The Cost of Routing Orders to High Frequency Traders

Robert Battalio
Mendoza College of Business
University of Notre Dame
email: rbattali@nd.edu

Brian Hatch
Lindner College of Business
University of Cincinnati
email: brian.hatch@uc.edu

Mehmet Sağlam
Lindner College of Business
University of Cincinnati
email: mehmet.saglam@uc.edu

This Version: November 2018

Abstract

Institutional brokers and high-frequency liquidity providers (HFLPs) engage in trading relationships that involve monetary incentives for brokers if they route their clients' orders to HFLPs. This paper examines whether the institutional clients benefit from these routing relationships. We analyze a novel dataset of institutional large-orders whose child-orders can be routed to the HFLPs with investor's consent and find that orders that are routed to HFLPs realize substantially higher trading costs. We obtain causal evidence by utilizing the trades of investors who oppose HFLP routing. The cost increase is due to early information leakage about the investor's liquidity need.

Keywords: High-Frequency Trading, Transaction Costs, Smart Routing

JEL Classification: G12, G14.

1. Introduction

Many retail and institutional brokers route their orders to high-frequency traders for execution. Usually referred to as "electronic liquidity providers," "internalizers" or "off-exchange market makers," these high-frequency liquidity providers (HFLPs) handle the routed orders in a fully automated fashion and decide algorithmically to execute the order at a particular price or not. This relationship has various benefits to both parties. For retail order flow, which is typically considered to be uninformed, the brokers directly receive payments from the HFLPs on the basis of executed marketable or non-marketable orders and the HFLPs profit from the spread between the fair value of the security and the transaction price.[1] This practice is called "payment for order flow" and has received considerable attention from policy-makers and academics with regards to execution quality. However, at least equally important case with institutional orders, which may be initiated for informational reasons about the value of the security, has been relatively unexplored in the academic literature partly due to lack of available data. In this paper, we aim to fill this gap.

In institutional order routing, the contract terms between the broker and the HFLP are more opaque. Marketing documents of the brokers suggest that they negotiate with HFLPs over the exact rebate or fee structure.[2] Recent settlement agreement between the State of New York and Bank of America Merrill Lynch explicitly writes that the broker does not pay the typical liquidity-take fees if the routed order gets executed by an HFLP.[3] Given that these fees (rebates) are directly paid (received) by the broker but not by the client submitting the order, this arrangement creates a significant incentive for the broker to route the orders to an HFLP.[4] On the other side of the agreement, HFLPs can benefit from these routed orders as they may still earn the spread between the fair value of the security and the transaction price. In theory, they can also take additional

[1] A large retail broker, Charles Schwab, states in its order routing report in 2018 Q1 that roughly 50% of the retail orders are routed to Citadel and Virtu and it receives payments for marketable orders executed through these firms, which averaged less than $.0009 per share.

[2] For example, Advanced Execution Services at Credit Suisse mentions in a marketing document that they negotiate rates with the HFLPs and the exact monetary benefit depends on the extent of the interaction with the venue, e.g., volume traded.

[3] `https://ag.ny.gov/sites/default/files/bofaml_settlement_agreement.pdf`

[4] Usually, clients pay for a bundled package of services from the broker that may include research reports, clearing, custody, margin accounts and securities lending rather than paying for a specific trade. Take fees depend on the venue and the security. For example, as of May 1, 2018, NYSE ARCA charges $0.0030 per share for each liquidity-taking order on Tape A security.

risk and unwind the accumulated position with passive limit orders in the next few seconds and earn the rebate associated with adding liquidity.[5] More importantly, since this trade is potentially part of a large order submitted by an institutional investor who has either information or an urgent liquidity need, the HFLPs also gain valuable order flow information that they can utilize for their market making activities. Overall, this arrangement between the broker and the HFLPs seems mutually beneficial but it is not clear whether the investor whose orders are routed benefits from this relationship via lower transaction costs.

There are broadly four competing theories about how this trading relationship can affect the investor's transaction costs. The first theory is based on predatory market making. Aït-Sahalia and Sağlam (2013) study dynamic liquidity provision in the context of market-making employed by a high-frequency trader. The high-frequency market maker is able to predict a low-frequency trader's liquidity need with an imperfect signal. When the market maker is confident that he is going to trade with an impatient trader, he widens his quotes before his order reaches to the market. The second theory is based on predatory trading (Brunnermeier and Pedersen, 2005). When an order is routed to an HFLP, he ultimately receives a signal about the urgency of the investor. Exploiting this liquidity need, the HFLP can start trading in the same direction with the investor and this early trading can make future trades of the investor more expensive. The HFLP can then unwind his position to the investor during the later stages of the execution and make profits without taking too much risk. The third theory is based on sharing informational rents (Yang and Zhu, 2015). If the investor initiates the trade based on private information and the HFLP learns about this from the past price impact of the trades, it would be better for the HFLP to trade in the same direction during the execution period. This same trading can increase the trading costs of the investor. The final theory is based on sunshine trading due to Admati and Pfleiderer (1991) and has opposite predictions compared to the earlier theories. In this scenario, if the HFLP can infer that the large-order is uninformed, he may be incentivized to provide liquidity to the large-order in the absence of adverse selection costs and decrease the overall execution costs.[6] All of these theories may be at

[5]For example, as of May 1, 2018, NYSE ARCA pays $0.0012 per share for each order adding liquidity on Tape A security.

[6]A related sunshine trading theory is provided by Bessembinder et al. (2016) by extending the model of Brunnermeier and Pedersen (2005) for resilient markets with transitory price impact. In this model, in addition to the

play here, hence, the net effect of the routing relationship with HFLPs on execution costs is not clear.

In this paper, we aim to answer this question empirically by examining the impact of the routing relationship between the broker and the HFLPs on the resulting transaction costs. We use a novel dataset of institutional large-orders which can be routed to the HFLPs after obtaining the permission of the investor in the pre-trade phase. The dataset includes more than 20,000 parent-orders and 2.5 million child-order trades occurring between January 1, 2011 and March 31, 2012. Average order size is sizable with around $1 million and corresponds to roughly 2% of the volume realized during the parent-order execution. The dataset exactly identifies the venue of each child-order trade from Financial Information Exchange (FIX) protocol tags and allows us to exactly identify the set of HFLPs who may provide liquidity to the large parent-order. These HFLPs include the largest high-frequency traders as of the data period: Citadel, D.E. Shaw, Getco, Knight, Sun Trading and Two Sigma. The average ratio of dollar volume executed by HFLPs is more than 5% and more than 60% of the parent-orders have at least one child-order filled by this group of HFLPs.

We first examine the timing of the HFLP trades and find a strong U-shaped pattern across the execution horizon. Specifically, we observe that a high-share of HFLP trades occur in the early stages of the execution which may be very valuable to HFLPs. Strikingly, the additional volume to the HFLPs in the early stages seems to be taken from the share of lit venues.

Using various transaction cost proxies, we then document robust evidence that parent-orders whose child-orders are routed to HFLPs suffer higher transaction costs at the univariate and multivariate level analyses. Specifically, at the univariate level, HFLP venues are the costliest and successive fills in them lead to larger spread costs. This finding is robust to controlling for informed trading proxy, and various control variables including stock, day and investor fixed-effects. In the most conservative case, if 1% of the parent-order is routed to HFLPs instead of lit venues, the net transaction costs increase by roughly 12%.

Since the routing decision and the child-order trades are all determined in equilibrium, quantify-

same-side trading before the liquidation, the strategic investors can trade against the liquidator and decrease his trading costs if the market is largely resilient.

ing a causal effect is empirically challenging. For example, the broker may be routing the most toxic orders to the HFLPs and if these variables are not properly controlled in the regression, this may bias the estimated impact of the HFLP exposure. Our dataset addresses this concern as it includes a set of investors whose child-orders are almost never routed to an HFLP. We exploit this variations in the dataset to test the causal relationship and find that parent-orders exposed to HFLPs have higher spread and price impact costs compared to not-exposed but matched parent-orders that are on the same stock and occurring in a similar period.

Further, there is an active investor in the dataset who has child-orders executed by HFLPs in the early periods of the data but then has no such child-orders in the remaining period suggesting that he strategically disallowed the broker to route his child-orders to the HFLPs after a particular date. We design a difference-in-difference regression to exploit this variation and find statistically significant decrease in this investor's transaction costs after eliminating the HFLP exposure.

Next, we investigate the underlying mechanism further in detail for the cost increase associated with HFLPs. First, we find that when 10% of the share volume goes from lit venues to HFLPs, the fraction of aggressive fills actually increases by 3.7% and the fraction of passive (mid-point) fills decrease by 2.9% (0.8%). This findings implies that HFLP trades do not substitute one-for-one with aggressive lit venue trades. Second, we examine the cost of exposing the child-orders to HFLPs in the early stages and find that early exposure is substantially costlier. Thus, one main driver of the cost increase is the price impact associated with early exposure. We also find that exposing the child-orders to multiple HFLPs also increases the price impact costs.

Finally, we evaluate our findings mainly in the light of the predictions of the predatory trading and back-running theory. Sunshine trading does not seem to align with our findings as the empirical findings uniformly point to increase in execution costs even in the presence of competing HFLPs. Back-running theory does not seem to support our findings either as we find that in the presence of an informed order, HFLPs do not significantly cut their liquidity provision in the late stages of the execution. The empirical findings support some predictions of the predatory trading. First, (successive fills) in HFLPs seem to have much larger spread costs aligning with the predatory market making theory presented in Aït-Sahalia and Sağlam (2013). Further, we show that the transitory component of the spread costs are also correlated with HFLP trading activity which also aligns with predatory trading motives rather than back-running. However, our set of HFLP trades is restricted to our data set, i.e., we do not observe how HFLPs trade in the lit markets. Thus, this analysis is limited to HFLP trading activity in the dark.

These contributions have important implications for the market structure. First, our empirical evidence largely imply that these routing relationships between HFLPs and the brokers are not in the best interest of the investors as HFLPs appear to take advantage of the liquidity needs of the investors. Second, our results could be beneficial to derive a set of disclosure requirements for routing practices. In July 2016, The Securities and Exchange Commission (SEC) voted to propose rules that for the first time would require broker-dealers to disclose the handling of institutional orders to customers. These rules mention a disclosure on execution quality about the routed orders and require brokers to provide statistics on percentage of total shares executed that were priced at the side of the spread more and less favorable to the institutional order. Although these statistics would be valuable, our findings suggest that the more important statistics would be on the price impact of the routed orders.

The rest of the paper is organized as follows: Section 2 provides a brief literature review. Section 3 reviews the relevant theoretical models that allow us to form our hypotheses regarding the impact of HFLP fills. In Section 4, we describe the dataset and provide its summary statistics. In Section 5, we specifically examine the trading patterns of HFLPs during the lifetime of the parent-order. Section 6 studies the correlation between trading cost proxies and HFLP trading activity. Section 8 utilizes the restricted coverage of the routing relationship and establishes causal link between HFLP exposure and execution costs. Section 9 investigates the underlying mechanisms for the documented increase in *IS* and evaluates the findings in terms of the summarized theories in Section 3. Finally, we conclude in Section 10.

2. Related Literature

This paper is related to three strands of the microstructure literature. First, there are a number of studies that examine trading costs and liquidity provision due to conflict of interests arising from payment for order flow and make-and-take fees. Battalio et al. (2016) study the limit order execution quality when retail brokers make order routing decisions that aim to maximize their rebate revenue. They find that retail brokers sell their market orders to increase order flow payments and send their limit orders to market exchanges with the highest rebates without considering the probability of execution. Easley et al. (1996) illustrate that payment for order flow may lead to fragmentation of the orders by information content and as the uninformed order flow is purchased and filtered out, this may worsen liquidity on the NYSE where informed trading takes place. Instead, examining purchased market orders of Knight Securities, Battalio et al. (2001) find no evidence that traders using brokers taking order flow payments are worse-off than traders using a broker that does not accept such payments.

Second, there is a growing literature studying the relationship between institutional trading costs and activities of high-frequency traders (HFTs). In this strand, the most closely related study is van Kervel and Menkveld (2017) who use child-order trading data reported from four institutional investors based in Sweden and infer parent-orders by stringing them together. They exclude 11.5% of institutional orders that does end up with purely buys or purely sells. They construct the net flows of HFT trading around the institutional order by using the trade reports provided by Nasdaq OMX, a public lit exchange, which has 65% of the market share. They do not have data on HFT trading activity in dark pools. There are significant differences in contribution between this paper and our study. First, van Kervel and Menkveld (2017) investigate the cost dynamics of large orders by studying the interaction between the aggregate HFT trades while we study the impact of the direct information leakage when a child-order is routed to an HFLP for execution. In van Kervel and Menkveld (2017), it is not clear how the HFT detects the presence of the large-order whereas the detection is almost immediate in our study due to the routing of the order. Second, van Kervel and Menkveld (2017) do not claim causality whereas we identify causal relationship by studying the variation through the investor's ability to disallow the broker to route his order to HFLPs. Third, we find empirical evidence supporting predatory trading motives which contrasts with the reported evidence for back-running in van Kervel and Menkveld (2017). Finally, our dataset provides more granular data at the institutional side and has information about HFLP child-order trades in the *non-lit* venues. On the contrary, van Kervel and Menkveld (2017) has high-quality data for HFT trading activity occurring in a *lit* venue. In this sense, these two papers complement each other. We have perfect knowledge of the parent-orders with exact start- and end-time which allows us to exactly compute the statistics around the execution interval. Further, our dataset involves the trading activity of a larger investor universe with 146 distinct clients that utilize a specific broker's single trading algorithm. The benefit of focusing on one algorithm is that it removes the variation in execution costs due to heterogeneity across brokers and their various algorithms. Larger investor base helps us to study the impact across a heterogeneous group of informed and uninformed traders.

Hirschey (2013) finds empirical evidence supporting that HFTs may increase the trading cost of non-HFTs by trading ahead of them. He obtains his dataset from NASDAQ which labels trading firms either an HFT or a non-HFT. The data is focused on this distinction and does not allow him to identify a specific instance of a large order execution and its corresponding cost. Brogaard et al. (2014) examine the shocks to the latency of London Stock Exchange that increase HFT activity over time and study the variation of institutional trading costs around these changes using data provided by Ancerno Ltd. (formerly known as Abel/Noser), a consulting firm specialized in the analysis of institutional trading costs. They cannot find any clear evidence of change in trading costs during these latency upgrades. Korajczyk and Murphy (2015) show that HFT liquidity provision is significantly lower for stressful large order executions by institutional investors using order-level data from Investment Industry Regulatory Organization of Canada, a regulatory organization for Canada's equity markets. As in the case of van Kervel and Menkveld (2017), this dataset does not exactly identify how a parent order is split into child orders. HFTs initially provide liquidity to the large order but then competes with it due to inventory management and back-running.

