Journal Of Investment Management, Vol. 18, No. 1, (2020), pp. 18–33

www.joim.com

DO HIGH-FREQUENCY TRADERS IMPROVE YOUR IMPLEMENTATION SHORTFALL?*

Robert A. Korajczyk[a] *and Dermot Murphy*[b,1]

We take advantage of a regulatory change that effectively imposed a "tax" on HFT order activity on Canadian equity venues to study the resulting effect on the execution costs of large institutional trades. We find that bid–ask spreads increase and price impact decreases for these trades following the regulatory change. The price impact effect is strongest for informed institutional traders. Our evidence indicates that this tax on high-frequency trading is associated with higher transaction costs for small, uninformed trades and lower transaction costs for large, informed trades. Hence, the tax increased the subsidy for informed traders from uninformed traders.

1 Introduction

High-frequency traders (HFTs) have largely assumed the market making role in modern equity markets. Extant research has shown that the presence of HFTs is associated with improvements in market quality, including lower bid–ask spreads and improved price efficiency (e.g., Menkveld, 2013; Brogaard *et al.*, 2014). However, it is still unclear whether lower spreads, clearly beneficial to small traders, lead to reduced transaction costs for large institutional traders. Institutional traders frequently buy or sell millions of dollars' worth of shares in short periods of time and their trading costs are potentially adversely influenced by HFTs' ability to quickly re-price standing limit orders and trade in competition with institutional orders. A number of proposals have surfaced to curb HFT activity, including taxes on message traffic, the replacement of continuous trading with batch trading, and increasing the latency of order and trade data feeds. In this paper, we study how one such proposal, which involves increasing the fees on HFT message traffic, affects HFT behavior and the execution costs of large institutional trades.

*This manuscript is based on Korajczyk and Murphy (2019), which was originally published in the Review of Financial Studies by Oxford University Press.
[a]Northwestern University, 2211 Campus Dr. #4357, Evanston, IL 60208-0898, USA. Tel.: (847) 491-8336; E-mail: r-korajczyk@kellogg.northwestern.edu.
[b]University of Illinois at Chicago, College of Business, Department of Finance, 601 South Morgan Street, Chicago, IL, 60607-7121, USA; Tel.: (312) 355-4372; E-mail: murphyd@uic.edu.
[1]Corresponding author.

A standard feature in the market making literature is that equilibrium quotes are set by competitive market makers such that uninformed investors, those with non-informational reasons for trading, subsidize informed investors (e.g., Glosten and Milgrom, 1985; Kyle, 1985). That is, market makers lose to informed traders and make up those losses from uninformed traders. The speed advantage conferred by market making HFTs might allow them to detect informed traders and adjust quotes in a way that allows them to reduce their exposure to adverse selection, thus increasing the trading costs for informed traders. This possibility is highlighted by Warren Buffett, Chairman and CEO of Berkshire Hathaway, who argues that the speed advantages gained by HFTs have made the "big orders" more costly while acknowledging that small investors have "never had it so good" (Crippen, 2014).

In theory, there are three primary mechanisms through which market-making HFTs might increase the execution costs of large institutional trades. The first mechanism is related to HFTs using their speed advantage to re-price standing limit orders in response to signals about incoming orders. This behavior is sometimes described as "phantom liquidity," where limit orders are modified before slower traders are able to complete their trades. The second mechanism is related to "back-running," where HFTs attempt to infer information from order flow to trade in competition with an institutional order. The third mechanism is based on classic microstructure models of inventory management (e.g., Ho and Stoll, 1981). A large institutional order will increase the magnitudes of market-making HFT inventory positions. If HFTs trade to reduce their inventory positions while the institutional trade is underway, the HFTs effectively compete with the institutional order. Although market maker inventory management is not new to the HFT environment, HFTs may use their speed and information processing advantages to reduce their inventories faster through the use of order modifications and trades that are in the same direction as the institutional order.

We study the implementation shortfalls of a sample of approximately 1.2 million institutional-sized directional trades on Canadian equities exchanges. The Investment Industry Regulatory Organization of Canada (IIROC) provides us with access to order-level data for all Canadian equities for the period from January 2012 to June 2013. For each of the approximately 60 billion messages, we are provided with a user ID, allowing us to track the order and trade activities for any of these IDs across time and in the cross-section of equities. In particular for this study, the information allows us to identify both HFTs and directional institutional orders.

On April 1, 2012, IIROC introduced the "integrated fee model," which changed the fee structure for traders in Canadian equities markets. In particular, traders would now be charged on a pro-rata basis for their number of executed trades and messages (i.e., orders, order modifications, and order cancellations) sent to Canadian marketplaces (see IIROC Notice 12-0043). Prior to this regulation, traders were only charged fees for their executed trades. HFTs were especially affected by this regulatory change because their strategies typically involve very high message traffic.

We find that daily HFT message traffic decreased approximately 20% following the fee structure change, which is consistent with the evidence in Malinova *et al.* (2018). We show that the average price impact—the component of implementation shortfall related to trade size—for large institutional trades decreases by 15% following the regulatory change, suggesting that the ability of HFTs to adjust their limit orders in the presence of large institutional trades was somewhat hindered by this new regulation. We also provide evidence that the average fixed cost for these trades increased by 3.0 basis points, suggesting that HFTs widened their bid–ask spreads to compensate for the new order submission fees. Taken together, this evidence indicates that execution costs increased for smaller-sized institutional trades and decreased for larger-sized institutional trades. We find a trade size break-even point of approximately $2 million.

We rank institutional traders by the profitability of their past trades and examine the differential effect of the regulatory change on transaction costs across profitability terciles. We confirm that traders from the "high-information group" (those with the largest average past profits) trade profitably out-of-sample, indicating that these traders exhibit some degree of skill. We find that the regulatory change has the strongest effect on the high-informed group, with a reduction in price impact of about 28%. In contrast, the price impact reductions for the remaining types are low and have weaker statistical significance. Thus, our results suggest that trading costs for informed traders are higher when HFTs can cheaply modify their orders and lower when the fee structure slows HFT activity.

