07)	(0.00)	(0.01)	(0.07)	(0.00)	(0.01)	(0.08)	(0.00)

Continued on next page.

Table 1 continued.

	All inst orders			Inst buy orders			Inst sell orders		
	Inst	HFTs	All	Inst	HFTs	All	Inst	HFTs	All
Avg duration (hours)	3.73			4.09			3.11		
	(3.20)			(3.19)			(3.13)		
Avg order size ($1,000)	1,940			1,758			2,447		
	(5129)			(4,442)			(5,128)		
Avg order size relative to ADV (%)	4.0			3.5			5.0		
	(7.6)			(7.2)			(8.2)		
Avg nr child trades	135			124			153		
	(231)			(221)			(246)		
Avg imp shortfall ($)	2,860			2,464			3,511		
	(33,949)			(34,711)			(32,653)		
Avg imp shortfall (bps)	7.4			6.5			8.8		
	(59.6)			(67.5)			(43.8)		
Wgt avg imp shortfall (bps)	14.7			14.7			14.7		
	(72.6)			(82.4)			(59.2)		
Avg gross trading revenue ($)		364			429			258	
		(3,013)			(3,418)			(2,187)	
Number of observations	5,910			3,675			2,235		

2.3 Summary statistics

Table 1 presents various summary statistics. The top panel shows the trading activity of the four institutional investors, the HFTs, and the market at large. We report statistics for all institutional orders combined, but also separately for institutional buy orders and institutional sell orders.

We observe that if an institution trades on a particular stock-day, it trades 140,000 shares on average. This corresponds to an average order size of $1.940 million. Expressed relative to average daily volume the order size is 4.0%. The meta-orders are extremely directional with an average of 1.00 for both buy and sell orders. These levels are therefore far above the 0.90 threshold we set as a filter.

HFT participation in shares is 22.2% on stock-days where an institution is active. It is 28.5% in terms of number of trades, which implies that their trade size is slightly smaller than average. In dollar terms, their trade size is $8,625. HFTs strongly mean-revert their positions intradaily as indicated by the low average HFT directionality of 0.08. The mean-reversion is even stronger across days, since the average daily HFT *net* flow is 900 shares out of the 1.70 million shares they trade on average. This is in sharp contrast to the institutions who are expected to have longer trading horizons.

The table further reveals various characteristics of the institutional orders. An order generates 135 child trades on average in a time span of 3.73 hours. The average implementation shortfall in dollars is $2,860, which corresponds to 14.7 basis points (an average dollar shortfall is effectively an order-size weighted relative shortfall). The equally-weighted average relative shortfall is only 7.4 basis points, which is not surprising given that larger orders are generally more expensive to execute.

3 HFT net flow in the lifetime 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

[9] It is inspired by the "imbalance" measure of Chordia and Subrahmanyam (2004). The precise definition of directionality is: |S-B|/(S+B), where B and S are buy and sell share volume, respectively.

to study whether HFT behavior is really due to the presence of an institutional order, or is simply the result of market conditions that prevailed on stock-days where institutions implemented their orders.

3.1 HFT net flow in the lifetime of an institutional order

In this subsection we simply plot how HFT net flow develops starting from the first child order execution of an institutional order. As end times we use increments of 30 minutes. For institutional buy orders, for example, we calculate HFT net flow from the start of the order until 30 minutes later for all orders that are "still alive." We then repeat this calculation but now for an interval that ends one hour later and so on. If for any interval HFT net flow is negative, then HFTs leant against the order on average. If it is positive then they "went with it." An equivalent analysis is done for institutional sell orders.

Figure 2 contains the result for buy orders (left panel) and sell orders (right panel). We observe that HFTs lean against buy orders in the first six hours of execution, as they go short for approximately $20,000. This result is statistically significant only in the first four 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 at a 10% significance level).

Strikingly, HFTs turn around and go with the order if it lasts more than six hours for buys, and more than two hours for sells. The result is statistically significant only if the order execution lasted more than seven hours. After eight hours HFTs are long $39,900 for buy orders and short $134,000 for sell orders. The with-wind behavior is stronger for sell orders, as HFTs switch earlier and obtain larger with-wind positions. Arguably, this is due to the higher execution "intensity" of sell orders. Indeed, Table 1 reveals that sell orders are 39.2% larger and are almost an hour shorter on average (3.11 instead of 4.09 hours).

The dashed green line shows how many institutional orders are still alive at each point in time. For buy orders, we observe significant with-wind flow for the 1,010 orders that are still alive after seven hours. This is about a third of all buy orders. For sell orders there is significant with-wind flow for 343 orders, which is about a sixth of all sell orders.

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

The figure plots average HFT net flow from the start of an institutional order until various time points after, sampled at 30-minute intervals. The average is taken across all orders that are still active at the end point of the interval. The size of the sample at each time point is indicated by the dashed green line (right axis). Statistical significance is established based on the *t*-value of the mean across stock-institution fixed effects (same as overall mean) with residuals clustered at a stock-day level.