Our paper complements these studies by directly studying an established routing relationship between a broker and a set of known HFLPs. Since some investors opt out of the routing relationship, we obtain causal evidence between exposure to HFLPs and execution costs. To our knowledge, this is the first paper that studies the impact of such institutional routing agreements on investors' trading costs.

Finally, our paper is related to the understanding the liquidity implications of dark venue executions. All of the HFLP trades in our dataset show up as dark pool trades in the publicly available TAQ dataset. Exchange code 'D' is used in the TAQ dataset to identify all trading within dark pools as well as internalized trades at the brokers. There are a few studies that use this classification to examine the impact of dark pool trades on market quality (e.g., O'Hara and Ye (2011), Hatheway et al. (2017) and Farley et al. (2018)). Given the special nature of the routing relationships, it is not clear how to aggregate them with the rest of the dark pool trading activity. The granularity of our dataset allows us to study the liquidity implications of a unique group of dark-labelled trades that was not possible to study in the past literature.

3. Theoretical Framework and Hypotheses Development

In this section, we review the relevant theoretical models that guide our empirical analysis on the impact of the routing relationship with the HFLPs. We summarize our expectations from the perspectives of four theories: payment for order flow, predatory trading, back-running, and sunshine trading with competing liquidity providers. Sunshine trading would imply lower trading costs whereas the remaining three theories imply higher trading costs.

Aït-Sahalia and Sağlam (2013) study dynamic liquidity provision in the context of market-making employed by a high-frequency trader. The high-frequency market maker is able to predict a low-frequency trader's liquidity need with an imperfect signal. When the market maker is confident that he is going to trade with an impatient trader, he lowers his bid quote before his order reaches to the market. In the context of our large order executions, the HFLP instantly learns that there is an impatient trader with a liquidity need thus this model also predicts that HFLPs can exploit the impatience of the investor by charging higher spread. Overall, this theory predicts larger spread costs for HFLP-routed orders as opposed to larger price impact.

When an order is routed to an HFLP, it immediately learns that there is a high probability that a large order is being executed. This information can be valuable for the HFLP to the extent of preying on it. Brunnermeier and Pedersen (2005) provide a model of this kind predatory trading in which a distressed large investor is forced to sell his position and other strategic adversaries aim to exploit from this liquidity need. Instead of providing liquidity, these traders exploit their information and trade in the same direction initially. These initial trading can cause the price to decrease further. At further depressed prices, predatory traders then buy from the investor, which ultimately increases the total liquidation cost of the investor. It is worth to note that in our setting this information leakage cannot lead to a riskless arbitrage opportunity.[7] Overall, this theory predicts larger price impact costs that is transitory.

The common implication of predatory trading is the temporary nature of the price impact. This implies that if predictable executions induce the stock price to go up during a large buy order, the stock price will be expected to revert back to its pre-execution level after a short period. However, predatory trading would not explain the persistence in the price impact. Such permanent price impact is consistent with the back-running theory presented in Madrigal (1996) and Yang and Zhu (2015). These papers use two-period Kyle models in which an informed investor trades on private information in the first period and strategic traders receive an imperfect signal about this in the second period. They then exploit this signal to trade in the same direction implied by the private information of the investor. In this sense, adversary traders are trying to "steal" the private information of the investor by inferring from the past order flow.

Admati and Pfleiderer (1991) provide an alternative theory of forced liquidation based on sunshine trading in which investors can potentially signal that their trade is not motivated by private information. This announcement can induce other strategic traders to provide liquidity and may actually reduce trading costs. In institutional large-order executions, it is unlikely that that all of the trades are uninformed so this theory may not directly apply in our context. However, a

[7]First, the routed order could be the last child order to be executed and thus there could be no further preying opportunity. Second, there is still fundamental risk in prices which may increase the risk in the predatory strategy. Third, there are many competing liquidity providers who may not be aware of this signal, thus it may be very risky to unwind the accumulated position. Finally, the broker can also walk away from the agreement if it starts losing investors for poor execution performance so the HFLP needs to take this into account while preying on this information.

related theory of sunshine trading is given by Bessembinder et al. (2016) who extend the model of Brunnermeier and Pedersen (2005) for resilient markets in which the immediate price impact of trades may be transitory. In this model, in addition to the same-side trading before the liquidation, the strategic investors trade in the opposite direction as the liquidator and decrease the liquidator's transitory price impact if the market is largely resilient. This benefit to the liquidator from strategic trading persists at any level of market resiliency if there are multiple strategic traders. This theory is more plausible in our context. Given that there are distinct HFLPs in the data, if the broker routes orders to multiple HFLPs, this theory would suggest a potential decrease in execution costs.

4. Data

We compile the data from several sources. Stock returns, volume, outstanding shares and prices come from the Center for Research in Security Prices (CRSP). Intraday trade and quote data come from the Trade and Quote (TAQ) database. Institutional large-order trades is provided by the global execution desk of a large investment bank. In the next section, we describe this dataset in detail and provide institutional details about the execution strategy.

4.1. Origins of the Data

For our empirical study, we use a detailed execution data from the historical order databases of a large investment bank ("the broker") in the United States. The broker is one of the top five providers of execution services by market share.

This data set was originally obtained to study the implications of investor heterogeneity on the estimation of the price impact. For this purpose, the executions in the sample have been selected from an active subset of the broker's clients. An investor is considered active if he has at least 100 and at most 500 VWAP parent-order executions on S&P 500 stocks between January 2011 and March 2012, each order is fully filled and lasts at least 10 minutes. These filters were chosen for the needs of the original research project.

- VWAP orders are passively managed so the performance of the trade would be a better indicator of the investor's timing ability rather than the broker's skill.

- S&P 500 stocks are very liquid and relatively homogeneous subset of the stock universe. Further, it is hard to have an investor trading with insider information.

- Clients with small number of executions (<100) are eliminated as their short-term trading skill may not be estimated reliably with fewer observations. Clients with high number of executions (>500) are filtered to have a balanced data set across clients and prevent any specific investor from fully driving the estimation procedure.

- Executions lasting less than 10 minutes are often very small in size.

With these criteria, we obtain 22,074 parent-orders from 146 unique clients in the original data set. Each parent-order is executed in multiple child-order trades. In the original dataset, we have 2.6 million child-order trades. Note that we only have access to information about the executed child-orders. We do not know whether the execution is due to a limit order or a market order. However, we can imperfectly infer the aggressiveness of the execution price by comparing the trade price to the prevailing best bid and offer prices in the market.

This dataset provides rich attributes at the parent- and child-order level. At the parent-order level, most of the statistics are based on the execution horizon. These statistics include order size, direction of the order (buy or sell), order start and end times, participation rate (the ratio of order size to the total volume during the trading interval), average execution price, proportional bid-ask spread and mid-quote volatility based on the duration of the execution. For each parent-order, we also have information at the child-order level. Child-order level statistics include the time (timestamped to the millisecond), size, venue (market center) and price of each child trade. All of the variables referred in the paper are explicitly defined in Table 1.

4.2. Final Data Set

We exclude executions which have less than 5 child-order trades or have value less than $50,000 at the arrival time of the order which correspond to approximately 1,500 parent-order executions.

We finally exclude additional 200 executions with missing entries of participation rate, spread, volatility, or duration.

The final sample consists of 20,335 executions coming from 9,856 buy and 10,479 sell orders on 498 stocks. There are 146 distinct investors submitting the orders.

4.3. Clients of the Broker

The orders originate from a diverse pool of investors, such as institutional portfolio managers, quantitative investment funds, internal trading desks and other brokers who aggregate their retail order flow. The dataset only reports the masked identity for each investor, thus it is not possible to know the underlying trading strategy they are following. For ease of reference, we will refer our investor universe as "institutional investors."

We do not know the compensation agreement with the individual clients and the broker. However, the broker informed us that there are two common practices: fixed commission versus pass-through. In the fixed commission scheme, the client pays a fixed fee for per share traded and any accumulated fees or rebates are the broker's responsibility. In the pass-through scheme, the client pays the fees and receives the rebates. In our data set, there is no indication that the VWAP algorithm is different across investors using different compensation schemes. Further, the broker explicitly noted us that some of their clients chose this broker to maximize their rebate revenue.

4.4. VWAP Strategy and Client Inputs

All of the large orders in this dataset are executed according to a single execution strategy: volume-weighted average price (VWAP) algorithm. The broker informed us that this trading algorithm is the most commonly employed strategy, constituting roughly 50% of all of the broker's execution volume. According to this strategy, parent-orders are executed in smaller child-order trades over the course of a trading day to achieve an average execution price that is as close as possible to the volume-weighted average price observed in the whole market during this trading period.

Based on our personal communication with the broker, the VWAP algorithm is implemented by taking a set of inputs from the client. First, the client submits a large order with either a target completion time or an urgency score in the broker's order submission platform. Second, the client can disallow the broker to route the child-orders to an HFLP or the dark pools by marking two separate check boxes. The broker refers to HFLPs as "Liquidity Partners" in its system. This opt-in feature seems to continue today in pre-trade instructions. Credit Suisse report a similar set of instructions in its order handling guidelines available at `https://www.credit-suisse.com/media/assets/sites/aes/doc/aes-us-order-handling-guidelines.pdf`. We do not have access to any of these pre-trade instructions. However, using the venue information from each child-order trades, we can infer whether a client consistently avoids HFLPs or the dark pools.

Note that even if the client allows for HFLP and dark pool trades, the algorithm ultimately decides on the routed venues. The algorithm then slices the large order using the historical volume curve over the past month between the initiation of the order and targeted completion time. In this set-up, the client directly affects the initiation of the VWAP algorithm with these choices including the order size, start-time and target end-time, but once the order is initiated, the clients do not have any control over how each child order size, price, or its timing is selected. In summary, the client may shape the parameters of the VWAP algorithm with a limited set of pre-trade instructions that may affect the ultimate execution cost.

We do not know how the VWAP algorithm chooses which venue to route to at any point in time. However, using the venue information from each child-order trades, we can infer some of the routing logic to dark pools versus lit markets. We undertake such analyses in Section 5

4.5. HFLPs and their Interaction with the Broker

The dataset exactly identifies the venue of each child-order trade from Financial Information Exchange (FIX) protocol tags and the broker has provided us a list where each venue is indicated by "HFLP", "Dark Pool" or "Exchange." Thus, we can exactly identify the set of HFLPs who may provide liquidity to the large-order execution without any assumption about the definition of high-frequency traders.

Table 3 shows the four-letter codes of these HFLPs that show up in the ExecBroker (Executing Broker) field in the FIX protocol. We match these codes to the HFLPs. The list of HFLPs includes the largest high-frequency traders as of the data period: Citadel, D.E. Shaw, Getco, Knight, Sun Trading and Two Sigma. We note that Knight operates two venues which may be different in their execution patterns. For our empirical analysis, we consider them as a single firm. This will not be an important distinction for our empirical analysis. Finally, we do not observe any executions from D.E. Shaw after January 7th, 2011, thus, during our data period, we have effectively 5 active HFLPs.

We do not know the exact monetary agreements between the broker and the HFLPs. We are informed that the HFLPs do not send bid or ask quotes to the broker but instead the broker routes the order to an HFLP and the HFLP agrees to fill the order with a particular price improvement compared to NBBO prices or the HFLP rejects to fill the order. Note that other order types beside VWAP can be also routed to HFLPs. For each filled order, the broker does not incur any explicit cost. Although we do not have any data on rejection rates, the broker explicitly told us that the rejection rates must be small per the routing agreements.

4.6. Summary Statistics

Table 2 provide additional summary statistics for our final execution data. On average, a parent-order has a market value of roughly \$1 million at the start of the execution. The order sizes range from \$0.05 million to \$62.8 million. The corresponding average participation rate is 1.8%. On average, a parent-order is executed in 128 child-order trades. Order duration ranges from 10 minutes to full trading day with a mean duration of roughly 3.4 hours. The average time-weighted bid-offer spread is approximately 4 bps and the corresponding volatility of the mid-quote price is 1.5% during the execution period.

The second panel in Table 2 reports that in terms of dollar size, 83.3% of the parent-order is executed in the lit exchanges ($LITdol$) in an average parent-order. Similarly, 10.8% and 5.4% of the average parent-order is executed in the broker's own dark pool, BODP, ($BODPdol$) and the HFLPs ($HFLPdol$), respectively. Getco and Citadel have the largest share of executions accounting for roughly 59.0% and 23.5% of all HFLP dollar volume, respectively. The share of other dark pools $OthrDPdol$ is considerably smaller accounting for only 0.4% of the parent-order. The large difference between the share of the BODP and other dark pools is very striking.

In the third and fourth panel, we report the ratios at the level of executed shares and number of trades which are approximately similar. When trade numbers are compared, there is a slight increase in the ratios for lit venues and HFLPs implying that typical dark pool trades are larger in dollar size.

Finally, in the fifth panel, we examine the routing behavior at the binary level. Let *HasVenue* be a binary variable taking value of 1 if the *Venue* has a child-order trade. We find that on an average basis 61.1% of all parent-orders has at least one child order executed by an HFLP as shown by the variable *HasHFLP*. This ratio drops to 48.3% for the indicator variable for the BODP, *HasBODP*. Almost all of the executions have a child-order trade in a lit venue.

4.7. Comparison of the Data with other Institutional Trading Data

In terms of order size and participation rate, our execution data is very similar to other datasets employed in the broader microstructure literature. For example, Anand et al. (2011) use Ancerno dataset and compute that the average daily participation of the institutions is 2.1% of the total market volume between 1999 and 2008 (1.0% in 2008). Compared to Ancerno dataset, our dataset has information on the exact start- and end-time of the execution (i.e., order duration), the interval volume during the execution (i.e., participation rate), the algorithm type (i.e., VWAP), and interval return.[8] More importantly, we have information on the executed child-orders, e.g., the price, quantity and execution venue of the child-order which enables to study the impact of HFLP trades.

Korajczyk and Murphy (2015) and van Kervel and Menkveld (2017) also employ institutional trading data (from Canada and Sweden, respectively) to examine the impact of high-frequency trading on execution costs. Korajczyk and Murphy (2015) report an average trade size of $2.2 million with a participation rate of 2.5%. van Kervel and Menkveld (2017) examine roughly 5,000 orders from Sweden during a similar time period as our data and report an average trade size of $2.2 million corresponding to a participation rate of 3.6%. Overall, these similar statistics support the representativeness of our institutional trading data.