Malinova *et al.* (2018) study the effect of the integrated fee model on the execution costs of both retail and institutional traders. They find that the effective spread for retail traders increases significantly by 0.9 basis points. They also find that implementation shortfall significantly increases by 4.9 basis points for institutional parent orders that only use marketable orders. For institutional orders that use both marketable and nonmarketable orders, implementation shortfall increases by a statistically insignificant 1.2 basis points. When we study average implementation shortfall, unconditional on trade size, we obtain results that are in line with their findings. We break our implementation shortfall into a spread component and a price impact component and find the former increases and the latter decreases with the integrated fee model. While the average trade has higher implementation shortfall, we find that the largest trades have lower implementation shortfall.

2 Data and classification methodology

We have been provided access to detailed order-level data by IIROC, a Canadian national self-regulatory organization that regulates securities dealers in Canada's equity markets.[1] Through the monitoring of the Canadian equities markets, IIROC collects detailed records on all orders submitted to Canadian exchanges. IIROC provides us with access to a data set that contains all trades, orders, order cancellations, and order modifications for the period from January 1, 2012 to June 30, 2013. Each record contains an "event" field that allows us to determine whether that observation is an order, trade, order cancellation, or order modification. We are also provided with the security ID and the price, quantity, date, and time associated with each record, where the time is reported at the millisecond level. Each record contains a masked identification for the trader submitting an order, allowing us to track the activity of any user ID over time, along with the direction (buy or sell) of that order. For every trade, we are provided with an "active/passive" indicator that identifies the party submitting the marketable limit order. In other words, this indicator tells us which side of the trade is "passive" (the standing limit order) and which side is "active" (the marketable order). Altogether, the data set comprises approximately 60 billion observations.

The IIROC data set contains a high level of detail that allows us to classify traders as market-making HFTs and identify large, institutional-sized trades.[2] The initial sample of stocks used for these identifications is based on IIROC's definition of "highly liquid securities," which are stocks that have an average of at least 100 trades and $10 million in dollar volume per day on Canadian marketplaces. Low-liquidity securities are mostly traded in markets that have very little HFT activity and are not applicable to our study. Of the approximately 4,200 publicly traded equity securities traded during the sample period, 295 stocks meet the criteria for highly liquid securities. This represents 7% of the publicly traded equity securities. On a dollar volume basis, the highly liquid securities comprise approximately $5 trillion during the sample period, whereas the remaining publicly traded equity securities comprise approximately $700 billion. Therefore, 88% of the dollar volume on Canadian exchanges is in the highly liquid securities. All dollar figures in this paper are reported in Canadian dollars. However, the exchange rate between U.S. and Canadian dollars was close to parity during our sample period.

Canadian market structure is fairly similar to U.S. market structure, and thus the findings of this study should be applicable to U.S. markets as well. In particular, the main exchanges in Canada and the U.S. are both completely electronic stock exchanges in which orders are submitted to their respective limit order books. Canadian investors also have access to a central source of consolidated Canadian equity market data via the TMX Information Processor (TMX IP), which is similar to the Securities Information Processor (SIP) in the U.S.[3] Both Canadian and U.S. markets have dark pools in addition to "lit" markets, although the level of dark pool trading in higher in the U.S. versus Canada (Devani *et al.*, 2015; CFA Institute, 2012). The "maker-taker" fee structure, in which traders are charged a per-share fee for taking liquidity from an exchange and receive a per-share rebate for providing liquidity to an exchange, is also common in both countries. (During our sample period, the TSX, the highest-volume exchange in Canada, had a taker fee of $0.0035 per share, which was slightly higher than the maximum possible taker fee in the U.S. of $0.0030 per share.) The Order Protection Rule in Canada is similar to Regulation NMS in the U.S. in that both regulations are designed to ensure that investors' marketable limit orders are executed at the best available prices and thus do not "trade through" orders at better prices. However, the regulations slightly differ about which orders to protect: while the Order Protection Rule is designed to prevent trade-throughs of visible limit orders at all price levels on an exchange, Regulation NMS is designed to only prevent trade-throughs of visible, best-priced limit orders on an exchange. Finally, the Universal Market Integrity Rules (UMIR) in Canada generally require that a market participant acting as a principal or agent for an order must submit that order to a marketplace. As a result, the internalization of order flow in Canada is less common compared to the U.S. Overall, our findings indicate that while there are some differences across Canadian and U.S. equity exchanges, the basic market structures are similar.

2.1 Classifying market-making HFTs

Following the methodologies of Comerton-Forde *et al.* (2018) and Malinova *et al.* (2018), we classify HFTs based on their operating speeds. For each user ID in our sample, we calculate the median time between submitting an order and canceling it. Neuroscience research suggests that the median reaction time of humans to external stimuli is approximately 250 milliseconds (Laming, 1968); we primarily classify a user ID as an HFT if the median order-to-cancel time is below this 250-millisecond threshold. We also follow these studies by classifying a user ID as an HFT if the corresponding trader submitted at least

1,000 orders within the first 500 milliseconds after 3:40 p.m. EST; this is when the Toronto Stock Exchange (TSX) makes announcements about net order imbalances for closing call auctions in TSX-listed securities. These announcements contain valuable information about closing prices in TSX-listed stocks that fast traders can exploit by trading before other traders. Thus, accounts that consistently submit many orders immediately following information releases at 3:40 p.m. EST are also classified as HFTs. Using these classification schemes, we identify 103 user IDs as HFTs. This number is close to the number of HFTs in Comerton-Forde *et al.* (2018) and Devani *et al.* (2014).

A number of the stocks in our sample barely make the cutoffs for highly liquid securities. For many of these stocks, there is little to no HFT presence, making them unsuitable for our analysis of the effects of HFTs on large institutional trades. Using a sample of stocks provided by NASDAQ, Brogaard *et al.* (2014) report that HFTs represent 42% of volume in large stocks and 18% in small stocks. Motivated by this, we exclude any stock from our sample if HFTs represent less than 15% of trading volume in that stock. Our results are not sensitive to alternative cutoffs of 5% and 25%. This leaves us with a final sample of 181 stocks. This sample represents approximately 77% of the total dollar volume of the publicly traded equity securities during the sample period.