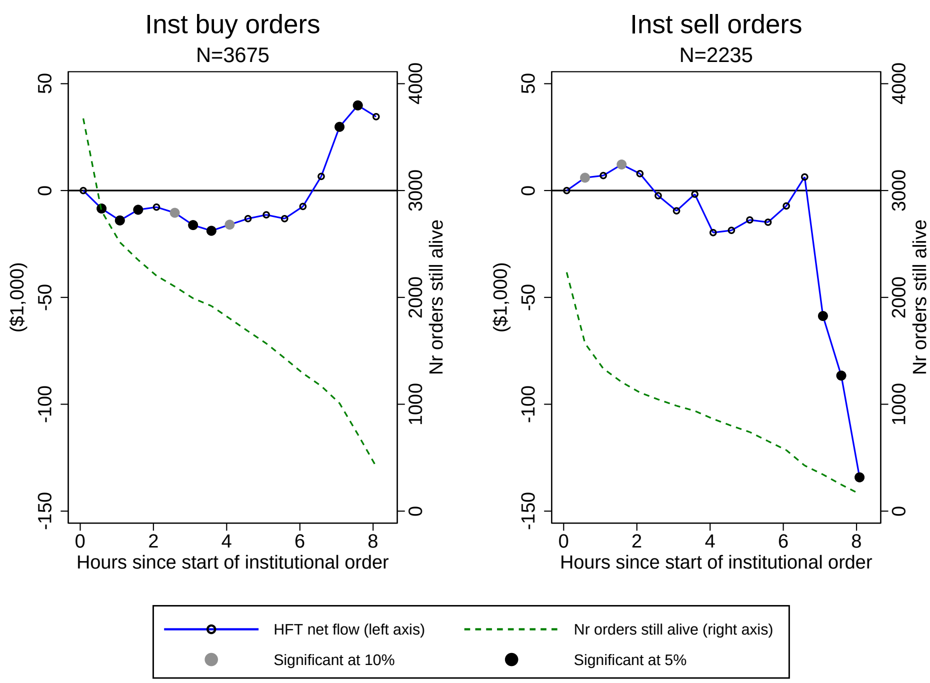

In an unreported figure we plot the institutional net flow over the lifetime of the order (the equivalent of the HFT net flow plot of Figure 2). It reveals that institutional positions build up at an almost linear rate for both buy and sell orders. It further shows that institutional sell orders execute more "aggressively" as the slope is about 25% larger in magnitude.

3.2 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-day for which one of our institutional investors was active to a "similar" stock-day for the same stock but where none of our four institutional investors was active. We proceed as follows. The placebo stock-day is selected based on matching four trade variables across two periods:

1. market open until the first child trade of the order in the "treated" sample
2. the lifetime of the order, i.e., the first to last child-order execution.

The open-till-first-execution period is added to account for potential endogenous timing by the institution when it comes to starting order execution (Hendershott, Jones, and Menkveld, 2013). The four trade variables used are: volume, market return, idiosyncratic return (with a beta obtained from Reuters), and realized volatility (based on one-minute midquote returns). A "nearest neighbor" matching procedure is used. We follow Davies and Kim (2009, p. 183) with one modification, i.e., the distance is measured in standard-deviation units as opposed to percentages.[10]

Note that the placebo analysis controls for a momentum based explanation of HFT behavior. If institutional orders have price impact, and if HFTs trade on momentum, then a with-wind pattern

[10]The 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 (almost) zero). We therefore choose to express distance in standard-deviation units. Specifically, for the eight matching variables x^k and for treated observation i, we select j as:

$$\underset{j\in N}{\operatorname{argmin}} \sum_{k=1}^{8} \left(\frac{|x_i^k - x_j^k|}{\sqrt{1/N \sum_j \left(|x_i^k - x_j^k|\right)^2}} \right).$$

Table 2: Quality of the match with a placebo sample

This table presents the mean and standard deviation of four variables that were used to construct a "placebo" sample. The main sample consists of stock-days where one of the institutions executed an order. Each such stock-day is matched to a stock-day without such order by any of the four institutional investors. The match is done by a "nearest neighbor" algorithm. It uses the average of four trade variables computed for two time intervals: (i) market open until the start of the order and (ii) start of the order until the end of the order. The four trade variables are dollar volume, index return, the stock's idiosyncratic return, and realized volatility (based on one-minute midquote returns). The distance is the average distance across all variables in a particular time period. Distance is expressed in standard-deviation units.

	Sample stock-days				Placebo stock-days				
	Volume	r^{Index}	r^{Idio}	Volatility	Volume	r^{Index}	r^{Idio}	Volatility	Distance
Panel A: Institutional buy orders									
Mean (open-start)	11.3	3.5	-0.2	6.3	10.5	2.3	0.4	5.3	1.4
St dev (open-start)	20.7	75.8	75.6	15.5	18.3	60.8	64.1	10.8	
Mean (start-end)	17.2	2.6	6.3	10.0	16.9	0.0	6.0	8.2	1.5
St dev (start-end)	25.5	93.8	73.0	23.6	24.3	69.0	58.5	18.1	
Panel B: Institutional sell orders									
Mean (open-start)	10.5	0.5	-2.6	3.9	10.2	0.9	-0.4	3.8	1.2
St dev (open-start)	17.9	46.0	71.0	6.7	16.1	41.2	62.5	5.6	
Mean (start-end)	12.4	-2.1	-9.6	4.5	12.2	-0.8	-7.5	4.2	1.2
St dev (start-end)	20.6	53.4	53.9	15.3	19.2	46.5	44.7	12.9	

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 were constructed to have the same price pattern, in terms of market return and in terms of idiosyncratic return.