[8] Hu et al. (2017) report that Ancerno dataset has client identifiers only through 2011.

5. Understanding HFLP Executions

In this section, we would like to uncover the drivers in HFLP fills in detail. Doe to the sensitivity of the topic, the data providers were not able to share the exact logic behind how the routing to HFLPs works. In this section, we run several analyses to understand the variation in HFLP fills.

5.1. When Do HFLP Executions Occur?

If HFLP fills occur very late in the lifetime of the parent-order, then the information leakage associated with them may be smaller. To test this, we compute the fraction of shares executed by HFLPs in various deciles of the parent-order, e.g., HFLPs' share of each 1000 shares for a parent-order consisting of 10,000 shares.

Figure 1 (top left chart) illustrates that the fractions of HFLP fills have roughly U-shaped pattern across execution deciles. We observe that HFLP trades occur substantially in the early stages of the execution as well. We also note high share of HFLP fills in the last few deciles of the execution as well. This figure suggests that the broker indeed routes orders to the HFLPs early in the execution and this may lead to significant information leakage.

Figure 1 (top right chart) shows the share of lit-venue executions across different volume deciles and we now observe an inverse U-shaped distribution implying the negative correlation between HFLP and lit-venue executions. Specifically, we observe that additional HFLP fills in the first and tenth deciles are offset by the decreases in lit-venue trades. Further, we find that the correlation between HFLP and lit-venue trades in the first and tenth deciles are very high in magnitude, -50% and -54%, respectively. The bottom panel in Figure 1 illustrates that dark pool trades in the broker are roughly uniform and smooth with a slight increasing trend with higher deciles. These statistics along with the visual evidence underscore that HFLP executions are borrowed from the share of lit venues.

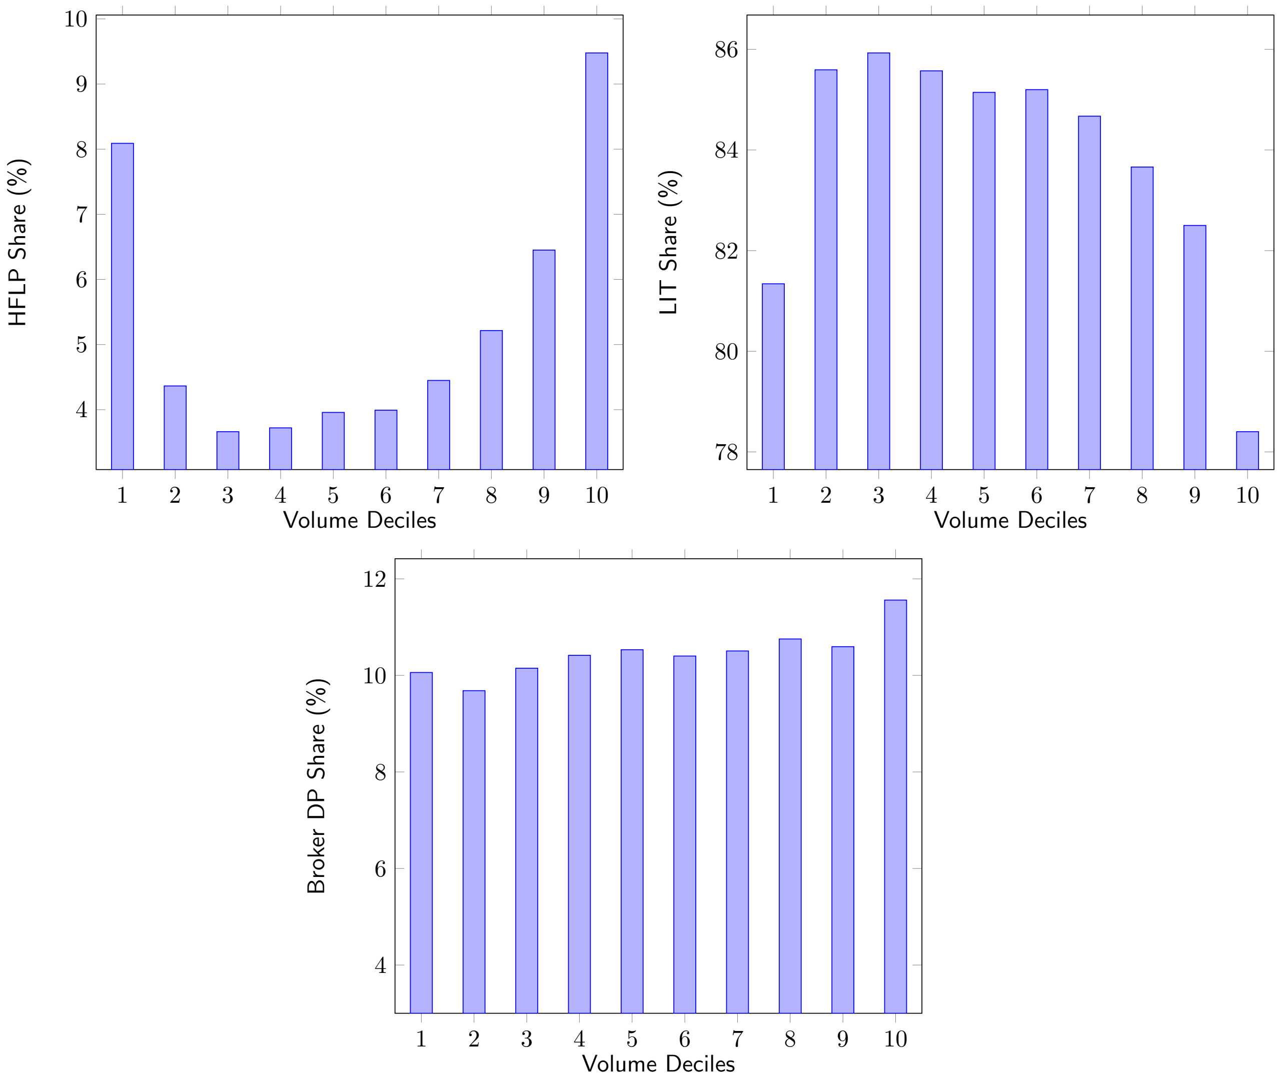

Figure 1: Shares of HFLP (top left), lit-venue (top right) and the broker's dark pool (bottom) executions in various deciles of executed volume, e.g., HFLPs' share of each 1000 shares for a parent-order consisting of 10,000 shares.

5.2. Client-Level Instructions and HFLP Trades

We do not have the pre-trade instructions of the investors so we cannot exactly identify the set of clients who disallows the broker to route their child-orders to HFLPs. However, we can infer this group of investors with a simple procedure. Analyzing the HFLP trades at the investor level, we find that 4 out of 146 investors have no child-order trades executed by an HFLP. These clients have individually 280, 228, 62 and 40 parent-orders consisting of 40,098, 19,997, 1,153 and 2,397 child-order trades, respectively. These 610 orders are executed on 224 distinct stocks on 217 different dates so it is almost certain that these investors have voluntarily opted out of HFLP trades. Further, except the client with 2,397 child trades, three of these clients have also opted out of dark pool executions implying a strategic desire to trade in lit venues.

Similarly, we have two additional clients who have only one child-order trade from an HFLP. These two clients have individually 180 and 166 parent-orders consisting of 64,251 and 13,136 trades. For the former investor, we also find that the HFLP fill occurs on the first parent-order of this investor and there is no HFLP fill in the remaining 179 parent-orders. Collectively, for these two clients, the ratio between HFLP-exposed parent-orders to the number of total parent-orders is only 0.6%. For the remaining 140 clients, this ratio is greater than 12% with a corresponding mean (median) of 61% (65%). Thus, HFLP exposure of these two clients is abnormally low. Consequently, we will also label their 344 parent-orders (excluding the two parent-order with HFLP trades) as *NotExposed*. Thus, in total, we obtain 954 parent-orders in the *NotExposed* group and 19,379 parent-orders in the *Exposed* group.

We now examine the differences between these two groups of parent-orders. First, we study the shares of lit venue and the BODP executions of the *NotExposed* group as a function of volume deciles. Figure 2 confirms our earlier conjecture about the substitution between HFLP trades and lit-venue trades when we compare these plots to the ones in Figure 1. When a parent-order is not exposed to HFLPs, we observe the first 9 volume deciles have roughly uniform shares of lit venue executions suggesting that in the presence of client's aversion to HFLPs, earlier trades are routed to lit exchanges.

Next, we look at the parent-order statistics for the *NotExposed* and *Exposed* data sets. Table 4 provides the order characteristics of the each client in *NotExposed* group and compares the average characteristics across two data sets. We specifically check participation rate, interval spread, volatility and turnover, dollar-based share of the *BODP* executions, and order duration. Out of these statistics, participation rate, *BODP* executions and order duration are statistically significant at 5% level (after clustering at the stock-day level). Given that three clients do not also have any dark pool executions, lower *BODP* execution rate is not surprising. Overall, the differences in other

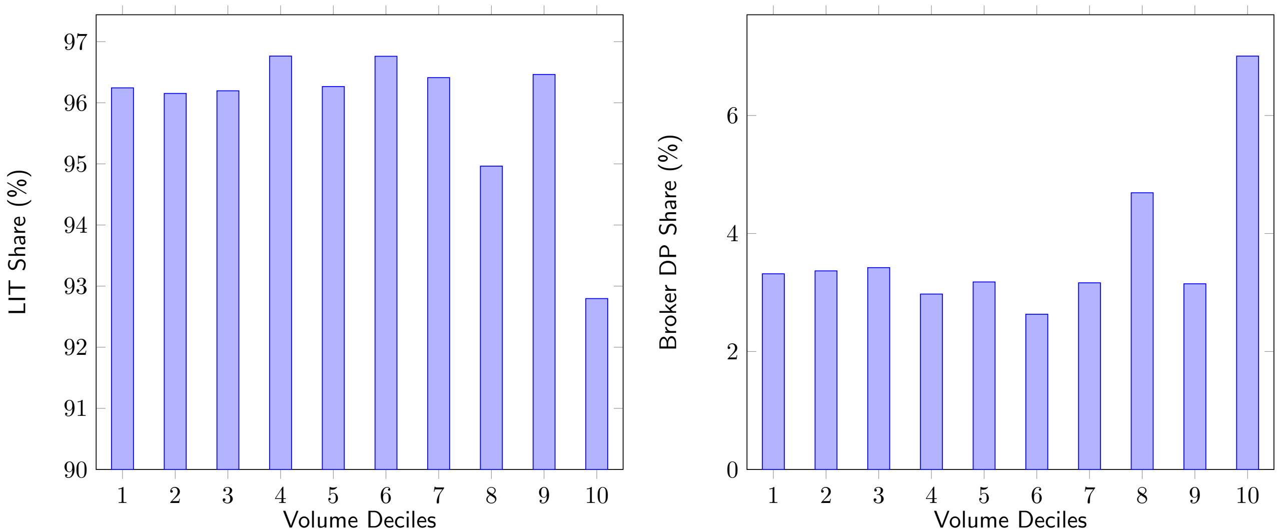

Figure 2: Shares of lit-venue (left) and the broker's dark pool (bottom) executions in various deciles of executed volume in the *NotExposed* group.

characteristics including participation rate and order duration are not very large. It is worth to highlight that there is no significant difference across two sets of executions with regards to the level of broad liquidity proxies given by spreads, volatilities and turnover. The insignificance in these implies that the investors in the *NotExposed* data set do not seem to time high-liquidity periods.

5.3. Relative Tick Size and HFLP Trades

Routing relationships may be very useful for HFLPs if the underlying stock has a relatively lower price. For these stocks, the queues at the best prices may be very large and if the HFLP does not have time-priority, it will not have execution priority. Consequently, routing relationship with the broker may allow the HFLPs to gain execution priority by offering price improvement to the routed order and the HFLP can make almost all of the half-spread. Further, these relationships will be helpful for the HFLP to manage its inventory more effectively as predicted in Aït-Sahalia and Sağlam (2013).

In our sample, Citigroup Inc. (NYSE:C) undergoes a 1-for-10 reverse stock-split on May 9, 2011 and its stock price jumps to roughly \$45 from \$4.5. This split can be useful for us to test the relationship between HFLP trades and nominal stock price visually. We have 109 parent-orders on

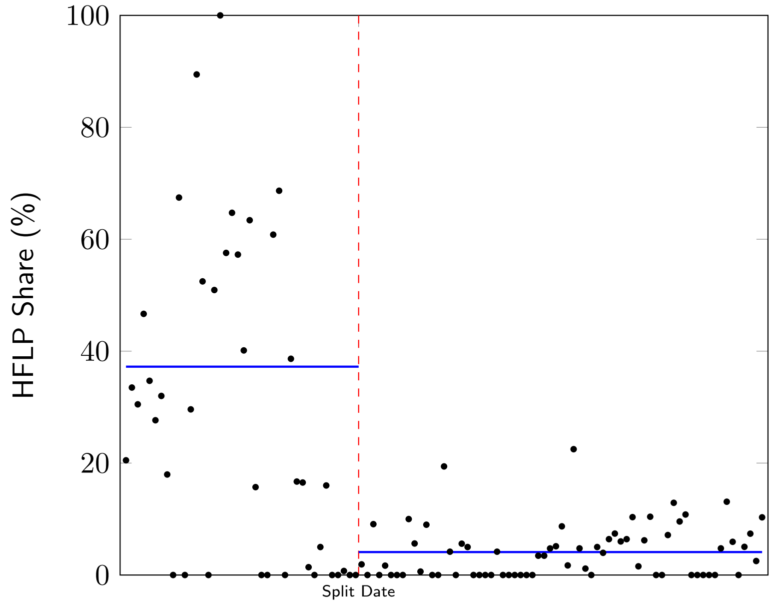

Figure 3: Shares of HFLP fills on Citigroup Inc. (NYSE:C) before and after the 1-for-10 reverse stock-split on May 9, 2011. We plot the average share of HFLP fills before and after the event.

this stock submitted by investors in the *Exposed* group. 40 of these are before the reverse split and the remaining ones are after the event. Figure 3 provides striking evidence that after the reverse split, the ratio of HFLP trades drop dramatically from 37% to 4%.

6. The Cost of HFLP Executions

In this section, we study the relationship between the proxies of execution costs and HFLP trading activity. First, we formally define our cost metrics and then test the relationship between the two in univariate and multivariate analyses.