Of our 103 accounts identified as HFTs, we identify 68 HFTs who act as market makers based on the "market maker index" (MMI) of Comerton-Forde *et al.* (2018). The MMI is the absolute value of the difference between an HFT's passive buy volume and passive sell volume, as a fraction of their total passive volume. By construction,

Table 1 Summary statistics for high-frequency trading activity.

	HFT Summary Statistics ($N = 67{,}787$)						
	Mean	Median	P5	P25	P75	P95	SD
Percentage of trade volume (%)	31.6	30.8	11.5	22.0	40.6	53.4	13.1
Percentage of orders (%)	55.4	56.0	21.0	41.3	69.2	85.9	22.9
Order-to-trade ratio	33.1	16.9	5.4	10.5	32.7	119.8	49.5
Aggressiveness (%)	27.8	26.9	7.8	18.2	36.2	50.9	13.3
Trade size (shares)	328	147	111	125	260	1,261	531
Trade value (dollars)	4,354	2,685	459	1,092	5,531	12,133	6,095
Inventory ($K)	3.7	1.3	−105.5	−16.9	23.6	119.4	72.9
Inventory (%)	2.5	0.2	−49.8	−3.3	5.8	63.6	52.1
Δ Inventory ($K)	0.0	0.0	−55.6	−7.5	7.5	55.8	48.7
Δ Inventory (%)	0.0	0.0	−100.0	−19.4	18.9	100.0	46.9

Percentage of trade volume is the ratio of HFT trading volume to total trading volume (double counted) for each stock-day. *Percentage of orders* is the ratio of total HFT limit-order size to total limit-order size for each stock-day. *Order-to-trade ratio* is the ratio of the number of HFT limit orders to the number of HFT trades. *Aggressiveness* is the percentage of HFT trades that are executed using marketable limit orders, as opposed to passive limit orders. *Trade size* is the number of shares composing a single HFT trade. *Trade value* is the dollar value of the shares composing a single HFT trade. *Inventory* ($\$K$) is the dollar inventory of the highest-volume market-making HFT in a stock and 15-minute period, and Δ *Inventory* ($\$K$) is the change in this variable across 15-minute periods. We also report summary statistics for these variables expressed as a percentage of dollar volume by the HFT in that stock and 15-minute period. For each variable, "PX" represents the Xth percentile of its distribution, and "SD" represents its standard deviation.

the MMI is bound between zero and one. For each HFT, we calculate the median value of MMI across all stocks and days in which that HFT is active. A median MMI close to zero indicates that the HFT consistently submits similar share quantities of buy and sell orders, suggesting that the HFT is a market maker. We find a structural break in the median MMI at 0.20, and classify an HFT as a market-making HFT if their median MMI is below this threshold.

Using all stock-days in our sample, we calculate summary statistics for our 68 market-making HFTs. These statistics are reported in Table 1. We find that market-making HFTs are responsible for 31.6% of total daily trading volume, on average. For 5% of the stock-days in our sample, market-making HFTs are responsible for at least 53.4% of daily trading volume. These numbers suggest that market-making HFTs have a significant presence in our sample of stocks. Thus, our sample of large, directional institutional traders will be trading with market-making HFTs fairly frequently. Furthermore, we find that market-making HFTs are responsible for 55.4% of total limit-order submission volume, on average, which is almost twice as high as their average percentage of trading volume. This is unsurprising: frequent limit-order submissions, cancellations, and modifications are hallmarks of HFT market-making strategies because they frequently adjust their limit orders to avoid the adverse-selection risk associated with trading with informed counterparties. The high average order-to-trade ratio of 33.1 in Table 1 further supports our claim that market-making HFTs frequently cancel their limit orders, and implies that market-making HFTs only execute about 3% of the orders that they post to the limit-order book. We also find that market-making HFTs execute approximately 27.8% of their trades using marketable limit orders, suggesting that liquidity provision is not the only component of their market-making strategies. Finally, we find that the median size for a trade involving a market-making HFT is 147 shares, with a median value of $2,685.

2.2 *Classifying large institutional trades*

To minimize execution costs, institutions commonly execute their large orders (parent orders) by executing a series of smaller-sized orders (child orders) over time. Doing so allows the institution to "hide" in the order flow of other traders, reducing the ability of other traders to detect the information content contained in their order flow. We classify a series of orders as a "large institutional trade" if the same user ID is on the same side of one or more transactions with a total parent order dollar volume of $100,000 or greater.[4] The $100,000 cutoff is based on institutional trade dollar size statistics reported in other studies. For example, Cready *et al.* (2014) report average institutional trade sizes that range from approximately $100,000 to $500,000, depending on the size of the institutional investor. The $100,000 cutoff used in our study is conservative and allows for a greater cross-section of institutional investors. Using this methodology, we identify 1,173,482 large institutional trades in our sample.

In addition to dollar trade size, we identify other important attributes of each institutional trade that could potentially influence the dynamics of how HFTs interact with that trade. "Aggressiveness" is defined as the percentage of the institutional trade that is executed using marketable orders, which are typically used by liquidity-demanding traders wishing to execute a trade quickly. "Time to Completion" is defined as the number of trading hours it takes to execute the large trade; a lower value is likely to reflect an institutional trade that is based on short-lived private information. Finally, we calculate the implementation shortfall (*IS*) of each institutional trade. This is based on the "implicit cost of interacting with the market" from Perold (1988) and calculated as follows:

$$IS_{i,t} = \frac{\sum_{n=1}^{N} p_n x_{i,n} - p_0 x_{i,N}}{p_0 x_{i,N}} \times (\mathbf{1}_B - \mathbf{1}_S), \tag{1}$$

where i represents the institutional parent order, t represents the date that the trade was initiated, p_n and x_n are the price and volume of child trade n within the parent order, p_0 is the bid–ask midpoint at the initiation of the parent order, $x_N \equiv \sum_{n=1}^{N} x_n$ is the total number of shares executed in the parent order, and $\mathbf{1}_B$ ($\mathbf{1}_S$) is an indicator variable that equals one if the institutional trade is a buy (sell). This is a standard approach used in other studies of institutional trading, such as Keim and Madhavan (1997) and Anand *et al.* (2012). We analyze the effect of the integrated fee model on the implementation shortfalls of large institutional parent orders.