Table 2 presents trade statistics for both the (treated) institutional trade 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 matched sample analysis. Figure 3 plots HFT net flow for the "treated" sample and the placebo sample. The treated sample line is the same as Figure 2. The only difference is that the dots now denote significance of the HFT net flow tested against the placebo HFT net flow as opposed to zero. In other words, it pertains to a test on whether the *differential* between the two

Figure 3: HFT net flow on placebo stock-days

The figure plots average HFT net flow for the placebo sample (and the main sample for reference). It echoes Figure 2 that was done for the sample of stock-days where one of the four institutional investors executed an order. The placebo sample consists of stock-days that match these "treated" stock-days in terms of trading conditions (volume, index return, idiosyncratic return, and volatility) but did not feature trades by any of these four investors. Statistical tests pertain to the differential across treated and placebo stock-days. They are done based on the *t*-value of the mean across stock-institution fixed effects (same as overall differential mean) with residuals clustered at a stock-day level.

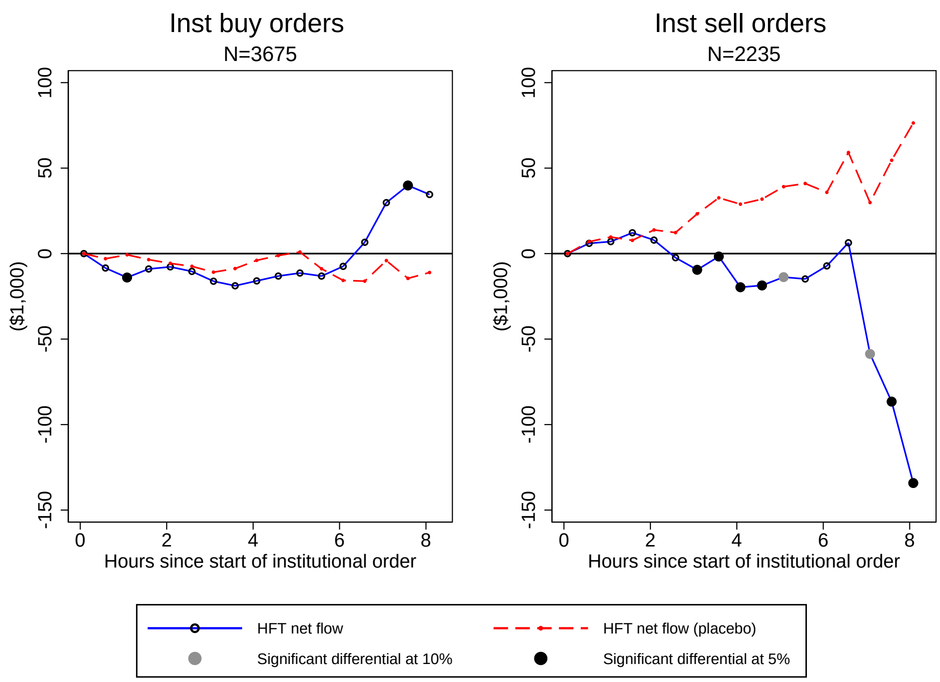

HFT net flows is zero.

In the placebo sample, HFT net flow mostly leans against the wind. It is negative for "buy" stock-days and positive for "sell" stock-days. When compared to the placebo sample, the HFT against-wind pattern in the treated sample becomes mostly insignificant whereas the against-wind becomes more often statistically significant. These findings suggest that the against-wind pattern is due to "market conditions," whereas the with-wind pattern seems truly related to 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 price impact (see Table 2), i.e., the average idiosyncratic return is positive for buy orders, 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 they do so initially. The wedge between the two lines suggests that HFTs switch from market-making to speculation when they detect a persistent, directional, long-lasting order. They stick to market-marking in the placebo sample.

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

3.3 Do HFTs detect large, long-lasting orders early?

In this subsection we repeat the HFT net flow analysis for various subsamples of the data. Specifically, we split the institutional-order sample into a small- and large-order sample to study whether HFTs have an ability to detect large, long-lasting orders early. The results in Table 3 show that this does not seem to be the case. HFTs also lean against large, long-lasting orders initially. Panel A shows that HFTs lean significantly against buy orders in the first hour. They lean against such orders also when they are either larger than $1 million, or when they are larger than median in terms of percentage of average daily volume (ADV). If anything, the against-wind pattern is larger

Table 3: Large, long-lasting orders

This table presents means of HFT net flows and their t-statistics for institutional buy orders and institutional sell orders. These average were calculated in the lifetime of the order, from the start until one hour later, from the start until two hours later, etc. It is similar to Figure 2, but adds to it by showing the pattern for various subsamples based on characteristics of the institutional order. Statistical significance is established based on the t-value of the mean across stock-institution fixed effects (same as overall mean) with residuals clustered at a stock-day level.

