The Cost of Exposing Large Institutional Orders to Electronic Liquidity Providers*

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

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

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

Initial Version: November 2018
This Version: June 2020

Abstract

We use a novel data set to examine the impact of exposing institutional orders to electronic liquidity providers (ELPs). We present empirical evidence that marketable pieces of large parent orders are routed to ELPs, seemingly to avoid paying liquidity fees on exchanges. This routing decision results in lower net effective spreads for these child orders, but leads to higher execution shortfall for the parent-order. We obtain causal evidence by utilizing the parent orders of investors who disallow the broker to route their child orders to ELPs. Our analysis suggests this cost increase is due to information leakage about the parent-order.

Keywords: Electronic Liquidity Providers, Transactions Costs, Order Anticipation

JEL Classification: G12, G14.

*We thank Jeff Bacidore, Yashar Barardehi (discussant), Andy Puckett (discussant), seminar participants at the Ohio University and University of Notre Dame, and participants at the AQR Capital Management, IEX Market Structure Conference, 6th Annual Conference on Market Regulation and the 2019 WFA meetings. An earlier version of this paper was circulated under the title, "The Cost of Routing Orders to High Frequency Traders."

1. Introduction

Jefferies Financial Group estimates that approximately 40% of institutional trading is in the form of orders directed to broker-dealer algorithms that use different strategies to achieve their trading goals.[1] Algorithms designed to execute orders at the volume-weighted average price (VWAP) over some period of time typically slice a large 'parent' order into child orders that are released into the market over time. Hendershott et al. (2011) write that "algorithms often use a mix of active and passive strategies, employing both market and limit orders" and they rely on smart order routers (SORs) to determine where each child order should be routed.[2] SORs can route child orders to three different types of trading centers: registered exchanges, alternative trading systems (ATSs) mainly operating as dark pools and electronic liquidity providers (ELPs) functioning as proprietary trading firms.[3] Sofianos (2007) writes that order routing logic should optimize execution across all of these trading centers to minimize information leakage.

Liquidity fees and rebates introduce another level of complexity that must be considered when evaluating the performance of smart routers. As the fees/rebates vary across exchanges, dark pools and ELPs, the Equity Market Structure Advisory Committee of the Securities and Exchange Commission (SEC) notes that broker-dealers who charge fixed commissions (e.g., do not pass back fees and rebates directly to their customers) have "an incentive to route to the venue with the highest rebate, rather than diligently search out the venue likely to deliver the best execution of its customer's order." Sofianos et al. (2011) note that brokers who minimize trading costs for client orders are penalized by a 0.04 basis point reduction in net revenue while clients of brokers that minimize all-in venue fees are penalized by a 0.26 basis point increase in trading costs.[4]

[1] See `https://www.sec.gov/comments/10-222/10222-321.pdf`

[2] Johnson (2010) writes that "an algorithm is a set of instructions for accomplishing a given task." He notes that the rules determining the type, the price and quantity for each of these child orders are often based on a mixture of historical and live market data. The algorithm is responsible for monitoring each child order, adjusting or cancelling as and when it becomes necessary. See Bertsimas and Lo (1998) for theoretical derivations of optimal trading strategies that minimize the expected cost of trading a large equity position over time.

[3] Registered exchanges are referred to as "lit markets" whereas all other trading centers constitute "dark markets." Throughout the paper, we will use the term "dark pool" and "ATS" interchangeably. For example, Credit Suisse operates an SOR that utilizes all of the thirteen exchanges, nineteen ATSs and eight ELPs. The full list can be accessed at `https://www.credit-suisse.com/media/assets/sites/aes/doc/aes-us-order-handling-guidelines.pdf`.

[4] In a comment letter dated May 24, 2018, BabelFish Analytics writes that "clients that specifically instruct brokers to remove rebate-driven trading behaviors from their algorithms achieve significantly lower trading costs that result in higher returns to their investors." In a May 16, 2018 website post, Clearpool notes that in a recent analysis of VWAP algorithms, they found that the fee-sensitive algorithm missed the VWAP and underperformed relative to the fee agnostic algorithm.

A broker operating a SOR can avoid liquidity fees by entering an agreement to send marketable child orders to one or more ELPs.[5] However, sending one or more child orders directly to an ELP can alert the ELP that a parent-order is being worked.[6] Because institutional orders have price impact (Keim and Madhavan, 1998), ELPs can profit from the early knowledge that an institution is seeking to buy or sell a large block of stock regardless of the institution's motivation for trading. In a memorandum by the Division of Trading and Markets of the SEC,[7] it is noted that by offering cheap liquidity, ELPs will have a "high position on taker routing tables" that will allow the ELPs "to interact with the first tranche of a large market order, thus allowing the traders to detect the earliest signs of a potential price move and quickly adjust their quoting or trading strategies on other markets." Once aware that a large parent-order is being worked, the trading of ELPs may exacerbate the price impact of the parent-order and lead to inflated trading costs. As described in SEC administrative proceedings against large U.S. broker-dealers in recent years, many institutions seek to avoid interacting with ELPs when executing large parent orders.[8] In this paper, we empirically examine the impact that routing child orders to ELPs has on the execution quality of the parent-order.