Summary statistics for the sample of large institutional trades are reported in Table 2. The average dollar size of a large institutional trade is \$720,000, with the upper 5% (1%) of these trades exceeding \$2.5 million (\$6.8 million). On average, an institutional trade is executed using 118 smaller trades and 234 limit orders. The discrepancy between the average number of trades and limit orders is the result of unexecuted limit orders that are canceled by the institution. The average and median order-to-trade ratios for an institutional trade are 4.9 and 1.0, respectively, further reflecting the use of order cancellations by some institutions. About 57% of the shares in an institutional order are bought or sold using marketable limit orders, indicating a 57/43 split between liquidity-demanding and liquidity-supplying limit orders. It takes an average of 3.0 hours to execute an institutional trade, and 95% of all institutional trades are executed within 6.5 hours. We find that 6.3% of the institutional trades span 2 or more days. Finally, we find that the average implementation shortfall of a large institutional trade equals 7.1 basis points. The large interquartile range of 31.8 basis points further suggests sizable variation in execution quality. We show below that some of this variation is based

Table 2 Summary statistics for large institutional trades.

	Institutional Trade Statistics ($N = 1{,}173{,}482$)						
	Mean	Median	P5	P25	P75	P95	SD
Trade size (\$M)	0.72	0.28	0.11	0.16	0.64	2.53	1.91
Number of orders	234	48	1	11	178	855	2,207
Number of trades	118	50	3	20	124	438	261
Order-to-trade ratio	4.9	1.0	0.1	0.4	1.8	6.4	36.2
Aggressiveness (%)	57.0	61.1	0.0	22.1	96.5	100.0	36.6
Time to completion (hours)	3.0	1.7	0.0	0.1	5.3	6.5	4.0
Implementation shortfall (bps)	7.1	2.5	−97.9	−8.8	23.0	119.3	81.9

Trade size is the dollar value of the institutional trade, measured in millions of dollars. *Number of orders* is the total number of orders submitted by the institution during the execution of the institutional trade. *Number of trades* is the number of smaller trades it takes to execute the large institutional trade. *Aggressiveness* is the percentage of the large institutional trade that is executed using marketable limit orders. *Time to completion* is the number of hours it takes to execute the institutional trade. *Implementation shortfall* is defined in the text and measured in basis points. For each variable, "PX" represents the Xth percentile of its distribution, and "SD" represents its standard deviation.

on the chance of the market moving in the same or opposite direction of the parent order.

3 Institutional trading costs and HFT

The first major step in this study is to establish a link between institutional trading costs and high-frequency trading activity. Theory suggests that HFTs can use their speed advantage to profitably modify their orders in the presence of a large institutional trade, thereby increasing the execution costs associated with that trade. An exogenous event that improves or impairs the ability of HFTs to modify their orders in the presence of large institutional trades would be ideal for establishing a link between HFT and institutional execution costs. We utilize the regulatory change that was implemented by IIROC on April 1, 2012, called the "integrated fee model." Prior to the integrated fee model, IIROC imposed a fee on its members based on the volume of shares traded. IIROC recognized that message traffic from Canadian exchanges was steadily increasing over time, increasing the burden on IIROC to monitor the traders on these exchanges. As a result, IIROC implemented a fee model in which traders would be charged for both the number of trades they execute and the number of messages they send to Canadian marketplaces. In developing the new fee model, IIROC stated that traders with a "greater share of messages or trades compared to their share of shares traded will incur higher fees under the proposed model compared to the current model" (IIROC Notice 10-0316). High-frequency traders were strongly affected by the new fee structure because their strategies typically involve high message activity and high order-to-trade ratios, as shown in Table 1. In contrast, the new message fees for our sample of institutional traders were largely trivial due to their low message activity relative to trades, and we do not observe any significant changes in the arrival rates of institutional orders after the regulation. Pro-HFT commenters noted that the proposed model would extend "an apparent bias against HFTs," further noting that "taxing message traffic will disproportionately hurt HFTs." In their responses, IIROC stated that they "developed the proposed fee model to be as neutral as possible between liquidity providers and liquidity takers."[5] The goal of this section is to provide a clearer picture of how this regulatory change affected the execution quality of the institutional trades in our sample.

As a preliminary test, we examine HFT trade, order, and cancellation statistics surrounding the fee change and do indeed see significant changes. For each stock, we calculate the average daily number of trades, orders, and cancellations during the 3-month periods immediately before and after the fee change, as well as the percentage change in these averages after the fee change. Table 3 reports these summary statistics. We find that the daily number of trades by HFTs declined by 14.7%, while the daily number of orders and cancellations by HFTs declined by 21.4% and 21.6%, respectively. Thus, the fee change seems to have had a significant effect on HFT behavior.

3.1 Baseline results

We examine the effect of the integrated fee model on the execution costs of large institutional trades by estimating the change in the fixed and trade-size related components of implementation shortfall through the following regression:

$$\begin{aligned} IS_{i,j,t} = \beta_1 \cdot \ln(TSize_{i,j,t}) + \beta_2 \cdot Fee_t \\ + \beta_3 \cdot (Fee_t \times \ln(TSize_{i,j,t})) \\ + \gamma \cdot X_{i,j,t} + \delta_j + \epsilon_{i,j,t}. \end{aligned} \tag{2}$$

In this specification, i represents the institutional trade, j represents the stock being traded, t represents the date that the trade was initiated, and δ_j denotes stock fixed effects. *Fee* is an indicator

Table 3 HFT activity statistics around regulatory fee changes.