6.1. Effective Spread (*ES*)

The first measure we consider is the percent effective spread. For the ith parent-order, the jth child-order trade has the following effective spread:

$$ESChild_{i,j} = \operatorname{sgn}(Q_i)\,\frac{P_{i,j} - M_{i,j}}{M_{i,0}},$$

where Q_i is the order size (in shares) with $Q_i > 0$ ($Q_i < 0$) for buy (sell) orders, $P_{i,j}$ is the execution price, $M_{i,j}$ is the mid-point of the best available quotes at the same second of the fill (i.e., NBBO mid-quote price) and $M_{i,0}$ is the mid-quote price of the security (arrival price) when the parent order starts being executed.

We define the following cost measure at the parent-order level by weighting each effective spread at the child-order level with its corresponding shares, i.e.,

$$ES_i = \frac{\sum_{j=1}^{N_i} ESChild_{i,j} Q_{i,j}}{Q_i},$$

where $Q_{i,j}$ is the executed number of shares in the jth child-order.

6.2. Implementation Shortfall (*IS*)

Perold (1988) introduced this measure to quantify the difference between the performance of a theoretical and the implemented portfolio. Over the years, *IS* has been extensively used as a proxy for institutional trading cost (Anand et al., 2011, 2013). It is computed as the normalized difference between the average execution price and the price of the asset prior to the start of the execution. Formally, the *IS* of the ith parent-order is given by

$$IS_i = \operatorname{sgn}(Q_i)\,\frac{\left(\frac{1}{Q_i}\sum_{j=1}^{N_i} P_{i,j} Q_{i,j}\right) - M_{i,0}}{M_{i,0}}. \tag{1}$$

Note that we can decompose *IS* into two terms where the first term is *ES* and the second term is a price impact term:

$$
\begin{aligned}
IS_i &= \operatorname{sgn}(Q_i)\,\frac{\left(\frac{1}{Q_i}\sum_{j=1}^{N_i} P_{i,j}Q_{i,j}\right) - M_{i,0}}{M_{i,0}} \\
&= \operatorname{sgn}(Q_i)\,\frac{\sum_{j=1}^{N_i}\left(\frac{P_{i,j}-M_{i,0}}{M_{i,0}}\right)Q_{i,j}}{Q_i} \\
&= \operatorname{sgn}(Q_i)\sum_{j=1}^{N_i}\left(\frac{P_{i,j}-M_{i,j}+M_{i,j}-M_{i,0}}{M_{i,0}}\right)\frac{Q_{i,j}}{Q_i} \\
&= \underbrace{\operatorname{sgn}(Q_i)\sum_{j=1}^{N_i}\left(\frac{P_{i,j}-M_{i,j}}{M_{i,0}}\right)\frac{Q_{i,j}}{Q_i}}_{\triangleq ES_i} + \underbrace{\operatorname{sgn}(Q_i)\sum_{j=1}^{N_i}\left(\frac{M_{i,j}-M_{i,0}}{M_{i,0}}\right)\frac{Q_{i,j}}{Q_i}}_{\triangleq PI_i}.
\end{aligned}
\tag{2}
$$

This decomposition will be useful for us to determine the major driver of the impact of HFLP executions.

6.3. Summary Statistics on Execution Cost Proxies

Table 5 provide the summary statistics for the execution cost metrics. The mean ES is 0.68 bps in the dataset whereas the mean IS is 3.12 bps in the dataset. Thus, ES (PI), on average, constitutes approximately 22% (78%) of the total IS. We note that PI and IS are very noisy whereas ES has relatively lower standard deviation. These average statistics will be useful for us to understand the economic impact of HFLP executions.

6.4. Univariate Analysis: Effective Spreads across Venues

We first utilize the complete child-order data set and compute the percent effective spread of each child order, *ESChild*. Table 6 illustrates the average *ESChild* for all of the venues in the data set with more than 200 child-order trades. We also provide the corresponding standard errors. Recall that Knight operates two different venues, TRIM and NITE, thus, we have seven HFLP venues. Strikingly, the top five venues with largest average effective spreads are all HFLPs. The bottom five venues with respect to this statistic consists of four lit venues and one dark pool. Overall, this univariate evidence suggests that HFLPs do not seem to provide better execution quality. Note that this analysis does not control for the *complexity* of the underlying child-order but provide unconditional average effective spreads for each venue.

One important characteristic of HFLP fills is the lack of anonymity due to direct routing of the order. Thus, one would expect that successive child-order trades in the HFLPs would lead to more information leakage. To test this hypothesis, we compute the average effective spread of the successive child-order fills at the venue level. Our conjecture is that compared to a lit venue, the effective spread associated with successive fills at the HFLP would be much larger. Figure 4 illustrates a visual evidence for this conjecture using three different venue types: HFLP, dark pool and a lit venue. We chose GFLO, *BODP* and BATS, respectively, for these venue types as these markets have roughly similar number of trades in terms of order of magnitude (GFLO: 108 thousand, *BODP*: 166 thousand, and BATS: 157 thousand). We observe that HFLP trades occur at higher spread initially and increase rapidly compared to the dark pool and the lit venue. Surprisingly, the lit venue has a pretty stable spread dynamics. We also fitted best fit lines to these observations and computed the slopes to be 0.014, 0.007 and -0.002. All of the slope coefficients are statistically significant.[9]

6.5. Information Leakage to HFLPs and Execution Costs

We use three proxies to measure the information leakage to HFLPs at the parent-order level. The first measure, *HFLPsh*, is based on the ratio between the number of shares executed by the HFLPs and the total shares executed. The second measure, *HFLPdol*, is based on the ratio between the dollar volume executed by the HFLPs and the total dollar volume. The third measure, *HFLPtr*, equals the fraction of the number of trades executed by the HFLPs compared to the total number of child-orders of the parent-order.

Execution costs can be a function of multiple trade-level and stock-level characteristics, thus, to test formally whether each HFLP proxy is associated with higher costs, we run the following

[9]Further, average trade size of fills by GFLO is the smallest which makes these findings more striking.

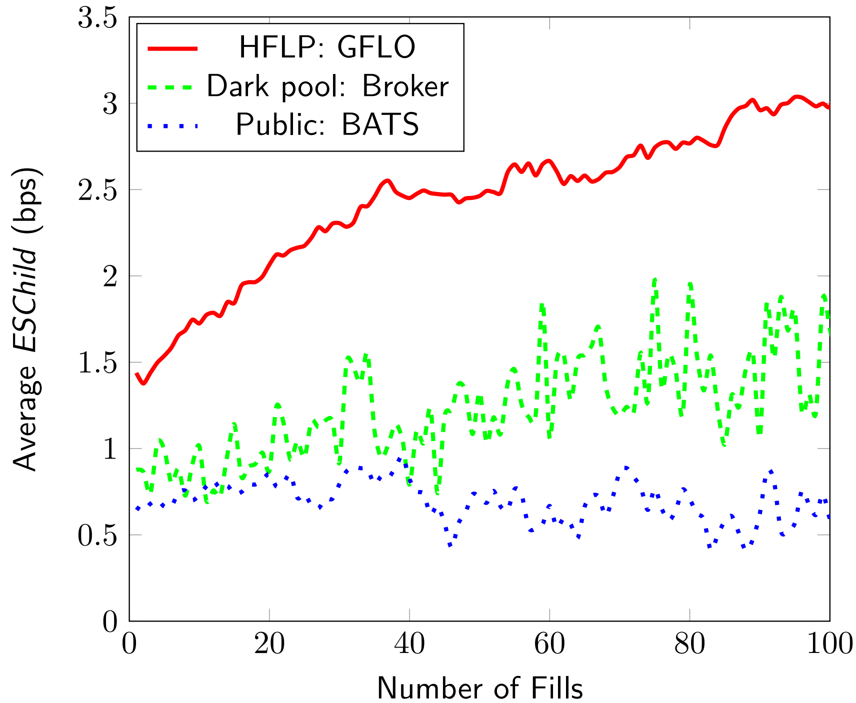

Figure 4: Effective spread of successive fills in different markets.

multivariate regression at the execution level with a rich set of control variables:

$$Cost_i = \alpha + \beta HFLPproxy_i + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i, \tag{3}$$

where *Cost* is either *ES*, *PI*, *IS* or *VS*, *HFLPproxy* is either *HasHFLP*, *HFLPsh* or *HFLPtr*. The mapping $i \xrightarrow{m} s$ is used to identify the executed stock s, the mapping $i \xrightarrow{c} k$ is used to identify the client k submitting the order, and the mapping $i \xrightarrow{d} t$ is used to identify the trading day. In addition to these stock, client and calendar day fixed-effects, we consider execution-level control variables including the ratio of *BODP* and other dark pool executions, participation rate, logarithm of the arrival price, mid-quote volatility, execution duration, and turnover.

For dark venue executions, we will use a control variable consistent with the *HFLPproxy* to better interpret the β coefficient. For example, when *HFLPproxy* is equal to *HFLPsh*, we will use *BODPsh* and *OthrDPsh* as control variables and the resulting β coefficient would measure the cost of increasing the share of HFLPs by 1% of total shares while keeping the share of the dark pools the same and decreasing the share of lit venues by the same 1%. Thus, in all proxies, our benchmark will be based on the cost of executing the same amount in the lit venues. This is the best benchmark to consider as we have already analyzed in the prior section that the distribution of HFLP trades across volume deciles are inversely correlated with that of lit venues (recall the U-shaped and inverse U-shaped patterns in Figure 1).

Participation rate and order duration can control for the urgency of the trade and client-fixed effects can control over the different trading strategies or the skill level of the investor that may be correlated with the price movements during the execution. Controlling for the ratio of dark venue executions is also important as clients can avoid dark pools or the algorithm may be strategic in routing the child orders between dark venues and the HFLPs. Finally, price level and volatility may also affect the execution rates of the HFLPs and can affect the total cost of the order, thus we also include them as controls. This set of control variables is also consistent with the prior literature in examining *IS*. For example, Almgren et al. (2005) includes participation rate and volatility in analyzing the variation in *IS*. van Kervel and Menkveld (2017) uses order duration, volatility, turnover, client and stock fixed effects.

Finally, in Table 7, we report the correlation matrix of the independent variables. The highest correlation in terms of magnitude is between *Duration* and Turnover with 0.46. We find that *HFLPsh* is negatively correlated with dark pool executions. As we have seen with the example on the reverse split of the Citigroup, *HFLPsh* is negatively correlated with the price level. We observe that all of the correlations are significant at 10% except the relationship between *LogPrice* and *Duration*.

Table 8 reports the regression results with standard errors clustered at the stock and calendar day level. For each proxy of execution costs and HFLP trading activity (12 cases in total), we observe that the estimated β coefficient is positive and highly statistically significant. Examining the value of the coefficients across three proxies of HFLP trading activity, we find that share- and dollar-based measures are almost identical. Based on number of trades, β coefficient goes down slightly. This makes sense as according to that metric HFLP activity is also slightly higher. Examining the value of the coefficients across four proxies of execution costs, we observe that *PI* and *IS* provide very similar coefficients.

The coefficients are also economically large. First, comparing the β coefficients with the coefficients on participation rate, we observe that they are very similar in terms of magnitude. This suggests that exposure to HFLPs seem as important as the size of the order. To understand the exact economic impact, consider the following example. Suppose that in the benchmark case, 100% of an execution is traded in lit venues and now you consider moving *HFLPsh* from 0% to 5.6% (average value of *HFLPsh* in the *Exposed* group). In this hypothetical example, compared to the benchmark case, *ES* would increase by 0.1 bps, and *PI* and *IS* would increase by 2.5 bps. All of these increases are substantial considering the summary statistics provided in Table 5.

Our regressions also analyze the relationship between other dark pool trading activity and execution cost proxies. Compared to the strong pattern we identified with HFLP fills, we do not find a consistent pattern for the impact of trades in the *BODP* and ODP. We observe that in the case of *PI* and *IS*, the coefficients on the share of *BODP* and ODP executions are negative implying that they reduce the price impact of the execution. Compared to our findings on HFLP transactions, this asymmetry is interesting as all of these trades, HFLP, *BODP* or ODP, have the same exchange code in the TAQ (i.e., 'D'). Overall, these different signs in the coefficients highlight the potential heterogeneity of the effects of dark pool executions and call for further research. For example, there are a few studies that use TAQ classification to examine the impact of dark pool trades on market quality (e.g., O'Hara and Ye (2011), Hatheway et al. (2017) and Farley et al. (2018)) but the TAQ data cannot differentiate between these two different types of fills. One takeaway from this analysis is that the potential positive effect of dark pools is biased downwards if one classifies the dark pool trades from the TAQ database.

6.6. Understanding HFLP Profits

In this section, we compute a lower bound on the aggregate profits of the HFLPs assuming that they follow a simple trading strategy. Since we observe the transaction price for their trades along with the number of shares executed, we just need to assume a trading strategy for how they offload this inventory. Consider the following simple offloading strategy. Suppose that the HFLPs buy (sell) the same amount of shares they sold (bought) to (from) the broker at the market VWAP price realized over the 30 seconds after the transaction with the broker. This is a lower bound given the universe of available trading strategies and the HFLP's advanced technology.

Formally, let $Profit_i$ be the gain or loss of the HFLPs for the ith execution following the simple offloading strategy.

$$Profit_i = \operatorname{sgn}(Q_i) \sum_{j=1}^{N_i} \mathbb{I}_{\{j \in \text{HFLP}\}} Q_{i,j} \left(P_{i,j} - VWAP_{i,j+30s}\right), \tag{4}$$

where $\mathbb{I}_{\{j \in \text{HFLP}\}}$ denotes whether the jth child-order trade is filled by an HFLP and $VWAP_{i,j+30s}$ is the volume-weighted average price realized over the 30 seconds after the jth child-order trade. Note that this analysis is very similar to computing the relative spread of each child-order trade which ultimately measures the profits to the market-maker.

We find that the HFLPs make $8.7 on average per parent-order execution in which they have at least 1 child-order fill. HFLPs make positive profits in roughly 75% of the parent-order executions using this simple strategy. We then aggregate their profits at the calendar day level by summing their profits across executions. During the 15-month horizon of the data set consisting of more than 300 trading days, HFLPs lose money only in 7 business days. The average daily profit is $362.1 with a corresponding standard deviation of $570.0. The skewness of the daily profit and loss distribution is positive with $7.03.

We also applied the same strategy for child-orders executed in BATS exchange which is a lit venue. BATS is a good benchmark case as the number of trades is similar in terms of order of magnitude.[10] In aggregate, using the same strategy as above, market-makers in BATS make $5.7 on average per parent-order execution in which they have at least 1 child-order fill. They make positive profits in only 58% of the parent-orders. Aggregating their profits to the calendar day level, we find that they lose money in 29 business days. Comparing the daily profits of HFLPs and

[10]This exchange is also referred to be HFT friendly in the popular media.