We use a novel data set of institutional parent orders which can be routed to the ELPs after obtaining the permission of the investor in the pre-trade phase. The data set includes more than 20,000 parent orders and 2.5 million child order trades occurring between January 1, 2011 and March 31, 2012. The average parent-order is sizable at around \$1 million and corresponds to roughly 2% of the volume realized during the parent-order execution. The data set identifies where each child order is executed and allows us to exactly identify the set of ELPs who provide liquidity

[5] See SEC Administrative Proceeding File No. 3-18549, which notes that "when orders were executed by ELPs, Merrill Lynch avoided the access fees typically charged by exchanges while receiving commissions from customers."

[6] van Kervel and Menkveld (2019) present evidence that HFTs operating on the Nasdaq OMX stock exchange in Sweeden are unable to detect the presence of a large parent-order early in its life. The fact that the HFTs do not have any information on the counterparty to their trades makes it more difficult to detect institutional order flow. Similarly, Baxter (2017) writes that "as dark pools are not required to immediately disclose execution or trade information, predatory firms are – without pinging orders – less able to front-run dark market trades." The author also notes that institutions can eliminate or reduce the effect of pinging orders by utilizing a minimum order size in dark pools.

[7] `https://www.sec.gov/spotlight/emsac/memo-maker-taker-fees-on-equities-exchanges.pdf`

[8] See Administrative Proceeding File No. 3-18766, which notes that "many market participants, particularly institutional firms, sought to avoid trading against HFT" in Citi Match, a dark pool run by Citigroup Global Markets. Also see Administrative Proceeding File No. 3-18549, which states that certain customers of Merrill Lynch "were concerned that orders routed to ELPs could be subject to information leakage" and that "certain customers specifically requested that their orders not be sent to ELPs." The proceeding also noted that "when orders were executed by ELPs, Merrill Lynch avoided the access fees typically charged by exchanges while receiving commissions from customers."

to each parent-order. These ELPs include the largest high-frequency traders during our sample period: Citadel, D.E. Shaw, Getco, Knight, Sun Trading and Two Sigma. The average ratio of trade volume executed by ELPs is 5.6% and more than 60% of the parent orders have at least one child order filled by this group of ELPs.

Our analysis of the trading costs incurred by individual marketable child orders suggests that, net of exchange fees, ELPs provide the cheapest liquidity when measured by spread costs. While this finding seemingly rationalizes the routing of marketable child orders to ELPs, it ignores how the underlying stock price moves during the life of a parent-order. Adverse price movements after the exposure of a parent-order to ELPs may increase the total trading costs by offsetting the savings in spread costs. We next examine when child orders are routed to ELPs and find that a high-share of ELP trades occur in the early stages of the execution, which provides the ELPs with knowledge that can be used to employ order anticipation trading strategies.[9] Consistent with assertion that the ELPs in our sample utilized their knowledge that a parent-order was being worked at the expense of that parent-order, we find the underlying price of the stock being purchased (sold) by a sample parent-order is more likely to rise (fall) while it is being worked when the order is exposed to ELPs. More specifically, despite the fact that ELPs offer inexpensive liquidity to individual child orders, parent orders that source more ELP liquidity have higher average trading costs, all else being equal. This finding is robust to controlling for venue fees and rebates, proxies for trade complexity, various control variables that can co-vary with liquidity and stock, day and investor fixed-effects. In terms of economic magnitude, a parent order execution with median trading cost would approximately move to 55th percentile of the cost distribution if its ELP exposure is increased from zero to 10% while reducing its aggressive exchange fills by the same amount.

Since the routing decision is determined in equilibrium, quantifying a causal effect is empirically challenging. For example, the broker may be routing the most toxic orders to the ELPs and if these variables are not properly controlled in the regression, this may bias the estimated impact of the ELP exposure. Our data set addresses this concern as it includes a set of investors whose parent orders are never routed to an ELP. We exploit this variation in the data set to test the

[9]SEC (2010) explains that "one example of an order anticipation strategy is when a proprietary firm seeks to ascertain the existence of one or more large buyers (sellers) in the market and to buy (sell) ahead of the large orders with the goal of capturing a price movement in the direction of the large trading interest (a price rise for buyers and a price decline for sellers)." See SEC Release No. 34-61358.

causal relationship and find that parent orders exposed to ELPs have higher implementation costs compared to matched parent orders whose child orders are not exposed to ELPs. Further, there is an active investor in the data set who has child orders executed by ELPs only in the first part of our sample period, suggesting the investor strategically disallowed the broker to route its child orders to ELPs after a particular date. We design a difference-in-differences regression to exploit this variation and document a statistically significant decrease in this investor's transaction costs after eliminating the ELP exposure.