	Pre-Regulation	Post-Regulation	Percentage Change (%)
A. Daily number of HFT trades			
Mean	5,220	4,451	-14.7^{***}
25th percentile	1,417	1,114	-21.4
Median	3,089	2,386	-22.8
75th percentile	6,586	5,587	-15.2
B. Daily number of HFT orders			
Mean	116,783	91,778	-21.4^{***}
25th percentile	29,463	21,444	-27.2
Median	60,590	50,989	-15.8
75th percentile	161,291	137,355	-14.8
C. Daily number of HFT cancellations			
Mean	112,611	88,250	-21.6^{***}
25th percentile	26,929	20,317	-24.6
Median	57,522	49,182	-14.4
75th percentile	156,566	125,434	-19.9

This table reports cross-sectional summary statistics for the average daily number of trades, orders, and cancellations submitted by all HFTs in each stock during the 3-month periods surrounding the integrated fee model. Cross-sectional statistics are also reported for the percentage changes in HFT trades, orders, and cancellations around the regulatory change. For the mean percentage changes, ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

variable that equals one if the trading day is on or after April 1, 2012, the date that the integrated fee model went into effect, and zero otherwise. $\ln(TSize)$ is the natural log of the standardized dollar value of the institutional order, calculated as the number of shares in the institutional order multiplied by the bid–ask midpoint at the initiation of the order and then divided by the trade size minimum of \$100,000. The β_2 coefficient on *Fee* represents the change in the average fixed cost (which we refer to as the "spread") for institutional trades following the fee change, while the β_3 coefficient on $Fee \times \ln(TSize)$ represents the change in the average price impact for institutional trades following the fee change. If the fee change makes it costly for HFTs to quickly modify orders in response to institutional order flow, then we would expect a negative β_3 and a positive β_2 coefficient. Finally, X is a vector of the following market-level control variables during the execution of the large institutional trade: (1) *Mret*, the contemporaneous S&P/TSX 60 market return multiplied by the direction of the large institutional trade, which is meant to control for the fact that the average execution price is also driven by market-wide price movements unrelated to this particular trade, and (2) $|Mret|$, the absolute value of *Mret*, which is meant to control for contemporaneous market-wide volatility.

The results from the regression in Equation (2) are reported in Column 1 of Table 4. In this test, we only focus on institutional trades executed during the 6-month period surrounding the fee change. Importantly, we find a significant decrease in the price impact coefficient following

Table 4 Institutional trading costs around regulatory fee changes.

	[−3, +3]	Months	All Months	Size > \$500 K	Size > \$1 M
	(1)	(2)	(3)	(4)	(5)
ln(*TSize*)	6.742***	8.938***	9.114***	13.141***	15.569***
	(22.07)	(26.6)	(29.67)	(20.69)	(15.72)
Fee	3.002***	3.615***	2.859***	2.057**	2.261**
	(4.25)	(4.50)	(5.26)	(2.46)	(2.20)
Fee × ln(*TSize*)	−0.981**	−1.135**	−1.467***	−1.976***	−2.778**
	(−2.16)	(−2.36)	(−4.46)	(−2.73)	(−2.40)
Mret	0.245***	0.247***	0.237***	0.282***	0.304***
	(21.21)	(21.53)	(28.18)	(28.01)	(25.42)
|*Mret*|	49.126	46.234	67.574	64.475	12.754
	(0.76)	(0.68)	(1.62)	(0.99)	(0.14)
Agg		0.185***	0.166***	0.134	0.098***
		(22.83)	(36.99)	(17.63)	(9.16)
Time		−0.942***	−1.080***	−1.135***	−1.086***
		(−5.76)	(−11.55)	(−8.76)	(−7.26)
ln(*Dvol*)		−4.059***	−4.405***	−5.483***	−4.262***
		(−6.15)	(−12.12)	(−8.59)	(4.65)
SE clustering	Stock-Date	Stock-Date	Stock-Date	Stock-Date	Stock-Date
Fixed effects	Stock	Stock	Stock	Stock	Stock
N	279,140	251,584	733,890	263,419	141,739
R-squared	0.061	0.071	0.063	0.077	0.085

This table reports results from OLS regressions that test the effect of the integrated fee model on the spread and price impact for large institutional trades. The dependent variable is the implementation shortfall of large institutional trades. Key dependent variables include *Fee* and *Fee* × ln(*Tsize*). The regressions in Columns 1 and 2 restrict the sample to the 6-month period surrounding the fee change. Column 3 uses the entire time sample. Columns 4 and 5 use the entire time sample and restrict the minimum institutional trade size to \$500,000 and \$1 million, respectively. All control variables are specified in the main body of the text. Standard errors are double clustered by stock and date. *t*-Statistics are reported in parentheses below the regression coefficients. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

the fee change. Specifically, the coefficient on ln(*TSize*) decreases by 14.5% (0.98/6.74), providing support for the hypothesis that the fee change led to lower price impact for institutional trades because it impaired HFT trading. We also find that the average spread for institutional trades increased by 3.0 basis points following the fee change. This latter result corroborates evidence in Malinova *et al.* (2018) showing that average execution costs for retail trades and institutional trades using only marketable limit orders increased after the fee change. For sufficiently large institutional trades, the reduced transaction costs from the deeper market can outweigh the increased transaction costs from the wider spread, leading to lower overall execution costs. We calculate the "break-even" trade size by setting the marginal cost from the increased spread equal to the marginal savings from the increased depth and find a break-even point of about \$2.1 million ($\exp(3.0/0.98) \times \$100{,}000$). In our sample, the top 7% of institutional trades by size exceed this \$2.1 million threshold and account for approximately 45% of total trading volume of institutional trades. Thus, it appears that the integrated fee model reduced implementation shortfall for approximately 45% of the large institutional trading volume and increased implementation shortfall for the remaining 55%.

Columns 2 to 5 of Table 4 demonstrate that the results from Column 1 are robust to alternative model specifications. In Column 2, we include additional trade and stock-level control variables: (1) the aggressiveness of the institutional trade (*Agg*), (2) the number of hours it takes to complete the institutional trade (*Time*), and (3) the natural log of stock dollar volume, expressed in millions of dollars, on the day that the institutional trade was initiated (ln(*Dvol*)). As one would expect, we find that more aggressive trades have higher implementation shortfall, reflecting the cost of taking liquidity from the limit-order book. Trades with shorter time spans also have higher implementation shortfall, suggesting that these traders are trading on short-lived information and pay a price for quick execution. We also find that implementation shortfall is lower on days with higher trading volume, since institutional traders can more easily "hide" within the noise trading when executing their trades. In Column 3, we use all of the trading days in the sample as opposed to the 6 months surrounding the fee change. In Columns 4 and 5, we redefine institutional trades as having a minimum institutional trade size of \$500,000 and \$1 million and standardize the *TSize* variable using the minimum institutional trade size in that subsample. For those tests, the results are similar to the baseline results, with spreads increasing and price impact decreasing for large institutional trades following the fee change.