	Hours since first child order execution							
	1	2	3	4	5	6	7	8
Panel A: HFT net flow mean in the lifetime of institutional buy orders								
All	-14.0***	-7.7	-16.2**	-16.0*	-11.4	-7.4	29.8**	34.6
t-stat	(3.9)	(1.5)	(2.4)	(1.9)	(1.1)	(0.6)	(2.1)	(1.4)
N	2,516	2,203	1,992	1,803	1,569	1,290	1,010	418
Size ≤ median %ADV	-9.7*	-9.6	-28.6***	-32.4**	-24.1	-7.9	28.0	32.0
t-stat	(1.7)	(1.2)	(2.7)	(2.5)	(1.5)	(0.4)	(1.3)	(0.8)
Size > median %ADV	-17.3***	-6.4	-6.9	-3.4	-2.2	-7.1	30.9*	35.9
t-stat	(3.6)	(1.0)	(0.8)	(0.3)	(0.2)	(0.5)	(1.7)	(1.1)
Size ≤ $1,000,000	-8.9**	-8.9	-24.3***	-30.0***	-26.3**	-23.1	8.5	8.7
t-stat	(2.0)	(1.4)	(3.0)	(2.9)	(2.1)	(1.6)	(0.5)	(0.3)
Size > $1,000,000	-21.4***	-6.1	-4.5	4.8	9.8	14.1	58.1**	60.7
t-stat	(3.5)	(0.7)	(0.4)	(0.3)	(0.6)	(0.7)	(2.4)	(1.5)
Duration < 4 hours	-14.3**	1.6	-16.6					
t-stat	(2.0)	(0.1)	(0.6)					
Duration > 4 hours	-13.9***	-9.8*	-16.1**	-16.0*	-11.4	-7.4	29.8**	34.6
t-stat	(3.3)	(1.8)	(2.3)	(1.9)	(1.1)	(0.6)	(2.1)	(1.4)
Panel B: HFT net flow mean in the lifetime of institutional sell orders								
All	7.0	7.9	-9.5	-19.7	-13.8	-7.2	-58.7**	-134.2**
t-stat	(1.3)	(0.9)	(0.8)	(1.4)	(0.7)	(0.3)	(2.1)	(2.4)
N	1,334	1,110	986	863	740	571	343	161
Size ≤ median %ADV	13.8	18.3	-16.1	-0.4	10.3	56.6	-46.4	13.1
t-stat	(1.4)	(1.2)	(0.7)	(0.0)	(0.3)	(1.0)	(0.7)	(0.1)
Size > median %ADV	4.6	4.5	-7.3	-25.9*	-20.8	-24.9	-61.7*	-176.6***
t-stat	(0.7)	(0.5)	(0.6)	(1.7)	(1.0)	(0.9)	(1.9)	(2.9)
Size ≤ $1,000,000	6.2	12.6	-7.8	-8.8	-5.4	33.0	10.1	3.8
t-stat	(0.8)	(1.0)	(0.4)	(0.4)	(0.2)	(0.8)	(0.3)	(0.0)
Size > $1,000,000	7.5	5.4	-10.4	-25.2	-17.6	-25.1	-86.2**	-191.1***
t-stat	(1.0)	(0.5)	(0.7)	(1.5)	(0.8)	(0.8)	(2.4)	(2.9)
Duration < 4 hours	18.9**	55.4***	41.0					
t-stat	(2.0)	(2.9)	(1.4)					
Duration > 4 hours	0.7	-5.1	-16.1	-19.7	-13.8	-7.2	-58.7**	-134.2**
t-stat	(0.1)	(0.6)	(1.4)	(1.4)	(0.7)	(0.3)	(2.1)	(2.4)

in magnitude for such orders, and more significant.

Panel A further shows that HFTs also lean against long-lasting buy orders in the first hour. Long-lasting is defined as orders with a lifespan of more than four hours (half a trading day). The result is statistically more significant, and equal in magnitude when compared to short-lived orders.

Panel B reveals that HFTs lean against sell orders, large or long-lasting, but the results are statistically insignificant. Notice however that for "full" hours, Figure 2 also shows that HFT net flow is insignificant for sell orders. The figure finds weak significance of the against-wind pattern only for the half-hour and the one-and-a-half-hour time point.

4 Implementation shortfall and HFT gross trading revenue

In this section we investigate the impact of HFT net flow — against-wind or with-wind — on institutional trading cost and HFT gross trading revenue.