BATS market-makers, we find that HFLPs make \$127.0 (t-stat = 3.6) additional per day. Overall, these statistics imply that HFLPs can benefit significantly from the routing relationships.

7. Robustness Tests

7.1. Controlling for Informed Trading

One important concern in the OLS regression in Equation (3) is that if the algorithm is routing informed orders to HFLPs, then there would be a mechanical positive relationship between *IS* and the volume of HFLP fills. Our main regression includes urgency, order size and client fixed-effects to control for the private information of the client. The clients may have execution-specific information, i.e., some orders from the same client may be informed whereas the remaining ones could be due to liquidity reasons. If the broker is successful in identifying the information content of these orders in early stages of the regression and then intentionally routes the informed orders to HFLPs, we may mechanically obtain positive correlation with executions costs and HFLP activity. We address this concern by utilizing a proxy for the permanent price impact of the order.

Motivated by the standard spread decomposition into realized spread (transitory) and adverse selection (permanent), van Kervel and Menkveld (2017) propose a simple measure to compute permanent price impact of a large order. In this setting, they compute the permanent price impact (PPI) of an execution by comparing the arrival mid-quote to the close price of the stock at the end of next trading day. Following this approach, we define the PPI of ith execution as follows:

$$PPI_i = \operatorname{sgn}(Q_i)\,\frac{X_{m(i),d(i)+1} - M_{i,0}}{M_{i,0}}, \tag{5}$$

where the mapping $i \overset{d}{\to} u$ is used to identify the date of the execution and $X_{j,u}$ is the close price of the asset j on day u.

By construction, *PPI* includes the execution period and has high correlation with *PI* and *IS*. We find that at the univariate level, the correlation between *PPI* and *IS* (*PI*) is 42.7% (42.8%). *PPI*'s correlation with *ES* is at −4.0%.

To examine whether PPI explains the significant regression coefficient on HFLP proxies, we include PPI as an additional control variable in our main regression:

$$Cost_i = \alpha + \beta HFLPproxy_i + \theta PPI_i + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i, \tag{6}$$

where *Cost* is either *ES*, *PI* or *IS* as in the main regression.

Table 9 reports the regression results for each proxy of HFLP participation. For every HFLP proxy and cost proxy, we again find that the estimated coefficient is positive and highly statistically significant. When the cost metric is *ES*, the coefficients on HFLP proxies increase slightly whereas for the remaining cost proxies, the coefficients decrease by roughly 20%. With regards to the share- or dollar-weighted HFLP proxy, a 1% increase in HFLP fills (with a corresponding 1% decrease in lit venue executions) results in additional *IS* of 0.38 bps at the minimum which is approximately 12% of the average *IS*. Overall, this test highlights the robustness of our results with respect to the strategic routing of informed orders to the HFLPs.

7.2. VWAP Slippage (*VS*)

IS and *PI* use an *ex ante* static benchmark and does not consider any permanent price changes during the execution period. For this reason, one popular *ex post* benchmark for the average execution price is the volume-weighted average price (VWAP) during the parent-order's trading interval. Formally, the VWAP during the ith parent-order is defined as follows. Suppose that there are K_i trades during this period. Given the sequence of trades at prices $P_1, \ldots, P_{K_i}$ with corresponding shares of $V_1, \ldots, V_{K_i}$, the VWAP is given by

$$VWAP_i = \frac{\sum_{k=1}^{K_i} P_k V_k}{\sum_{k=1}^{K_i} V_k}. \tag{7}$$

Using this definition, the VWAP slippage for the ith parent-order is given by

$$\mathrm{VS}_i = \mathrm{sgn}\left(Q_i\right) \frac{\left(\frac{1}{Q_i}\sum_{j=1}^{N_i} P_{i,j}Q_{i,j}\right) - VWAP_i}{VWAP_i}. \tag{8}$$

The average (median) *VS* in the data set is 1.56 (1.13) bps. We find that at the univariate level, the correlation between *VS* and *IS* (*PI*) is −6.0% (−6.3%). On the contrary, *VS* is positively correlated with *ES* at 6.3%. Overall, these numbers point that *VS* may measure a different dimension of execution costs.

To investigate the robustness of our findings with regards to another cost proxy, we run our main regression in which *VS* is the dependent variable.

$$VS_i = \alpha + \beta HFLPproxy_i + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i. \tag{9}$$

Table 10 reports the regression results for each proxy of HFLP participation. For every HFLP proxy, we again find that the estimated coefficient is positive and highly statistically significant. With regards to the share- or dollar-weighted HFLP proxy, a 10% increase in HFLP trading activity results in additional VWAP slippage of 0.56 bps. Given that the mean *VS* is 1.56 bps, these coefficients are economically large. Overall, this test highlights the robustness of our results with respect to a different cost proxy.

7.3. Robustness to Excluding The Most Active HFLP

Recall that one HFLP (GFLO) has the highest share of HFLP fills. 59% of the HFLP dollar volume belongs to this single HFLP. In this section, we formally investigate whether our findings are driven by this single HFLP. We decompose each of our HFLP trading activity proxy into two components: shares of GFLO and non-GFLO trades. For example, we define *HFLPdolGF* and *HFLPdolOthr* where the sum of them equals *HFLPdol*.

To examine the robustness of our findings to the GFLO's active share in HFLP fills, we re-run our main specification with the decomposed variables:

$$
\begin{aligned}
Cost_i = \alpha + \beta_{GF} HFLPproxyGF_i + \beta_{Othr} HFLPproxyOthr_i \\
+ \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i.
\end{aligned} \tag{10}
$$

Table 11 reports the summary of the regression results. For each cost proxy, we again find that the estimated coefficients on both HFLP proxies, β_{GF} and β_{Othr}, are positive and highly statistically significant. This result implies that our findings are not driven by a specific HFLP. We find that in terms of magnitude, β_{GF} is slightly higher than β_{Othr} when the cost metric is *ES*. However, we observe the opposite is true for the price impact component. Thus, for *IS*, β_{GF} is substantially lower than β_{Othr} and this difference is statistically significant for each HFLP proxy.

8. Evidence from Clients' Aversion to HFLPs

Our findings in the previous section point to a strong positive correlation between proxies of HFLP trading activity and execution costs which may not directly imply causal relationship. There may be an omitted variable that is correlated with execution costs and HFLP trading activity which may bias the regression coefficients. In order to mitigate this concern, we pursue two identification strategies. First, we utilize the set of 954 executions from six investors analyzed earlier in Section 5.2. Second, we compare the cost of executions of a particular client before and after he seemingly disallows the broker to route orders to HFLPs.

8.1. Evidence from Clients without HFLP Exposure

Recall that Table 4 provides the order characteristics of the distinct six clients who have zero (C1–C4) or abnormally small exposure to HFLPs (C5–C6). First, we observe that each client has quite different order characteristics. For example, C6 has very large participation rate and relatively small duration which may lead to higher execution costs. On the opposite side, C2 has a relatively small participation rate but larger duration that can lead to lower trading costs. C1–C3 have zero exposure to *BODP* whereas C4–C6 have pretty substantial exposure. One takeaway from this heterogeneity is that there does not seem to be a common pattern across these six investors beside their zero exposure to HFLPs. However, since the clients endogenously choose not to expose their orders to HFLPs, we will employ an exact matching procedure to obtain a balanced sample of treated and control groups.

For each parent-order i in the *NotExposed* group, we search for a matched parent-order in the *Exposed* group (without replacement) using the following algorithm. First, the executed stock and the trade direction (buy or sell) must be the same. Second, the dates of the executions of the ith parent-order and the matched parent-order must be within two trading days. We screen the *Exposed* group using these criteria. If there is no match, ith parent-order will not be matched. If there are multiple matches, we use the parent-order with the closest number of shares to be executed. We were able to match 285 (out of 954) parent-orders using this exact procedure. We will label these groups of parent-orders, *NoHFLP* and HFLP, respectively. This exact matching procedure addresses the potential timing and stock-selection ability of the clients in the *NotExposed* group. By matching on the same stock and close trading days, the two groups mainly differ in their HFLP exposure.

Next, we compare the characteristics of the *NoHFLP* and *HFLP* matched samples. Table 12 provides the detailed comparison. We specifically check participation rate, interval spread, volatility and daily turnover, dollar-based share of the *BODP* executions, and order duration. We observe that the differences in spread and volatility are extremely small implying the success of the exact matching at the stock and calendar date level. We observe that executions in the *NoHFLP* group have larger participation rate and smaller duration which would typically lead to higher execution costs. Given that we conjecture lower execution costs for the *NoHFLP* group, this difference in urgency may bias our findings downwards in the univariate child-order level analyses.

We first perform our analysis at the child-order level trades. We have 38,768 (38,435) child-order trades in the *NoHFLP* (*HFLP*) group. Given the high-number of trades, we can compare the average cost trajectories of these groups of parent-orders. We compute the percent effective spread and price impact of each child order in the *HFLP* and *NoHFLP* groups and compute the median

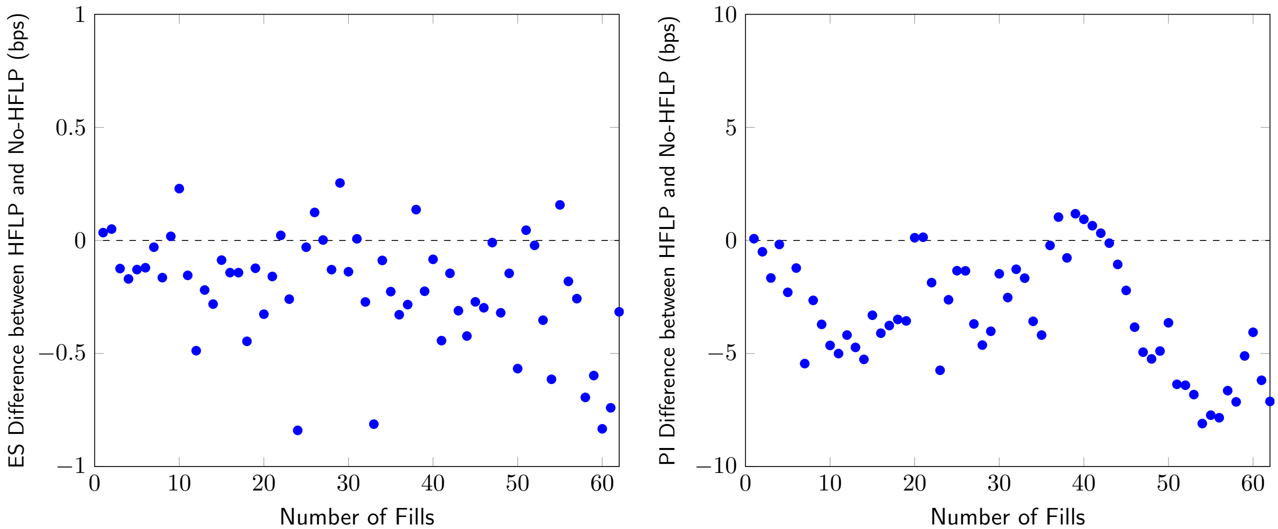

Figure 5: Difference between the effective spread and price impact at the child-order level. We plot $Cost(NoHFLP) - Cost(HFLP)$ for both effective spread and price impact. We plot the difference up to 62nd child-order after which the number of observations drop below 150 for both proxies.

cost corresponding to the jth child order. Figure 5 plots the difference in these cost measures across *HFLP* and *NoHFLP* groups. That is to say, we plot $Cost(NoHFLP) - Cost(HFLP)$ for both cost proxies. We plot the difference up to 62nd child-order after which the number of observations drop below 150 for both proxies. We find statistically significant evidence that the mean difference across these groups is negative implying that child orders that are exposed to HFLPs have worse execution. For effective spread, the mean difference is -0.22 bps (t-stat = -6.8) and for price impact, the mean difference is -3.25 bps (t-stat = -10.2) which are similar to the estimates obtained from the OLS regressions in the previous section. Recall that orders in the *NoHFLP* group were actually more urgent and would be expected to be costlier. Thus, this visual and statistical evidence highlight the cost effect of exposure to HFLPs.

8.2. Evidence from Switching to No Consent

In theory, clients should experiment with pre-trade instructions to maximize their execution performance. They can examine the reports on post-trade transaction cost analysis and may also detect the cost increase associated with HFLP fills.[11] By examining the clients' exposures to HFLPs over time, we find that there is one investor who seemingly consents for his orders to be routed to HFLPs in early 2011 but then does not have any exposure to HFLPs. This sharp switching implies that this client has disallowed the broker from routing the child-orders to the HFLPs. We feel confident about this inference because this client has no executed child-order from an HFLP in more than 40,000 child-order trades in more than hundred distinct parent-orders submitted on different stocks and dates. Furthermore, the client continues to trade in the BODP and other dark pools. This switching behavior with regards to consenting for HFLP fills is consistent with the cost increase due to HFLPs and provides us another opportunity for causal identification.[12]

Between January 2011 and March 2011 this client has 40 parent-orders (5,239 child-order trades) that has sporadic exposure to HFLPs. On March 18, 2011, in his last parent-order execution during this period, there is a child-order executed by an HFLP suggesting that the client consents for HFLP trades. During this period, 67.5% of the client's executions have HFLP exposure and 5.5% of the dollar volume is executed in BODP or other pools. The client's next execution after this period occurs on April 6, 2011 and starting from this execution, the client's child orders have never been routed to HFLPs in 111 distinct parent-orders consisting of 40,479 child-order trades. The client still continues to trade in the dark pools and 7.4% of the volume is executed in dark pools in this no-HFLP period. These statistics strongly imply that the client may have switched its instructions on HFLP routing after April 2011.

Given that the same client switches from active HFLP exposure to zero HFLP exposure, we can design a difference-in-difference framework to formally test the impact of HFLP exposure on execution costs. Let *Post* be a binary variable that takes a value of 1 if the execution occurs after April 2011 and *Switched* be an indicator variable that is equal to 1 for executions of the client who starts having zero exposure to HFLPs after April 2011. In the control group, we can use the parent-order executions in the *Exposed* group that have HFLP exposure throughout the data period. We then run the following diff-in-diff regression using the same set of control variables and

[11] All of the clients receive a detailed report from the broker on execution venue analysis.

[12] Consistently, there is no client who has the opposite switching behavior in the data.

stock, day and client fixed effects as in Equation (3):

$$Cost_i = \alpha + \beta Post_i \times Switched_i + \kappa Post_i + \upsilon Switched_i + \sum_j \delta_j Control_{j,i} \tag{11}$$
$$+ \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i,$$

where *Cost* is either *ES*, *PI*, *IS* or *VS*. In this framework, β measures the impact of eliminating the exposure to HFLPs on execution costs.