3.2 Price impact analysis by institutional trader type

It is possible that our results in the previous subsection are being driven by changes in overall market conditions or other unobserved outcomes around the time of the fee change. We address this concern by exploiting heterogeneities in the trading motivations of institutional traders. Some institutions allocate resources toward obtaining an informational advantage, and thus are motivated to trade on their information. Other institutions often trade for reasons unrelated to information about fundamental security values, such as the need to accommodate fund inflows or redemptions. Theory suggests that HFTs compete with informed institutional traders, who trade with greater urgency because of their often short-lived information advantage. If the integrated fee model makes it more costly for HFTs to compete with these traders, as our previous results suggest, then we would expect a stronger reduction in price impact for informed institutional traders compared to uninformed institutional traders.

To classify informed institutional traders, we calculate the return for each institutional trade. For institutional buy (sell) orders, the return is calculated as (the negative of) the percentage difference between the closing price 5 days after the trade has been executed and the share-weighted average price of the parent order. The 5-day window for calculating returns for institutional trades is also used in Chan and Lakonishok (1995). This represents a balance between a shorter window, which is more likely to produce return estimates that are biased by the transitory price impact of the trade, and a longer window, which produces noisier return estimates of trade performance. We calculate the average return for each institutional trader on a monthly basis. Institutions are placed into terciles based on their average historical returns up to the end of the previous month, where the terciles are denoted by $g \in \{H, M, L\}$ (high, medium, low). We consider institutions in the highest tercile as most likely to be trading on private information.

We first examine whether past institutional trader profits are related to out-of-sample profits. For each institutional parent order, we calculate the return for a number time of windows relative to the price at the beginning of the trade. The time windows we consider are the end of the trade at day t and the close of days $t + 1$, $t + 5$, and $t + 20$. For example, if an institution initially bought for \$100.00 and the closing price 5 days later is \$105.00, then the return equals 5%. If institutions in the high-informed group are truly informed, then we should expect these institutions to incur positive and significant average returns in future time periods. Figure 1 displays the cumulative returns for institutional trades originating from each of the informed terciles. For the highly informed group, we find that the mean 5-day return relative to the price at the beginning of the institutional trade is about 17 basis points, which is significantly higher than the 5-day returns for the low and medium informed groups. The results are even stronger for mean 20-day returns, with the highly informed group averaging about 26 basis points and the low and medium informed groups averaging about 10 basis points. If we factor in the transaction costs from the implementation shortfall of the trade, then we find that the positive 20-day returns are only significant for the highly informed group. Overall, our evidence indicates that institutions in the highly informed group trade profitably out-of-sample.

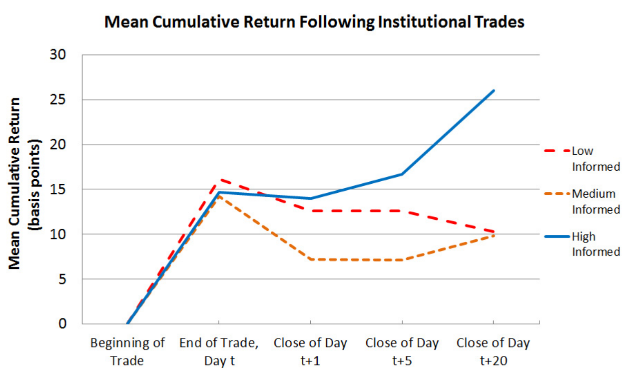

Figure 1 Price evolution around institutional trade by informed type. This graph displays the mean cumulative return relative to the price at the beginning of an institutional trade ($\boldsymbol{p}_0$). For institutional buy (sell) orders, we calculate the cumulative return at t as (the negative of) $(\boldsymbol{p}_t - \boldsymbol{p}_0)/\boldsymbol{p}_0$. The y-axis is reported in basis points.

We examine the differential effect of the integrated fee model on the execution costs for the three informed institutional trader groups. For each informed group $g \in \{H, M, L\}$, we estimate the baseline regression model in the previous subsection. If our hypothesis is correct about the integrated fee model having the largest decrease in price impact for the high-informed group, then we would expect a more negative β_3 for group H compared to the other two groups (M and L).

The results are reported in Columns 1 to 3 of Table 5. Institutional trades from the high-informed group are analyzed in Column 1. We find that, following the fee change, the price impact of large institutional trades from the high-informed group decreases by approximately 27.5% (2.69/9.79); this change is statistically significant at the 1% level. In contrast, Columns 2 and 3 indicate a 10.1% decrease (0.89/8.80) in price impact for the medium-informed group (significant at the 10% level) and no statistically significant decrease in price impact for the low-informed group. The regression in Column 4 formally tests the change in price impact for the high-informed group relative to the changes in price impact for the other two groups. Specifically, we use a pooled sample of institutional trades from all three groups and construct indicator variables representing the medium-informed and low-informed groups ($\mathbf{1}_M$ and $\mathbf{1}_L$) and interaction

Table 5 Institutional trading costs around fee change by informed type.