4.1 Implementation shortfall and HFT net flow

A standard measure of institutional trading cost is implementation shortfall. It is defined as:

$$ImpShortfall_{ijt} = D_{ijt} \times (\log \overline{P_{ijt}} - \log P_{ijt}^{\text{Start}}), \tag{1}$$

where i indexes institutional investors, j indexes stocks, t indexes days, D_{ijt} is a buy-sell indicator which equals one for institutional buy orders and minus one for institutional sell orders, $\overline{P_{ijt}}$ is the average trade price on the order, and P_{ijt}^{Start} is the midquote price that prevailed at start of the order. Implementation shortfall is expected to be positive on average as buy orders have positive price impact and sell orders have negative price impact that gets multiplied by minus one in the definition. Implementation shortfall is defined in relative terms, but is often expressed in dollar terms by multiplying it with the dollar size of the order.

Before turning to regression analysis, it is useful to inspect whether there is any correlation between the two variables of interest: implementation shortfall and the nature of HFT net flow. To that end, HFT net flow in the lifetime of an institutional *sell* order is first multiplied by minus one.

Figure 4: Implementation shortfall by HFT net flow terciles

This figure plots, for the against-wind, neutral, and with-wind HFT net flow terciles, the average implementation shortfall, the order-size weighted average implementation shortfall, and the average HFT net flow. The terciles are created by first multiplying HFT net flow during institutional sell orders with minus one. All HFT net flow observations are then sorted and put into equal-size bins. The tercile with 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 highest values contains strong with-wind HFT net flows.

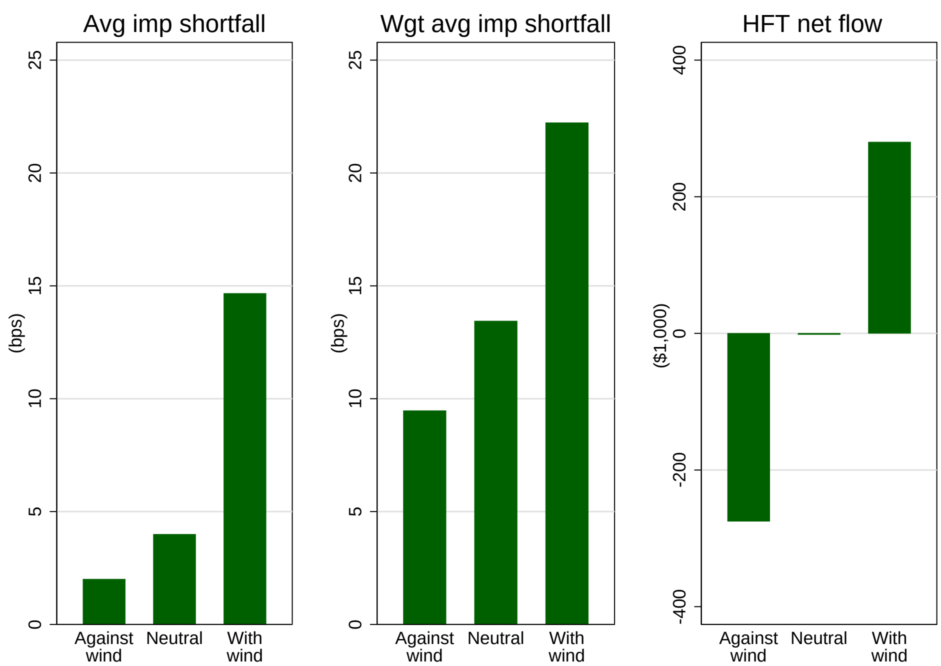

HFT net flows for these orders can then be added to HFT net flows for buy orders in a meaningful way. Negative HFT net flow can then be interpreted as against-wind trading and positive net flow as with-wind trading. These net flows are sorted and binned into equal-size terciles. The average implementation shortfall is then calculated for each tercile.

Figure 4 illustrates that implementation shortfall is lower when HFTs lean against the order, but higher when they go with the order. Implementation shortfall increases monotonically across the three terciles. It is 3.5 basis points for the against-wind tercile, 4.2 for the neutral tercile, and 14.5 basis for the with-wind tercile. The order-size weighted implementation shortfall exhibits a similar pattern. One noteworthy observation is that the increase for with-wind flow is larger in magnitude then the decrease for against-wind flow. This is not simply due to HFT net flow being larger in magnitude for the with-wind bin as the right-most panel in the figure shows that they are equal in magnitude; both are about $275,000. Institutional investors seem to suffer more from with-wind HFT net flow than they benefit from against-wind HFT net flow.

For the regression analysis we separate the two types of HFT net flow by creating the following two variables:

$$AgainstWindHFTNetFlow_{ijt} = 1_{\{\text{sgn}(HFTNetFlow_{ijt}) \neq \text{sgn}(InstOrder_{ijt})\}} \times \left| HFTNetFlow^{*}_{ijt} \right| \tag{2}$$

and

$$WithWindHFTNetFlow_{ijt} = 1_{\{\text{sgn}(HFTNetFlow_{ijt}) = \text{sgn}(InstOrder_{ijt})\}} \times \left| HFTNetFlow^{*}_{ijt} \right|, \tag{3}$$

where 1_A is the indicator function, i.e., it equals one when A is true and zero otherwise, sgn(A) is the sign function, i.e., it is plus one if A is positive, zero if A is zero, and minus one if A is negative, and $HFTNetFlow^{*}$ is standardized HFT net flow.[11] These HFT net flow variables appear on the right-hand side of regressions either in dollar terms or expressed relative to the size of the institutional order.