Table 13 reports the regression results with our main execution cost metrics. We find that β is negative and statistically significant for *PI* and *IS*. These findings provide further causal support for the cost increase in executions with exposure to HFLPs.

9. How do HFLPs Increase Transaction Costs?

In this section, we further explore the underlying mechanisms that can explain the higher execution costs with HFLP-exposed parent-orders. We also evaluate our findings in the light of the predictions of the theoretical models described in Section 3.

9.1. Do HFLP Trades Substitute Passive Orders?

We first compare the transaction prices of each child order to the prevailing mid-quote price and examine the relationship between the aggressiveness of the parent-order and HFLP trades.

First, we classify each child order we define a buy (sell) child-order trade as an aggressive order if transaction price occurs at a price higher (lower) than the mid-point of the National Best Bid and Offer (NBBO) realized during the same second. Similarly, passive orders correspond to child order executions if the transaction price is lower (higher) than the NBBO mid-point when the investor has a large buy (sell) order. Finally, mid-point fills are transactions happening at the mid-point of the NBBO. We then create the corresponding parent-order level statistics, passive order ratio (PO), aggressive order (AO) ratio and mid-point order (MO) ratio by computing the fraction of the shares executed via each order type.

We regress the aggregate fraction of orders filled via aggressive, mid-point, and passive orders at the parent-order level on HFLP trading activity and our standard set of execution controls and fixed-effects. Formally, we run the following regression:

$$FracOrderType_i = \alpha + \beta HFLPproxy_i + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i,$$

where *FracOrderType* is either *FracAO*, *FracMO* or *FracPO* and *HFLPproxy* is *HFLPsh*. The findings are similar with other HFLP proxies.

Table 14 reports the regression results. We find that when 10% of the share volume goes from lit venues to HFLPs, the fraction of aggressive fills increases by 3.7% and the fraction of passive (mid-point) fills decrease by 2.9% (0.8%). These statistics suggest that increase in HFLP trades is associated with decrease in passive fills, which usually occur in lit venues. This findings implies that HFLP trades do not substitute one-for-one with aggressive lit venue trades.

9.2. Early versus Late HFLP Trades

In Figure 1, we observe the U-shaped pattern for HFLP fills. Given the positive relationship with price impact proxies and HFLP participation, we expect that early exposure to HFLPs would lead to higher execution costs (especially higher price impact). This finding would be consistent with the information leakage channel. To investigate this explanation formally, we decompose our HFLP trading activity proxy into three components: share of HFLP trading activity in the first decile, tenth decile and the remaining deciles. For example, we define *HFLPshD1*, *HFLPshD10* and *HFLPshDRest* where the sum of them again equals *HFLPsh*.

To test this channel, we re-run our main specification with the decomposed three variables:

$$Cost_i = \alpha + \beta_{\text{Early}} HFLPproxyD1_i + \beta_{\text{Late}} HFLPproxyD10_i + \beta_{\text{Rest}} HFLPproxyDRest_i + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i. \quad (12)$$

We expect that $\beta_{\text{Early}} > \beta_{\text{Late}}$ for each cost metric if there is higher information leakage with regards to HFLPs.

Table 15 reports the regression results when we use the share-based HFLP trading activity. The findings are identical with the other HFLP proxies. We find that our conjecture holds in all cases. Using a Wald test, we confirm that β_{Early} is statistically greater than β_{Late} for all of the cost metrics. Consistent with our hypothesis, we also observe that $\beta_{\text{Rest}} > \beta_{\text{Late}}$ with statistical significance. Overall, these findings point to higher information leakage with HFLP fills.

9.3. Consistency with Back-running and Predatory Trading Theories

Given that we do not observe the trading of the HFLPs in other venues outside of our data set, we cannot exactly test the back-running or predatory trading theories. However, it is still informative to examine the predictions of these models assuming that their activities in our data set are representative of their aggregate activity.

Back-running theory would predict that HFLPs would not execute orders arriving from informed traders especially at the later stages of the execution. Here, the idea is that they would share the informational rents of the informed investor by trading in the same direction with the large-order. Thus, the same-side trading motive in later stages should show up in our dataset in the the form of sharp drop in HFLP fills. Building on this intuition, we analyze whether our empirical findings are consistent with this theory.

First, we would expect that the share of HFLP fills to drop in the later stages of the execution if back-running theory were to hold. To test this simple insight, we compare the HFLP participation across the first and second half of the execution in terms of the volume executed. We use the share-weighted measure as other measures also give identical results. We let *HFLPshH1* and *HFLPshH2* be the share-weighted ratio of HFLP-executed shares in the first and second half, respectively, and compute $\mathit{HFLPDiff} = \mathit{HFLPshH1} - \mathit{HFLPshH2}$. At the univariate level, we find that the average *HFLPDiff* is 1.35% implying higher participation in the second-half. This difference is also evident in Figure 1. Overall, higher participation in the second-half aligns better with predatory trading.

We also examine the sensitivity of late-stage HFLP participation with regards to the level of informed trading proxy, *PPI*. Back-running theory would suggest negative relationship between late-stage HFLP trading and *PPI* as HFLPs would be exploiting the informative signal by trading in the same direction of the parent-order. Using our standard regression framework, we regress the proxies of late-stage HFLP participation on *PPI* and various controls and fixed-effects:

$$\begin{aligned} HFLPproxy_i = {} & \alpha + \beta PPI_i + \sum_j \delta_j \, Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} \\ & + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i, \end{aligned} \tag{13}$$

where *HFLPproxy* is either *HFLPshH2* or *FracHFLPD10*.

Table 16 reports the regression results. We find that HFLP trading in the second half or in the tenth decile is not significantly correlated with *PPI*. The signs of coefficients are also mixed. Overall, our findings do not strongly align with the back-running theory.

If HFLPs follow predatory trading motives, one would expect this additional cost to be transitory. We can test this hypothesis by computing the transitory component of the effective spread, e.g., realized spread. We can compute the normalized realized spread (RS) at the child-order level and then aggregate it to the parent-order level:

$$RSChild_{i,j} = \operatorname{sgn}(Q_i) \frac{P_{i,j} - VWAP_{i,j+30s}}{M_{i,0}}, \qquad RS_i = \frac{\sum_{j=1}^{N_i} RSChild_{i,j} Q_{i,j}}{Q_i}. \tag{14}$$

To investigate the correlation between *RS* and HFLP trading proxies, we run our main regression in which *RS* is the dependent variable.

$$\begin{aligned} RS_i = {} & \alpha + \beta HFLPproxy_i + \sum_j \delta_j \, Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} \\ & + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i. \end{aligned} \tag{15}$$

Table 17 reports the regression results for each proxy of HFLP participation. For every HFLP proxy, we find that the estimated coefficient is positive and highly statistically significant. Comparing the estimated coefficients to those of *ES* in Table 8, we find that more than 50% of the cost is transitory. Overall, this test highlights the positive correlation between transitory execution costs and HFLP trading activity.

9.4. Competition between HFLPs

Bessembinder et al. (2016) extends the model of Brunnermeier and Pedersen (2005) for resilient markets in which the immediate price impact of trades may be transitory. In this model, in addition to the same-side trading before the liquidation, the strategic traders trade in the opposite direction compared to the direction of the parent-order. This theory illustrates that this opposite-side trading can decrease the liquidator's transitory price impact. This benefit to the liquidator from strategic trading persists at any level of market resiliency if there are multiple strategic traders. Given that there are six distinct HFLPs in the data, if the broker routes orders to multiple HFLPs, this theory would suggest a potential decrease in execution costs compared to the case where there is a single HFLP taking advantage of the order flow information.

Recall that 61.1% of all parent-orders has at least one child order executed by an HFLP. We compute the breakdown of this statistic with regards to distinct HFLPs. We find that 25.3% has only one HFLP, 21.4% has two distinct HFLPs, 11.7% has three distinct HFLPs, 2.4% has four distinct HFLPs and 0.25% has five distinct HFLPs executing trades in that parent-order. There is no execution with all of the HFLPs participating.

To test the impact of the competition between HFLPs, we run the following regression:

$$\begin{aligned} Cost_i = \alpha + \beta_1 HFLPsh_i + \beta_2 Has2Plus_i + \beta_3 Has3Plus_i + \beta_4 Has4Plus_i + \beta_5 Has5Plus_i \\ + \sum_j \delta_j Control_{j,i} + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i, \end{aligned} \tag{16}$$

where *HasXPlus* denotes the binary variables that takes a value of 1 if the execution has X or more distinct HFLPs executing a trade. Here, we are controlling for the share of the HFLP trades and test for the incremental effect of routing to multiple HFLPs.

Table 18 reports the regression results. We find that 10 out of 12 coefficients are positive. Specifically, the coefficients on *Has2Plus* and *Has3Plus* are positive and significant. Hence, these findings suggest that after controlling for the ratio of HFLP fills, having additional competing HFLPs does not decrease execution costs. These findings are instead more consistent with higher information leakage when the orders are routed to multiple HFLPs.

10. Conclusion

This paper examines the costs of routing relationships between a broker and high-frequency liquidity providers (HFLPs) from the perspective of the investors submitting the orders to the broker. In this arrangement, the broker is either directly receiving a payment for the routed and executed orders from the HFLPs or has the opportunity to execute the trade without any commissions or fees. The brokers benefit from this relationship as the rebates and fees realized during the execution are not passed to the investors. The HFLPs can profit from this relationship by making the spread between the transaction price and fair value and they also gain valuable order flow information from the broker which is potentially less noisier compared to the market data.

Analyzing a large order execution data with such relationship, we obtain robust evidence that executions that are routed to HFLPs suffer higher transaction costs using different measures of HFLP trading activity. This finding is robust to controlling for various order-level characteristics including informed trading proxy, stock, day and investor fixed-effects. In the most conservative case, a 1% increase in HFLP participation rate in the execution corresponds to 3.7 bps increase in execution costs.

We establish causal relationship by examining executions from a set of investors who strategically opt out of the routing agreement with the HFLPs. We exploit this variation in the dataset to test the causal relationship between HFLP exposure and execution costs. Further, we also exploit a variation around an investor who removes his initial consent to route his orders to HFLPs in the later periods of the data. Using both of these identification strategies, we find statistically significant differences in execution costs between HFLP-included and -excluded executions. Executions that are potentially exposed to HFLPs through the routing relationship cost approximately 10 bps higher when compared to HFLP-excluded executions.

Investigating the underlying mechanism for the cost decrease in detail, we find that roughly 90% of the HFLP fills are filled at worse prices compared to the mid-quote price at the time of the fill. For other trades in the dataset, this ratio is only 57%. However, these spread costs can only explain a very small portion of the total execution costs. Decomposing the total *IS* into spread and price impact costs, we find that roughly 95% of the cost increase is due to price impact. This contrasts with the unconditional decomposition of the *IS* in which only 78% is due to the price impact.

Finally, we evaluate our findings primarily through the lenses of the predatory trading motives and back-running theory. Back-running theory does not seem to support our findings as we find that execution costs of informed large-order trades that are exposed to HFLPs have incrementally lower trading costs compared to the non-routed informed orders. Further, we do not find evidence of same-side trading in the later stages of the execution. Instead, HFLPs seem to provide more liquidity to the large-order in the later stages. The empirical findings are mostly consistent with predatory trading. We show that the transitory component of *IS* also loads on the HFLP exposure implying that the cost increase is largely temporary.

These contributions have significant market structure implications. First, our empirical evidence largely implies that these routing relationships between HFLPs and the brokers are not in the best interest of the investors. Thus, our findings simply advise investors to opt out of these routing relationships with HFLPs. More importantly, our empirical findings could be beneficial in deriving a set of disclosure requirements for institutional routing. In July 2016, The Securities and Exchange Commission (SEC) voted to propose rules that would require broker-dealers to disclose more information to their clients about how they handle the parent-order execution with regards to detailed routing practices. These rules mention a disclosure on execution quality and ask brokers to provide statistics on percentage of total shares executed that were priced at the side of the spread less favorable to the institutional order. Although these statistics would certainly be beneficial for clients, our findings imply that the more important statistics would be on the price impact of the routed orders.

References

Admati, A.R., Pfleiderer, P., 1991. Sunshine trading and financial market equilibrium. Review of Financial Studies 4, 443–481.

Aït-Sahalia, Y., Sağlam, M., 2013. High frequency market making: Optimal quoting .

Almgren, R., Thum, C., Hauptmann, E., Li, H., 2005. Direct estimation of equity market impact. Risk 18, 58–62.

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2011. Performance of institutional trading desks: An analysis of persistence in trading costs. The Review of Financial Studies 25, 557–598.

Anand, A., Irvine, P., Puckett, A., Venkataraman, K., 2013. Institutional trading and stock resiliency: Evidence from the 2007–2009 financial crisis. Journal of financial Economics 108, 773–797.

Battalio, R., Corwin, S.A., Jennings, R., 2016. Can brokers have it all? on the relation between make-take fees and limit order execution quality. The Journal of Finance 71, 2193–2238.

Battalio, R., Jennings, R., Selway, J., 2001. The relationship among market-making revenue, payment for order flow, and trading costs for market orders. Journal of Financial Services Research 19, 39–56.

Bessembinder, H., Carrion, A., Tuttle, L., Venkataraman, K., 2016. Liquidity, resiliency and market quality around predictable trades: Theory and evidence. Journal of Financial Economics 121, 142–166.

Brogaard, J., Hendershott, T., Hunt, S., Ysusi, C., 2014. High-frequency trading and the execution costs of institutional investors. Financial Review 49, 345–369.

Brunnermeier, M.K., Pedersen, L.H., 2005. Predatory trading. The Journal of Finance 60, 1825–1863.

Easley, D., Kiefer, N.M., O'Hara, M., 1996. Cream-skimming or profit-sharing? the curious role of purchased order flow. The Journal of Finance 51, 811–833.

Farley, R., Kelley, E.K., Puckett, A., 2018. Dark trading volume and market quality: A natural experiment. Available at SSRN .

Hatheway, F., Kwan, A., Zheng, H., 2017. An empirical analysis of market segmentation on us equity markets. Journal of Financial and Quantitative Analysis 52, 2399–2427.

Hirschey, N., 2013. Do high-frequency traders anticipate buying and selling pressure. Available at SSRN 2238516 .

Hu, G., Jo, K., Wang, Y.A., Xie, J., 2017. Institutional trading and abel noser data. Available at SSRN 3090150 .

van Kervel, V., Menkveld, A.J., 2017. High-frequency trading around large institutional orders. Journal of Finance .