	Trader Informativeness			
	High	Medium	Low	Pooled
	(1)	(2)	(3)	(4)
ln(*TSize*)	9.790***	8.797***	8.581***	9.067***
	(14.98)	(19.17)	(10.50)	(24.02)
Fee	0.424	2.129***	4.579***	0.156
	(0.35)	(3.07)	(3.43)	(0.14)
$Fee \times \ln(TSize)$	−2.693***	−0.889*	−1.069	−2.342***
	(−3.74)	(−1.86)	(−1.26)	(−4.97)
$Fee \times \mathbf{1}_{\mathrm{M}}$				2.138*
				(1.80)
$Fee \times \ln(Tsize) \times \mathbf{1}_{\mathrm{M}}$				1.203***
				(3.56)
$Fee \times \mathbf{1}_{\mathrm{L}}$				4.202***
				(2.81)
$Fee \times \ln(TSize) \times \mathbf{1}_{\mathrm{L}}$				1.333***
				(3.54)
$Fee \times (1 + \mathbf{1}_{\mathrm{M}})$				2.294***
p-Value				0.001
$Fee \times \ln(TSize) \times (1 + \mathbf{1}_{\mathrm{M}})$				−1.139***
p-Value				0.007
$Fee \times (1 + \mathbf{1}_{\mathrm{L}})$				4.358***
p-Value				0.000
$Fee \times \ln(TSize) \times (1 + \mathbf{1}_{L})$				−1.009**
p-Value				0.027
SE clustering	Stock-Date	Stock-Date	Stock-Date	Stock-Date
Controls	Yes	Yes	Yes	Yes
Fixed effects	Stock	Stock	Stock	Stock-Group
N	168,056	327,499	165,766	650,492
R-squared	0.068	0.057	0.061	0.064

The first three columns of this table report results from OLS regressions that examine the effect of the integrated fee model on the spread and price impact of large institutional trades for each informed group $g \in \{H, M, L\}$. Column 4 reports the results of a pooled regression containing all informed types. The dependent variable in all of the regression tests is the implementation shortfall of large institutional trades. All control variables are specified in the main body of the text. Standard errors are double clustered by stock and date. t-Statistics are reported in parentheses below the regression coefficients. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

terms between each of these indicator variables and the key dependent variables (*Fee* and *Fee* × ln(*TSize*)). Similar to the tests in Columns 1 to 3, we find a much stronger negative effect on the price impact of large institutional trades for the highly informed group than for the other two groups. The differences are statistically significant. We also provide graphical evidence of the implementation shortfall changes for all traders and each of the informed groups in Figure 2 by plotting the average *IS* for each week. Panels A through C indicate that the mean *IS* for the pooled sample and the low and medium informed groups increased, and panel D indicates that the mean *IS* for the highly informed group decreased, a finding consistent with the evidence from our regressions. Furthermore, these graphs indicate that our results are not being driven by general trends in institutional execution costs. Overall, the results from the regressions and figures indicate that the integrated fee model led to higher informational rents for informed institutional traders through the price impact channel. Furthermore, this evidence suggests that some traders could be deterred from acquiring costly information in the presence of HFT, which is consistent with the evidence in

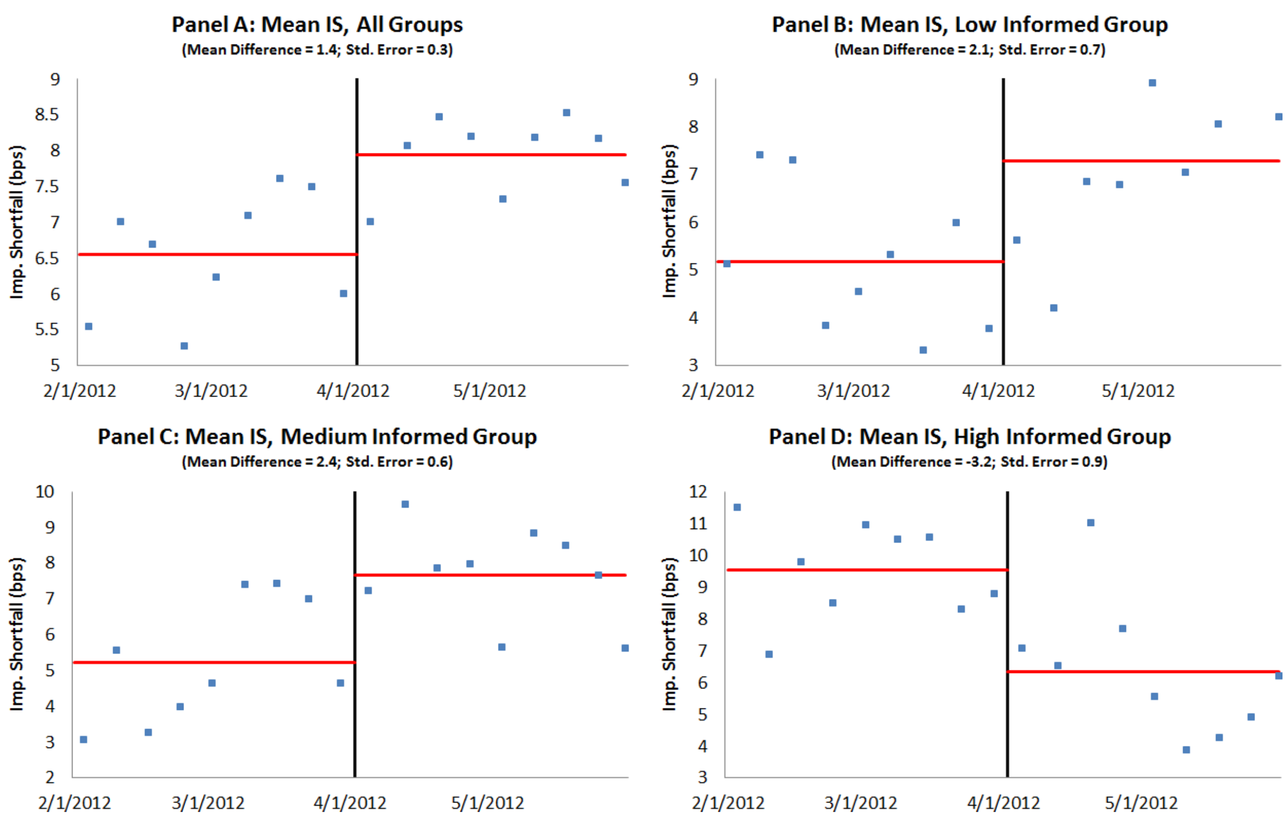

Figure 2 Average implementation shortfall by informed type. This figure displays the weekly mean implementation shortfall (***IS***) for all traders (panel A) and the three informed groups (panels B to D). The vertical line in each panel indicates the day that the integrated fee model went into effect. The horizontal lines indicate the average ***IS*** before and after the fee change. In the caption of each panel, we report the mean difference in the horizontal lines and the standard error of the difference using the the daily average ***IS***. The *y*-axis is reported in basis points.