The following panel regression is run to verify whether the general pattern of Figure 4 holds up when standard control variables are added. The general model specification is

[11]We standardize all right-hand side variables to make coefficients more easily interpretable. The sign of almost none of the HFT net flow observations changes as the overall average HFT net flow is close to zero.

$$ImpShortfall_{ijt} = \alpha_{ij} + \beta_1 AgainstWindHFTNetFlow_{ijt} + \beta_2 WithWindHFTNetFlow_{ijt} + \gamma' X_{ijt} + \varepsilon_{ijt}, \quad (4)$$

where α_{ij} is shorthand notation for the addition of an institution fixed effect and a stock fixed effect, X_{ijt} is a vector with control variables, and ε_{ijt} is a residual that is allowed to exhibit correlation within a stock-day, but not across stock-days (standard errors are clustered at the stock-day level). We use two model specifications in the regressions. The first specification expresses both implementation shortfall and HFT net flow relative to the size of the order. The second specification expresses both variables in dollar terms.

The control variables are standard market-condition variables and order-specific variables (see, e.g., Anand, Irvine, Puckett, and Venkataraman, 2012). In particular, we add the size of the institutional order expressed relative to average daily volume (ADV), the duration of the order, realized volatility, and stock volume.

Table 4 presents the regression results. We observe that against-wind HFT flow reduces implementation shortfall significantly, both for relative shortfall and the dollar shortfall specifications. A one standard-deviation against-wind HFT flow reduces *relative* implementation shortfall by 2.9 basis points, a reduction of 39% relative to the sample average of 7.4. The coefficient of the *dollar* implementation shortfall is insignificant, but the point estimate is of similar magnitude.

With-wind HFT flow increases implementation shortfall significantly, both for relative shortfall and dollar shortfall. A one standard-deviation HFT net flow increases *relative* implementation shortfall by 4.7 basis points, an increase of 64% relative to the sample average. A one standard-deviation HFT net flow increases *dollar* implementation shortfall by $2,965, an increase of 104% relative to the sample average of $2,860. The marginal effect of with-wind seems larger than that of against-wind flow, but we can only reject equality for the dollar shortfall specification with a p-value of 0.01.

The statistically significant control variables carry the expected signs. The statistically strongest covariate is order size as a percentage of average daily volume. A one standard-deviation increase raises shortfall by 6.5 basis points or $5,207. In the dollar-shortfall models, the duration of the

Table 4: Implementation shortfall regressed on HFT net flow and control variables

This table presents panel regression results where implementation shortfall is the dependent variable. The main explanatory variable is HFT net flow measured over the lifetime of the institutional order. It is first standardized, then the absolute value is taken, and finally it is labeled "with-wind" if the standardized value had the same sign as the institutional order and "against-wind" otherwise. Various variables are added as controls. All these variables have been standardized. ADV is average daily volume based on the full sample. Order size and ADV are measured in shares. Stock volume and volatility are measured from the start to the end of the order. Volatility is measured as realized volatility based on one-minute midquote returns. Also reported are p-values of (i) a test whether the against-wind coefficient equals minus the with-wind coefficient and (ii) whether indeed the model coefficients are equal for institutional buy and sell orders (this table shows the pooled regression results). The 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 and reported right after variable names. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

	Implementation shortfall (bps)	Implementation shortfall ($1,000)
Against-wind HFT net flow (%)	-2.946***	
	(-3.2)	
With-wind HFT net flow (%)	4.660***	
	(4.0)	
Against-wind HFT net flow ($)		-0.621
		(-1.4)
With-wind HFT net flow ($)		2.965***
		(4.7)
Order size relative to ADV (%)	6.446***	5.207***
	(6.9)	(9.5)
Order duration (hours)	-0.658	-0.651*
	(-0.6)	(-1.6)
Stock volatility (%)	-0.554	-0.056
	(-0.7)	(-0.3)
Stock volume ($)	1.691	0.748
	(1.1)	(1.0)
Number of observations	5,910	5,910
R^2	0.021	0.114
p-val "*With-wind* $= -1 \times$ *Against-wind*"	0.28	0.01
p-val "*Buys* = *Sells*"	0.35	0.38

order also significantly affects implementation shortfall. A three-hour longer order (one standard-deviation) reduces shortfall by \$651, a reduction of 23%. All else equal, spreading the order over a longer horizon reduces execution costs. The coefficients on volatility and stock volume are insignificant.

Finally, we test whether the regression coefficients differ between institutional buys and sells, i.e., whether we are allowed to pool institutional buy and sell orders in one regression (as we did). The p-values in the table reveal that, indeed, the null of equal coefficients cannot be rejected. p-values are 0.35 and 0.38 for the "relative" and the "dollar" model, respectively.

4.2 HFT gross trading revenue and HFT net flow

The previous section showed that HFT net flow strongly affects institutional trading cost. Do HFTs make money off of such behavior and, if so, how much? To this end, we repeat the regressions of the previous section, but use HFT gross trading revenue (GTR) as a dependent variable to proxy for their gross profit.