We next investigate the specific channels of the increase in trading costs when child orders are exposed to ELPs. First, we examine whether there is additional cost when child orders are exposed to ELPs in the early stages and find that all else equal, early exposure to ELPs is costlier than a comparable late exposure to ELPs. Second, we directly study the trading costs around child order executions. By comparing the execution prices to recent market prices, we find that ELP transactions occur at significantly costlier prices compared to other venues. Third, we find that successive fills by the ELPs are costlier when compared to successive fills in other venues, implying that the ELPs may take advantage of the urgency of the institutional order. Finally, we find that exposing the child orders to multiple ELPs increases the implementation costs when compared to a same level of exposure with a single ELP. Both of these empirical findings highlight the information leakage channel associated with ELP exposure.

We provide empirical evidence that suggests routing marketable child orders in a manner that reduces venue fees is associated with elevated trading costs for parent orders. For brokers that do not pass through venue fees directly to their clients, this result presents a dilemma as routing orders to minimize parent-order trading costs reduces their net revenue. While Battalio et al. (2016) provide results that suggest that the use of differential fee schedules by different equity trading venues creates agency problems for retail brokers, our results indicate this agency problem extends to institutional brokers as well. Thus, our results provide empirical justification for the SEC's recent decision to extend broker routing information to institutional orders. Our results also provide economic justification for the SEC's decision to allow broker-dealers to exclude certain types of liquidity providers from their dark pools.

The remainder of this paper is organized as follows. Section 2 provides a brief literature review. In Section 3, we describe the data set and provide its summary statistics. Section 4 specifically examines the patterns of the liquidity provided by the ELPs during the lifetime of the parent-order. Section 5 provides a brief rationale on why ELPs are utilized. In Section 6, we present multivariate regression results on the relationship between overall trading costs and ELP exposure. Section 7 utilizes the client aversion to ELPs to establish a causal link between ELP exposure and execution costs. Section 8 investigates the underlying mechanisms that can explain the increase in trading costs and presents empirical evidence consistent with the information leakage channel. Finally, we conclude in Section 9.

2. Related literature

This paper is related to three strands of the microstructure literature. First, there are a number of studies that examine how the potential conflicts of interest arising from payment for order flow and make-and-take fees impact trading costs and liquidity provision. Battalio et al. (2016) find that several retail brokers monetize their order flow by selling market orders to wholesalers and routing their limit orders to exchanges that offer large liquidity rebates. They present evidence that routing orders in this manner adversely impacts retail limit order execution quality. Bacidore et al. (2011) compare the performance of the broker smart routers and find that their performances do not vary significantly for marketable orders. However, they find that for nonmarketable orders, routers that aim to maximize rebates perform worse than routers that maximize fill rates.[10]

In a contemporaneous paper, Anand et al. (2019) use FINRA's proprietary Order Audit Trail System (OATS) database to empirically examine the relation between broker routing, venue ownership and execution outcomes. They examine over 350 million institutional orders in 300 stocks that were placed by 43 active institutional brokers in October 2016. The average-sized order in their analysis is 1,348 shares. Their data do not indicate, however, the identity of the trader or the strategy (e.g., VWAP, implementation shortfall, percentage of volume) utilized to execute the order. Among other things, the authors present evidence that the larger institutional brokers route a disproportionate number of child orders to their own dark pool and that this type of order routing leads to inflated trading costs. In contrast, our study examines 20,335 parent orders seeking to trade an average of 26,000 shares that were placed by 146 unique investors and were executed via

[10] The role of the brokers in order flow signal leakages is also studied in Barbon et al. (2019) in which best clients of the broker front-run the liquidating funds working with the same broker.

a single broker's VWAP algorithm. While we are unable to compare and contrast the effectiveness of the routing strategies used by different institutional brokers, we are able to hold broker and algorithm constant to examine whether exposing parent orders to ELPs inflates trading costs.

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 (2019), who use child executions obtained from four institutional investors and trade level data from the Nasdaq OMX in Sweden that identifies the counterparties to every execution. They present evidence that HFTs first provide liquidity to and later trade alongside large, informed institutional orders and show that this type of behavior increases institutional trading costs. In contrast to van Kervel and Menkveld (2019), our data set 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 the aggressiveness in which their algorithms work orders. van Kervel and Menkveld (2019) construct the net flows of HFT trading around institutional orders by using the trade reports provided by Nasdaq OMX, a public lit exchange which has 65% of the market share. The authors exclude 11.5% of their institutional orders because they generated both buy and sell orders and they are unable to observe HFT activity in dark pools. 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.