Weller (2018) showing that algorithmic trading can lead to reduced price efficiency because it deters information acquisition.

4 Conclusion

The introduction of message fees in Canadian equities markets in April 2012 had a strong effect on HFTs because HFT strategies typically rely on high message activity. In this paper, we show that price impact for large institutional trades decreased by 15% and spreads increased by 3.0 basis points after the introduction of the message fees. A break-even analysis indicates that institutional trades above the $2.1 million threshold had a net reduction in total transaction costs after the fee introduction. For institutions that are more likely to be trading on information about future price movements, the price impact reduction was even larger at 27.5%. Taken together, our results suggest that the change in the fee structure made it more costly for HFTs to adjust their limit orders in response to the information they infer from large institutional orders. An additional interpretation is that the latency advantage for HFTs relative to institutional traders was mitigated by the message fees. A consequence of these message costs was that institutional traders above a certain trade size threshold could earn higher profits from their information-based trades, even after taking into account the evidence that HFTs widened their spreads in response to the message fees.

Overall, our results provide an important link between HFT activity and institutional trading costs, and suggest that high-frequency trading is associated with higher execution costs for large and information-based institutional orders and lower costs for small, uninformed orders. The common practice of focusing on how HFT activity affects bid–ask spreads only looks at one dimension of a multidimensional problem. The other dimension is the change in price impact across traders with different informational advantages. Regulators should incorporate this multifaceted effect of HFT on institutional execution costs when contemplating regulatory changes, particularly those that mitigate the latency advantages incurred by HFTs.

Notes

1 This information and additional details can be found at www.iiroc.ca/about.

2 Henceforth, we refer to large, institutional-sized trades as "institutional trades," although it is possible that some of these trades are coming from high net worth individuals.

3 The TMX IP is operated by the TMX Group, which operates some of the largest stock exchanges in Canada, including the Toronto Stock Exchange (TSX), the TSX Venture Exchange, and the TSX Alpha Exchange. The TMX IP is also required to meet standards that are approved by Canadian regulators.

4 We link an institutional trade across days if its execution involves at least one trade in both the last half-hour of day t and the first half-hour of day $t + 1$.

5 For more information about the integrated fee model and feedback from marketplace members about the proposed model, see IIROC Notices 11-0125 and 12-0043.

Acknowledgments

We thank the Investment Industry Regulatory Organization of Canada (IIROC) for providing us with access to the data used for this study and Victoria Pinnington and Helen Hogarth of IIROC for answering questions regarding the data and Canadian market structure details. We also thank Gifford Fong, Bradley Katsuyama, an anonymous referee, and seminar participants at IEX Group and the Commodities Futures Trading Commission for their helpful comments.

References

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

Brogaard, J., Hendershott, T., and Riordan, R. (2014). "High Frequency Trading and Price Discovery," *Review of Financial Studies* **27**, 2267–2306.

CFA Institute (2012). "Dark Pools, Internalization, and Equity Market Quality," https://www.cfainstitute.org/-/media/documents/article/position-paper/dark-pools-internalization-equity-market-quality.ashx.

Chan, L. K. C. and Lakonishok, J. (1995). "The Behavior of Stock Prices Around Institutional Trades," *Journal of Finance* **50**, 1147–1174.

Comerton-Forde, C., Malinova, K., and Park, A. (2018). "Regulating Dark Trading: Order Flow Segmentation and Market Quality," *Journal of Financial Economics* **130**, 347–366.

Cready, W., Kumas, A., and Subasi, M. (2014). "Are Trade Size-Based Inferences about Traders Reliable? Evidence from Institutional Earnings-Related Trading," *Journal of Accounting Research* **52**, 877–909.

Crippen, A. (2014). "Buffett, Gates and Munger Criticize High-Frequency Trading," *CNBC*, May 5. www.cnbc.com/2014/05/05/buffett-gates-and-munger-criticize-high-frequency-trading.html.

Devani, B., Anderson, L., and Zhang, Y. (2015). "Impact of the Dark Rule Amendments: Trading Review and Analysis," Discussion Paper, IIROC.

Devani, B., Tayal, B., Anderson, L., Zhou, D., Gomez, J., and Taylor, G. W. (2014). "Identifying Trading Groups—Methodology and Results," Discussion Paper, IIROC.

Glosten, L. R. and Milgrom, P. R. (1985). "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics* **14**, 71–100.

Ho, T. and Stoll, H. R. (1981). "Optimal Dealer Pricing under Transactions and Return Uncertainty," *Journal of Financial Economics* **9**, 47–73.

Keim, D. B. and Madhavan, A. (1997). "Transactions Costs and Investment Style: An Interexchange Analysis of Institutional Equity Trades," *Journal of Financial Economics* **46**, 265–292.

Korajczyk, R. A. and Murphy, D. (2019). "High-Frequency Market Making to Large Institutional Trades," *Review of Financial Studies* **32**, 1034–1067.

Kyle, A. S. (1985). "Continuous Auctions and Insider Trading," *Econometrica* **53**, 1315–1335.

Laming, D. R. J. (1968). *Information Theory of Choice-Reaction Times* (Oxford, England: Academic Press).

Malinova, K., Park, A. and Riordan, R. (2018). "Do Retail Investors Suffer from High Frequency Traders?" Working Paper, University of Toronto.

Menkveld, A. J. (2013). "High Frequency Trading and the New Market Makers," *Journal of Financial Markets* **16**, 712–740.

Perold, A. (1988). "The Implementation Shortfall: Paper versus Reality," *Journal of Portfolio Management* **14**, 4–9.

Weller, B. M. (2018). "Does Algorithmic Trading Deter Information Acquisition?" *Review of Financial Studies* **31**, 2184–2226.

Keywords: High-frequency trading; institutional trades; implementation shortfall; transaction costs; transaction tax

JEL codes: G1, G11, G18, G24