HFT gross trading revenue (GTR) over the lifetime of the order is calculated as in Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010). It is a simple accounting exercise where HFTs start off with no position in the stock and zero cash, they buy and sell the stock along the way (where they can borrow at zero cost), and at the end time any nonzero position in the stock is converted into cash using the end-of-day stock price. HFT GTR is defined as the cash position that they then have at the end of the period. GTR is measured in dollars, but it can also be expressed relative to total amount they traded in the stock. It then indicates how much they make on each dollar they trade. Both dollar GTR and relative GTR will be used in the regressions (in parallel to what was done for the implementation shortfall regressions).

Table 5 presents the regression results. Relative GTR does not seem affected by HFT net flow, but dollar GTR is significantly higher the more HFTs engage in either with-wind or against-wind trading. Taken together, it seems that the institutional order gives them additional trading opportunities, i.e., they trade more. These trading opportunities, however, are not more profitable to them in terms of the margin they make on each dollar they trade.

Table 5: HFT gross trading revenue regressed on HFT net flow and control variables

This table presents panel regression results where the dependent variable is HFT gross trading revenue in dollar terms or expressed relative to HFT dollar volume. The main explanatory variable is HFT net flow measured over the lifetime of the institutional order. It is first standardized, then the absolute value is taken, and finally it is labeled "with-wind" if the standardized value had the same sign as the institutional order and "against-wind" otherwise. Various variables are added as controls. All these variables have been standardized. ADV is average daily volume based on the full sample. Order size and ADV are measured in shares. Stock volume and volatility are measured from the start to the end of the order. Volatility is measured as realized volatility based on one-minute midquote returns. Also reported are p-values of (i) a test whether the against-wind coefficient equals minus the with-wind coefficient and (ii) whether indeed the model coefficients are equal for institutional buy and sell orders (this table shows the pooled regression results). The 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 and reported right after variable names. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

	HFT gross trading revenue ratio (bps)	HFT gross trading revenue ($1,000)
Against-wind HFT net flow (%)	-0.389	
	(-0.7)	
With-wind HFT net flow (%)	-0.257	
	(-0.7)	
Against-wind HFT net flow ($)		0.701**
		(2.4)
With-wind HFT net flow ($)		1.164***
		(3.2)
Order size relative to ADV (%)	0.019	-0.046
	(0.1)	(-0.4)
Order duration (hours)	0.232	-0.189
	(0.6)	(-1.4)
Stock volatility (%)	-0.108	-0.144
	(-0.7)	(-1.0)
Stock volume ($)	0.103	0.597***
	(0.4)	(2.8)
Number of observations	5,910	5,910
R^2	0.000	0.034
p-val "*With-wind* = $-1 \times$ *Against-wind*"	0.37	0.00
p-val "*Buys = Sells*"	0.04	0.18

Figure 5: Marginal impact of HFT net flow on implementation shortfall and HFT revenue

This figure plots the marginal impact of HFT net flow on implementation shortfall and HFT gross trading revenue. The marginal impact is estimated in Table 4 and 5. The impact therefore controls for standard covariates, e.g., order size, duration, volume, and volatility. The domain spans the interval from minus to plus one standard-deviation of HFT net flow.

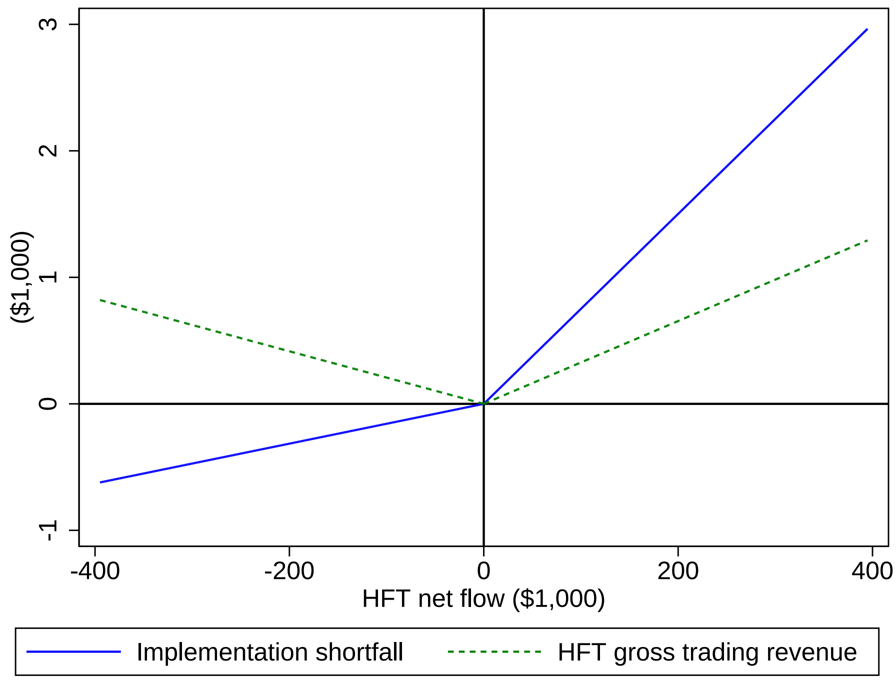

The coefficient is largest for with-wind flow. A one-standard-deviation HFT with-wind flow raises HFT gross profit by $1,164, which is an increase of 320% relative to the sample average of $364. This increase is large, also in terms of the additional cost institutional investors incur in such case, i.e., it is $100\% * (\$1,164/\$2,965) = 39\%$ of their cost increase (see Section 4.1). The coefficient for against-wind flow is $701 per standard deviation. This coefficient is 40% smaller than the with-wind effect, but the difference is not significant.