Korajczyk, R.A., Murphy, D., 2015. High frequency market making to large institutional trades. Available at SSRN 2567016 .

Madrigal, V., 1996. Non-fundamental speculation. The Journal of Finance 51, 553–578.

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality. Journal of Financial Economics 100, 459 – 474. doi:`http://dx.doi.org/10.1016/j.jfineco.2011.02.006`.

Perold, A.F., 1988. The implementation shortfall: Paper versus reality. The Journal of Portfolio Management 14, 4–9.

Yang, L., Zhu, H., 2015. Back-running: Seeking and hiding fundamental information in order flows .

Table 1: Definitions of the variables extensively used in the empirical analysis.

Variable	Description
BODPdol	The fraction of the parent-order executed in the BODP in terms of dollar volume.
BODPsh	Share-based ratio of the parent-order executed in the broker's dark pool.
BODPtr	The fraction of the parent-order executed in the BODP in terms of trade count.
Duration	Order duration is expressed as a fraction of full trading day. The duration of a full trading day in U.S. equity markets is 6.5 hours.
ES	Effective spread of the parent-order expressed in basis points
ESChild	Effective spread of the child-order expressed in basis points
FracAO	Fraction of aggressive fills (share-based) in a parent-order. A buy (sell) fill is aggressive if its price is higher (lower) than the mid-point of the NBBO at the same second.
FracMO	Fraction of mid-point fills (share-based) in a parent-order. Mid-point fills are transactions happening at the mid-point of the NBBO.
FracPO	Fraction of passive fills (share-based) in a parent-order. A fill is passive if it is not aggressive or mid-point.
HasBODP	Binary variable taking a value of 1 if the parent-order execution has a child-order fill in the BODP.
HasHFLP	Binary variable taking a value of 1 if the parent-order execution has a child-order fill executed by an HFLP.
HasLIT	Binary variable taking a value of 1 if the parent-order execution has a child-order fill executed in a lit venue.
HasOthrDP	Binary variable taking a value of 1 if the parent-order execution has a child-order fill executed in a dark pool other than the BODP.
HasXPlus	This variable takes a value of 1 if the execution has fills of X or more distinct HFLPs.
HFLPdol	The fraction of the parent-order executed by the HFLPs in terms of dollar volume.
HFLPdolGF	The fraction of the parent-order executed by GFLO in terms of dollar volume.
HFLPdolOthr	The fraction of the parent-order executed by HFLPs other than GFLO in terms of dollar volume.
HFLPsh	Share-based ratio of the parent-order executed by the HFLPs.
HFLPshD1	Fraction of shares executed by the HFLPs in volume decile 1.
HFLPshD10	Fraction of shares executed by the HFLPs in volume decile 10.
HFLPshDRest	Fraction of shares executed by the HFLPs in volume deciles 2 through 9.
HFLPshGF	Share-based ratio of the parent-order executed by GFLO.
HFLPshOthr	Share-based ratio of the parent-order executed by HFLPs other than GFLO.
HFLPtr	The fraction of the parent-order executed by the HFLPs in terms of trade count.
HFLPtrGF	The fraction of the parent-order executed by GFLO in terms of trade count.
HFLPtrOthr	The fraction of the parent-order executed by HFLPs other than GFLO in terms of trade count.
IS	Implementation shortfall of the parent-order expressed in basis points
IvlTvr	The ratio of share volume during the execution interval to shares outstanding scaled by 1000.
LITdol	Dollar-based fraction of the parent-order executed in lit venues.
LITshr	Share-based ratio of the parent-order executed in lit venues.
LITtr	Trade-count-based fraction of the parent-order executed in lit venues.
OthrDPdol	The dollar-based fraction of the parent-order executed in dark pools other than the BODP.
OthrDPsh	Share-based ratio of the parent-order executed in dark pools excluding the BODP.
OthrDPtr	The trade-count-based fraction of the parent-order executed in dark pools other than the BODP.
PcpRate	Participation rate is equal to the ratio of executed volume to total volume realized during the lifetime of the parent-order.
PI	Price impact of the parent-order expressed in basis points
PPI	Permanent price impact of the parent-order expressed in basis points. We use the next day's close price and compare it to the arrival price of the parent-order. See Equation 5.
RS	Relative spread of the parent-order expressed in basis points
Spread	The bid-ask spread is normalized using the mid-quote price and time-weighted during the execution horizon. It is computed in basis points.
Volatility	Volatility is computed using the mid-quote returns based on five seconds.
VS	VWAP slippage of the parent-order expressed in basis points. See Equation 8.

Table 2: Summary statistics for the main attributes in our execution data. Value equals the dollar value of the parent-order at the time of the order submission. Participation rate (PcpRate) is equal to the ratio of executed volume to total volume realized during the lifetime of the parent-order. Number of child-orders (NumChild) equals to the number of child-order trades in a given parent-order. The bid-ask spread is normalized using the mid-quote price and time-weighted during the execution horizon. Order duration is expressed as a fraction of full trading day. The duration of a full trading day in U.S. equity markets is 6.5 hours. Volatility is computed using the mid-quote price observed every second.

Statistic	Mean	Min	Pctl(25)	Median	Pctl(75)	Max
Value ($ M)	1.015	0.050	0.131	0.343	1.001	62.864
PcpRate	0.018	0.00001	0.002	0.006	0.019	0.521
NumChild	127.8	5	26	60	148	4,533
Duration	0.52	0.03	0.16	0.52	0.90	1.00
Spread	3.95	0.71	2.35	3.22	4.62	45.13
Volatility	0.015	0.001	0.008	0.012	0.018	0.274
LITdol (%)	83.36	0.00	79.16	90.71	96.97	100.00
BODPdol (%)	10.82	0.00	0.00	0.00	11.18	100.00
HFLPdol (%)	5.41	0.00	0.00	2.65	7.69	100.00
OthrDPdol (%)	0.42	0.00	0.00	0.00	0.00	65.80
LITsh (%)	83.39	0.00	79.17	90.75	96.98	100.00
BODPsh (%)	10.85	0.00	0.00	0.00	11.25	100.00
HFLPsh (%)	5.36	0.00	0.00	2.63	7.65	100.00
OthrDPsh (%)	0.41	0.00	0.00	0.00	0.00	66.00
LITtr (%)	84.51	0.00	80.00	91.38	97.22	100.00
BODPtr (%)	9.57	0.00	0.00	0.00	10.00	100.00
HFLPtr (%)	5.60	0.00	0.00	2.63	7.69	100.00
OthrDPtr (%)	0.33	0.00	0.00	0.00	0.00	48.00
HasLIT (%)	99.27	0.00	100.00	100.00	100.00	100.00
HasBODP (%)	48.27	0.00	0.00	0.00	100.00	100.00
HasHFLP (%)	61.10	0.00	0.00	100.00	100.00	100.00
HasOthrDP (%)	8.48	0.00	0.00	0.00	0.00	100.00

Table 3: ExecBroker (Executing Broker) codes in the FIX protocol for the HFLPs in the data set.

Venue Code	Firm
CDRG	Citadel
SHAW	D.E. Shaw
GFLO	Getco
NITE	Knight
TRIM	Knight
FSOM	Sun Trading
SOHO	Two Sigma

Table 4: Statistics of various parent-order characteristics of the six clients in the *NotExposed* group and the aggregate-level statistics of the *NotExposed* and *Exposed* groups. Standard errors of the differences are adjusted by double-clustering on stock and day. All variables are defined in Table 1.

Statistic	C1	C2	C3	C4	C5	C6	*NotExposed*	*Exposed*	Diff
# of Parent-Orders	280	228	62	40	179	165	954	19,379	18,425
HasHFLP (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	64.2	-64.2***
									(1.49)
HFLPDol (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.62	-5.62***
									(0.26)
PcpRate (%)	1.82	0.84	0.95	1.02	2.03	7.50	2.52	1.76	0.76**
									(0.34)
Spread	7.22	3.95	3.76	3.84	3.78	3.55	4.79	3.92	0.88
									(0.59)
BODPdol (%)	0.00	0.00	0.00	15.66	14.27	7.06	4.55	11.16	-6.61***
									(0.69)
Volatility (%)	1.81	1.38	1.84	1.89	2.53	1.40	1.78	1.49	0.28
									(1.62)
Turnover	7.1	8.3	6.7	7.7	2.7	3.9	5.6	5.9	-0.3
									(0.07)
Duration	0.31	0.70	0.04	0.48	0.18	0.25	0.36	0.53	-0.18***
									(0.04)

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 5: Summary statistics for the execution cost metrics.

Statistic	Mean	Std. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
ES	0.68	2.02	-57.00	0.16	0.56	1.10	61.00
IS	3.12	77.71	-678.30	-27.89	2.65	35.24	698.10
PI	2.44	77.72	-680.10	-28.61	2.06	34.64	698.30

Table 6: Average percent effective spread (in bps) for active venues in the execution data set. We include venues for which we have at least 200 child-order trades.

Rank	Mean *ESChild*	Std. Error	Venue Type
1	2.52	0.12	HFLP
2	2.50	0.11	HFLP
3	2.11	0.06	HFLP
4	1.94	0.01	HFLP
5	1.76	0.05	HFLP
6	1.66	0.03	Other DP
7	1.56	0.02	HFLP
8	1.48	0.01	HFLP
9	1.48	0.05	Lit
10	1.22	0.02	Lit
11	1.16	0.02	Lit
12	1.13	0.04	Lit
13	1.05	0.02	Broker DP
14	0.99	0.09	Lit
15	0.73	0.01	Lit
16	0.69	0.03	Lit
17	0.68	0.01	Lit
18	0.64	0.00	Lit
19	0.60	0.00	Lit
20	0.46	0.04	Lit
21	0.38	0.01	Lit
22	0.29	0.09	Lit
23	-0.28	0.05	Other DP

Table 7: The lower triangular part of this table presents the pairwise Pearson correlations between the main independent variables in the multivariate regressions the upper triangular part reports the corresponding p-values. All variables are defined in Table 1.

	HFLPSh	BODPSh	OthDPSh	PcpRate	LogPrice	Volatility	Duration	Turnover
HFLPSh	1.00	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
BODPSh	-0.17	1.00	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
OthDPSh	-0.12	-0.06	1.00	[0.00]	[0.00]	[0.00]	[0.10]	[0.00]
PcpRate	-0.05	-0.11	0.11	1.00	[0.02]	[0.00]	[0.00]	[0.00]
LogPrice	-0.25	-0.05	0.06	-0.02	1.00	[0.00]	[0.87]	[0.00]
Volatility	-0.07	-0.03	-0.06	-0.07	-0.13	1.00	[0.00]	[0.00]
Duration	-0.11	0.09	-0.01	-0.27	0.00	-0.09	1.00	[0.00]
Turnover	-0.07	0.05	-0.03	-0.17	-0.06	0.26	0.46	1.00

Table 8: Regression of execution cost proxies on HFLP trading activity proxies. We add dark pool executions in the broker's and other venues (adjusted for the proxies of share, dollar or trade counts), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:								
	ES			*PI*			*IS*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
HFLPdol	1.69***			46.07***			47.75***		
	(0.30)			(8.42)			(8.39)		
BODPdol	-0.19			-9.76***			-9.98***		
	(0.19)			(3.70)			(3.69)		
OthrDPdol	1.11**			-39.31**			-38.62*		
	(0.53)			(20.04)			(19.90)		
HFLPsh		1.69***			46.01***			47.69***	
		(0.30)			(8.42)			(8.39)	
BODPsh		-0.19			-9.77***			-10.00***	
		(0.19)			(3.71)			(3.69)	
OthrDPsh		1.11**			-39.24**			-38.54*	
		(0.53)			(20.00)			(19.86)	
HFLPtr			1.39***			40.65***			42.03***
			(0.27)			(8.50)			(8.50)
BODPtr			-0.24			-9.92**			-10.21**
			(0.19)			(4.06)			(4.04)
OthrDPtr			1.48**			-43.89			-42.83
			(0.68)			(28.86)			(28.73)
PcpRate	1.46**	1.46**	1.43**	53.00***	52.99***	51.68***	54.53***	54.52***	53.17***
	(0.59)	(0.59)	(0.58)	(19.58)	(19.58)	(19.66)	(19.62)	(19.62)	(19.70)
LogPrice	-1.01***	-1.01***	-1.03***	5.72	5.71	5.37	4.67	4.67	4.31
	(0.30)	(0.30)	(0.32)	(4.74)	(4.75)	(4.77)	(4.75)	(4.75)	(4.80)
Volatility	9.46	9.46	9.31	52.44	52.45	50.68	58.08	58.09	56.16
	(12.15)	(12.15)	(12.16)	(318.36)	(318.36)	(318.65)	(314.42)	(314.41)	(314.71)
Duration	0.10	0.10	0.10	-4.25	-4.25	-4.22	-4.13	-4.13	-4.09
	(0.08)	(0.08)	(0.08)	(8.67)	(8.67)	(8.67)	(8.69)	(8.69)	(8.68)
Turnover	-0.01	-0.01	-0.01	0.38	0.38	0.38	0.38	0.38	0.38
	(0.004)	(0.004)	(0.004)	(0.38)	(0.38)	(0.38)	(0.38)	(0.38)	(0.38)
Observations	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335
Adjusted R^2	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 9: Regression of execution cost proxies on HFLP trading activity proxies when informed trading is controlled. We add dark pool executions in the broker's and other venues (adjusted for the proxies of share, dollar or trade counts), permanent price impact (PPI), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:								
	ES			*PI*			*IS*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
HFLPdol	1.71***			36.35***			38.05***		
	(0.30)			(7.40)			(7.37)		
BODPdol	-0.19			-8.58***			-8.81***		
	(0.19)			(3.21)			(3.20)		
OthrDPdol	1.06**			-18.20			-17.54		
	(0.54)			(17.33)			(17.16)		
HFLPsh		1.71***			36.34***			38.04***	
		(0.30)			(7.40)			(7.36)	
BODPsh		-0.19			-8.60***			-8.82***	
		(0.19)			(3.21)			(3.20)	
OthrDPsh		1.06**			-18.10			-17.43	
		(0.54)			(17.31)			(17.14)	
HFLPtr			1.41***			32.49***			33.88***
			(0.27)			(6.89)			(6.89)
BODPtr			-0.24			-8.90**			-9.19***
			(0.19)			(3.51)			(3.49)
OthrDPtr			1.44**			-26.75			-25.71
			(0.68)			(23.92)			(23.80)
PPI	-0.0003***	-0.0003***	-0.0003***	0.13***	0.13***	0.13***	0.13***	0.13***	0.13***
	(0.0001)	(0.0001)	(0.0001)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
PcpRate	1.40**	1.40**	1.38**	77.57***	77.56***	76.88***	79.06***	79.05***	78.33***
	(0.59)	(0.59)	(0.58)	(17.87)	(17.87)	(17.93)	(17.85)	(17.85)	(17.91)
LogPrice	-1.00***	-1.00***	-1.02***	3.76	3.76	3.51	2.72	2.72	2.44
	(0.30)	(0.30)	(0.32)	(4.74)	(4.74)	(4.74)	(4.73)	(4.73)	(4.75)
Volatility	9.39	9.39	9.25	82.63	82.65	81.17	88.23	88.25	86.61
	(12.11)	(12.11)	(12.13)	(279.67)	(279.67)	(279.70)	(276.08)	(276.08)	(276.11)
Duration	0.09	0.09	0.09	-2.91	-2.91	-2.87	-2.79	-2.79	-2.75
	(0.08)	(0.08)	(0.08)	(7.98)	(7.98)	(7.98)	(8.00)	(8.00)	(8.00)
Turnover	-0.005	-0.005	-0.01	0.35	0.35	0.35	0.35	0.35	0.35
	(0.004)	(0.004)	(0.004)	(0.28)	(0.28)	(0.28)	(0.28)	(0.28)	(0.28)
Observations	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335
Adjusted R^2	0.10	0.10	0.10	0.26	0.26	0.26	0.26	0.26	0.26