We further find that dollar GTR is larger at times of more overall volume, all else equal. Relative GTR does not seem to depend on volume. Again, more overall volume enables them to make more money, but they do not earn more on a per dollar-traded basis.

Summary. Figure 5 summarizes this paper's main results. It plots the estimated impact of HFT net flow on (i) implementation shortfall of an institutional buy order and (ii) on HFT gross trading revenue. The estimated impact is *ceteris paribus* because it is taken from panel regressions that included standard control variables (see the model specification in (4)). We observe that HFT against-wind flow lowers institutional trading cost whereas HFT with-wind flow raises it. The with-wind effect is substantially stronger. HFTs benefit either way, but their trading revenue is slightly for with-wind trading. The plots further show that changes in HFT revenue are smaller in magnitude than changes in institutional trading cost.

5 Conclusion

This paper is the first to document how trading by HFTs affects the trading cost of end-users. The latter are known to execute their large orders through a series of small "child" order transactions. A sample of order executions by four large institutional investors was studied for Swedish stocks in 2011-2013. We find that HFTs initially lean against an order (trade in the direction opposite to it), but if the order lasts more than a couple of hours then they turn around and go with the order. HFT gross profit is positive either way. Institutional investors' cost is lower for against-wind HFT net flow, but disproportionately larger for with-wind HFT flow.

The results are inconsistent with "front-running" in the sense of HFTs who detect a large, long-lasting order right from the start and trade along with it (Brunnermeier and Pedersen, 2005). Rather, we speculate that HFTs eventually feel the imbalance caused by it. In response, they trade out of their position as they understand that leaning against such order as a market maker requires a long-lasting inventory position. HFTs prefer to be flat at the end of the day.

The finding that HFTs not only unwind their position when they detect a long-lasting order but decide to trade along with it is somewhat surprising. One would only do so if the order is an informed one. For an institutional buy order, for example, one wants to join on the way up when "late to the party" only if the price stays at the high level, i.e., the order is informed.[12]

[12]This appears true for our sample as the idiosyncratic price impact tends to be permanent as judged by where the price is not only at the end of the day, but also one day later.

An alternative explanation is that both institutional investors and HFTs received the same private information, with the institutional investor trading on it hours before HFTs do. We consider the latter explanation less likely given HFTs' edge in information technology.

We believe the market structure debate should recenter around 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 (as was done for this study, for example). Alternatively, regulators might demand more data granularity from data centers, 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 did, brokers had to report whether it was a principal or an agency trade and, if agency, 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.[13]

References

Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2012, "Performance of Institutional Trading Desks: An Analysis of Persistence in Trading Costs," *Review of Financial Studies*, 25(2), 557–598.

——— , 2013, "Institutional Trading and Stock Resiliency: Evidence From the 2007-2009 Financial Crisis," *Journal of Financial Economics*, 108(3), 773–797.

Boehmer, E., K. Fong, and J. Wu, 2014, "International Evidence on Algorithmic Trading," *Working Paper EDHEC Business School*.

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.

[13] See https://www.sec.gov/divisions/marketreg/rule613-info.htm.

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.

Chordia, T., and A. Subrahmanyam, 2004, "Order Imbalance and Individual Stock Returns: Theory and Evidence," *Journal of Financial Economics*, 72(1), 485–518.

Comerton-Forde, C., T. Hendershott, C. M. Jones, P. C. Moulton, and M. S. Seasholes, 2010, "Time Variation in Liquidity: The Role of Market-Maker Inventories and Revenues," *Journal of Finance*, 65(1), 295–331.

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.

Gao, C., and B. Mizrach, 2013, "Market Quality Breakdowns in Equities," *Working Paper Rutgers University*.

Hagströmer, B., and L. Nordèn, 2013, "The Diversity of High-Frequency Traders," *Journal of Financial Markets*, 16(4), 741–770.

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*.

Jones, C. M., 2014, "What Do We Know About High-Frequency Trading?," *Columbia Business School Research Paper No. 13-11*.

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, 2014, "High Frequency Market Making to Large Institutional Trades," *Working Paper Northwestern University*.

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.

Menkveld, A. J., and B. Z. Yueshen, 2015, "Anatomy of the Flash Crash," *Working Paper VU University Amsterdam*.

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.

Ye, M., C. Yao, and J. Gai, 2013, "The Externalities of High Frequency Trading," *Working Paper University of Illinois at Urbana-Champaign*.