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 10: Regression of VWAP slippage (*VS*) on HFLP trading activity proxies. We add dark pool executions in the broker's and other venues (adjusted for the proxies of share, dollar or trade counts), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	VS		
	(1)	(2)	(3)
HFLPdol	5.61^{***}		
	(1.47)		
BODPdol	-0.08		
	(0.49)		
OthrDPdol	8.95		
	(7.58)		
HFLPsh		5.61^{***}	
		(1.47)	
BODPsh		-0.08	
		(0.49)	
OthrDPsh		9.01	
		(7.56)	
HFLPtr			5.41^{***}
			(1.24)
BODPtr			-0.27
			(0.55)
OthrDPtr			10.96
			(7.62)
PcpRate	2.78	2.78	2.84^{*}
	(1.72)	(1.72)	(1.69)
LogPrice	0.44	0.44	0.41
	(1.03)	(1.03)	(1.06)
Volatility	143.07^{*}	143.08^{*}	142.98^{*}
	(82.80)	(82.80)	(82.86)
Duration	-0.28	-0.28	-0.25
	(0.78)	(0.78)	(0.78)
Turnover	0.01	0.01	0.01
	(0.06)	(0.06)	(0.06)
Observations	20,335	20,335	20,335
Adjusted R^2	0.08	0.08	0.08

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 11: Regression of execution cost proxies on HFLP trading activity based on GFLO and other HFLPs. We add dark pool executions in the broker's and other venues (adjusted for the proxies of share, dollar or trade counts), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:								
	ES			*PI*			*IS*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
HFLPdolGF	1.73^{***}			35.63^{***}			37.25^{***}		
	(0.39)			(9.93)			(9.92)		
HFLPdolOthr	1.62^{***}			65.69^{***}			67.48^{***}		
	(0.32)			(11.97)			(11.96)		
BODPdol	-0.19			-9.71^{***}			-9.93^{***}		
	(0.19)			(3.71)			(3.69)		
OthrDPdol	1.11^{**}			-39.33^{**}			-38.64^{*}		
	(0.53)			(20.04)			(19.90)		
HFLPshGF		1.73^{***}			35.60^{***}			37.22^{***}	
		(0.39)			(9.93)			(9.92)	
HFLPshOthr		1.62^{***}			65.59^{***}			67.39^{***}	
		(0.32)			(11.97)			(11.96)	
BODPsh		-0.19			-9.72^{***}			-9.95^{***}	
		(0.19)			(3.71)			(3.69)	
OthrDPsh		1.11^{**}			-39.26^{**}			-38.56^{*}	
		(0.53)			(20.01)			(19.87)	
HFLPtrGF			1.39^{***}			31.21^{***}			32.52^{***}
			(0.27)			(8.90)			(8.92)
HFLPtrOthr			1.43^{**}			84.45^{***}			86.15^{***}
			(0.59)			(15.35)			(15.31)
BODPtr			-0.24			-9.55^{**}			-9.84^{**}
			(0.19)			(4.06)			(4.05)
OthrDPtr			1.48^{**}			-42.35			-41.28
			(0.68)			(28.84)			(28.71)
PcpRate	1.46^{**}	1.46^{**}	1.43^{**}	52.40^{***}	52.39^{***}	54.11^{***}	53.92^{***}	53.91^{***}	55.61^{***}
	(0.59)	(0.59)	(0.58)	(19.56)	(19.56)	(19.70)	(19.60)	(19.60)	(19.74)
Observations	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335	20,335
Adjusted R^2	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 12: Summary of various parent-order characteristics in the *HFLP* and *NoHFLP* groups. Standard errors are adjusted by double-clustering on stock and calendar day.

Statistic	*NoHFLP*	*HFLP*	Difference	p-value
# of Parent-Orders	285	285	0	
# of Child-Orders	38,768	38,435	333	
HasHFLP (%)	0.00	72.3	-72.3	<0.01
HFLPdol (%)	0.00	6.75	-6.75	<0.01
PcpRate (%)	1.99	1.28	0.72	<0.01
Spread	3.98	3.89	0.10	0.70
BODPdol (%)	5.21	9.47	-4.27	<0.01
Volatility (%)	1.56	1.56	0.00	0.96
Turnover	5.2	7.1	-1.9	0.10
Duration	0.35	0.53	-0.18	<0.01

Table 13: Difference-in-difference regression for a client who is inferred to switch to zero HFLP exposure after April 2011. *ES*, *PI*, and *IS* are the dependent variables. *Post* takes a value of 1 for executions after April 2011. *Switched* takes a value of 1 for executions belonging to the switching client. The main independent variable is *Post* × *Switched*, a binary variable that takes a value of 1 if the execution is from the switching client and occurs after April 2011. Executions are restricted to the set in the *Exposed* group. We add dark pool executions in the broker's and other venues in terms of dollar share, participation rate, log price, volatility, duration, and turnover. All of these variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	ES	*PI*	*IS*
Post × Switched	0.26	-32.32**	-31.99**
	(0.46)	(15.80)	(15.78)
Switched	-0.26	-1.58	-1.83
	(0.20)	(10.98)	(10.98)
Post	0.55	48.64***	49.31***
	(0.52)	(11.98)	(11.85)
BODPdol	-0.35*	-13.24***	-13.63***
	(0.18)	(3.39)	(3.35)
OthrDPdol	0.61	-50.87**	-50.70**
	(0.56)	(21.17)	(21.10)
PcpRate	1.37**	44.99***	46.44***
	(0.58)	(15.84)	(15.84)
LogPrice	-1.17***	2.60	1.39
	(0.34)	(4.27)	(4.32)
Volatility	2.90	67.90	67.75
	(9.08)	(255.63)	(251.97)
Duration	0.05	-7.65*	-7.58*
	(0.08)	(3.96)	(3.95)
Turnover	-0.01	0.50*	0.49*
	(0.004)	(0.29)	(0.29)
Observations	19,386	19,386	19,386
Adjusted R^2	0.09	0.10	0.10

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 14: Regressing the proxies of the aggressiveness of the parent-order with respect to HFLP trading activity. We define a buy (sell) child-order trade as an aggressive order if transaction price occurs at a price higher (lower) than the mid-point of the National Best Bid and Offer (NBBO) realized during the same second. Similarly, passive orders correspond to child order executions if the transaction price is lower (higher) than the NBBO mid-point when the investor has a large buy (sell) order. Finally, mid-point fills are transactions happening at the mid-point of the NBBO. We then create the corresponding parent-order level statistics, passive order ratio (*FracPO*), aggressive order (*FracAO*) ratio and mid-point order (*FracMO*) ratio by computing the fraction of the shares executed via each order type. We add dark pool executions in the broker's and other venues (share-based), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	FracAO	*FracMO*	*FracPO*
HFLPsh	0.37***	-0.08***	-0.29***
	(0.02)	(0.01)	(0.02)
BODPsh	-0.11***	0.03***	0.08***
	(0.01)	(0.01)	(0.01)
OthrDPsh	0.23***	0.11**	-0.34***
	(0.08)	(0.05)	(0.06)
PcpRate	0.06	-0.11***	0.04
	(0.04)	(0.02)	(0.05)
LogPrice	-0.03***	-0.01	0.04***
	(0.01)	(0.01)	(0.01)
Volatility	-0.15	-0.13	0.28
	(0.20)	(0.09)	(0.18)
Duration	-0.01*	0.01***	0.001
	(0.01)	(0.003)	(0.01)
Turnover	-0.001*	-0.0003***	0.001**
	(0.0004)	(0.0001)	(0.0004)
Observations	20,335	20,335	20,335
Adjusted R^2	0.33	0.20	0.28

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 15: Regression of execution cost proxies on HFLP trading activity during early and late stages of the execution. *HFLPshD1* (*HFLPshD10*) measures the share-based fraction of HFLP trades in the first (tenth) volume decile. *HFLPshDRest* measures the share-based fraction of HFLP trades in volume decile 2 through decile 9. We add share of dark pool executions in the broker's and other venues (share-based), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	ES	*PI*	*IS*
HFLPshD1	1.79***	62.84***	64.67***
	(0.38)	(10.57)	(10.53)
HFLPshD10	1.64***	34.37***	36.06***
	(0.34)	(9.62)	(9.57)
HFLPshDRest	1.78***	42.71***	44.55***
	(0.38)	(11.07)	(11.02)
BODPsh	-0.35*	-9.47**	-9.85**
	(0.21)	(3.96)	(3.94)
OthrDPsh	0.84	-40.32*	-39.47*
	(0.58)	(21.80)	(21.67)
PcpRate	1.12*	61.84***	62.96***
	(0.58)	(19.71)	(19.75)
LogPrice	-1.07***	7.23	6.13
	(0.29)	(5.03)	(5.00)
Volatility	4.17	162.08	162.78
	(12.23)	(328.74)	(327.04)
Duration	0.10	-4.48	-4.35
	(0.08)	(7.62)	(7.65)
Turnover	-0.01*	0.52	0.51
	(0.004)	(0.32)	(0.32)
Observations	19,247	19,247	19,247
Adjusted R^2	0.11	0.11	0.11

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 16: Regression of HFLP trading activity during late stages of the execution on informed trading proxy. *HFLPshH2* measures the share-based fraction of HFLP trades in the second half of the execution. *HFLPshD10* measures the share-based fraction of HFLP trades in volume decile 10. We add share of dark pool executions in the broker's and other venues (share-based), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:	
	HFLPshH2	*HFLPshD10*
PPI	0.038	-0.015
	(0.025)	(0.061)
BODPsh	-0.08***	-0.11***
	(0.01)	(0.01)
OthrDPsh	-0.33***	-0.51***
	(0.03)	(0.06)
PcpRate	-0.11***	-0.29***
	(0.03)	(0.05)
LogPrice	-0.06**	-0.05***
	(0.02)	(0.02)
Volatility	-0.47***	-0.32
	(0.14)	(0.22)
Duration	-0.03***	-0.06***
	(0.004)	(0.01)
Turnover	-0.0001	0.0001
	(0.0002)	(0.001)
Observations	20,335	20,335
Adjusted R^2	0.30	0.26

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 17: Regression of realized spread (RS) on HFLP trading activity proxies. We add dark pool executions in the broker's and other venues (adjusted for the proxies of share, dollar or trade counts), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	RS		
	(1)	(2)	(3)
HFLPsh	0.89^{***}		
	(0.29)		
BODPsh	-0.55^{***}		
	(0.19)		
OthrDPsh	0.59		
	(0.76)		
HFLPdol		0.89^{***}	
		(0.29)	
BODPdol		-0.55^{***}	
		(0.19)	
OthrDPdol		0.59	
		(0.76)	
HFLPtr			0.87^{***}
			(0.23)
BODPtr			-0.62^{***}
			(0.18)
OthrDPtr			0.67
			(1.02)
PcpRate	3.05^{***}	3.05^{***}	3.05^{***}
	(0.57)	(0.57)	(0.55)
LogPrice	-0.90^{**}	-0.90^{**}	-0.90^{**}
	(0.38)	(0.38)	(0.39)
Volatility	52.04^{**}	52.04^{**}	51.97^{**}
	(23.21)	(23.21)	(23.17)
Duration	0.09	0.09	0.09
	(0.11)	(0.11)	(0.11)
Turnover	−0.004	−0.004	−0.004
	(0.01)	(0.01)	(0.01)
Observations	20,335	20,335	20,335
Adjusted R^2	0.18	0.18	0.18

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 18: The impact of competing liquidity providers on execution costs. The main independent variables are proxies of HFLP fill activity and binary variables, *Has2Plus*, *Has3Plus*, *Has4Plus*, and *Has5Plus* denoting the number of distinct HFLPs executing an order routed by the broker. We add share of dark pool executions in the broker's and other venues (share-based), participation rate, log price, volatility, duration, and turnover. All of the variables are defined in Table 1. All regressions include stock, client and calendar day fixed-effects. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:		
	ES	*PI*	*IS*
Has2Plus	0.05	4.99***	5.05***
	(0.03)	(1.47)	(1.47)
Has3Plus	0.03	5.86*	5.85*
	(0.05)	(3.02)	(3.02)
Has4Plus	0.001	−1.03	−1.04
	(0.09)	(8.13)	(8.13)
Has5Plus	0.25	21.09	21.42
	(0.23)	(34.86)	(35.06)
HFLPsh	1.58***	33.61***	35.18***
	(0.31)	(8.31)	(8.27)
BODPsh	−0.18	−8.12**	−8.33**
	(0.19)	(3.69)	(3.68)
OthrDPsh	1.15**	−34.79*	−34.05*
	(0.53)	(19.69)	(19.54)
PcpRate	1.43**	49.18**	50.69**
	(0.59)	(19.69)	(19.73)
LogPrice	−1.00***	6.29	5.25
	(0.30)	(4.75)	(4.78)
Volatility	9.36	39.75	45.31
	(12.17)	(316.79)	(312.78)
Duration	0.08	−6.14	−6.03
	(0.08)	(8.57)	(8.58)
Turnover	−0.01	0.38	0.38
	(0.004)	(0.38)	(0.37)
Observations	20,335	20,335	20,335
Adjusted R^2	0.10	0.10	0.10

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$