There are significant differences in contribution between this paper and our study. First, van Kervel and Menkveld (2019) 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 directly routed to an ELP for execution. In van Kervel and Menkveld (2019), 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 (2019) 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 ELPs. Finally, our data set provides more granular data at the institutional side and has information about ELP child order trades in the non-lit venues. Conversely, van Kervel and Menkveld (2019) has high-quality data for HFT trading activity occurring on a stock exchange.

Korajczyk and Murphy (2019) show that HFT liquidity provision is significantly lower for stressful large institutional orders 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 (2019), this data set does not exactly identify how a parent-order is split into child orders. HFTs initially provide liquidity to the large order but then compete 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 ELPs. Since some investors opt out of the routing relationship, we obtain causal evidence between exposure to ELPs 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 market quality implications of aggregate dark venue executions (e.g., Buti et al., 2011; Comerton-Forde and Putniņš, 2015; Foley and Putniņš, 2016). All of the ELP trades in our data set show up as dark pool trades in the publicly available TAQ data set. Exchange code 'D' is used in the TAQ data set to identify all trading within ATSs as well as internalized trades at the broker-dealers. 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; 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 data set allows us to study the liquidity implications of a unique group of dark-labeled trades that was not possible to study in the past literature.

3. Data and summary statistics

3.1. Description of the institutional order 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 parent orders and the corresponding child order executions are provided by the execution desk of a large investment bank. We next describe this data set in detail and provide institutional details about the execution strategy.

The investment bank, hereafter referred to as the broker dealer, 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 (see Sağlam et al., 2019). For this purpose, the parent-order executions in the sample have been selected from an active subset of the broker-dealer's clients. An investor is considered active if he has at least 100 and at most 500 VWAP parent-order executions in S&P 500 stocks between January 2011 and March 2012 that are fully executed and take at least 10 minutes to execute. The rationale behind these filters are explained in detail in Sağlam et al. (2019) but we briefly summarize them here for completeness. First, VWAP orders are passively managed so the performance of the trade would be mainly based on investor's timing ability rather than the broker's skill. Second, S&P 500 stocks are very liquid and relatively homogeneous subset of the stock universe and it would be unlikely to have an investor trading with insider information. Third, investors with small number of executions (< 100) are eliminated as their short-term trading skill may not be estimated reliably with fewer observations. Investors with high number of executions (> 500) are also filtered to have a balanced data set across clients and prevent any specific investor from fully driving the estimation procedure. Finally, short executions are often very small in size that may get executed with a single market order.

With these criteria, we obtain 22,074 parent orders and 2.6 million child order trades. At the parent-order level, most of the variables 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), arrival price (NBBO mid-quote at the start time of the parent-order execution), average execution price, proportional quoted spread and mid-quote volatility. Further, we have detailed information on the child order executions generated by each parent-order. For each child order execution, the data set includes the trade time (timestamped to the millisecond), trade size, executing venue, and the trade price. We do not know whether the executed child order is a market or a limit order and we do not have unexecuted child orders.

We exclude parent-order executions which have less than 5 child order executions or have value less than $50,000 at the arrival time of the order which correspond to approximately 1,500 parent-order executions. We also exclude an additional 200 executions with missing entries of participation rate, spread, volatility, or duration. The final sample consists of 20,335 parent orders, 9,856 buy and 10,479 sell orders placed by 146 distinct investors.

The orders originate from a diverse pool of investors, such as institutional portfolio managers, quantitative investment funds, internal trading desks and brokers who aggregate their retail order flow. The data set 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 between the individual clients and our data provider. However, the broker-dealer informed us that there are two common practices: fixed commission and pass-through. In the fixed commission scheme, the client pays a fixed fee 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 to us that some of their clients chose this broker to maximize their rebate revenue. Each of the sample parent orders is executed by a 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. The algorithm slices the parent-order into child orders based on the historical volume curve realized over the past month during the planned execution horizon. Institutional clients can shape the parameters of the VWAP algorithm with a limited set of pre-trade instructions that may affect the ultimate execution cost. Clients select the size and side of the parent-order as well as the start time and targeted order completion time. Clients can also instruct the algorithm to avoid routing child orders to an ELP or to dark pools. The broker-dealer refers to ELPs as "Liquidity Partners" in its user interface. While we do not have access to any of these pre-trade instructions, we can examine where a client's child order executions occur to infer whether the client consistently avoids ELPs or specific dark pools. Once the order is initiated, the clients do not have any control over the creation, duration, and routing of each child order.

Our data provider furnished us with information that allows us to use the Financial Information Exchange (FIX) protocol tags in the data set to identify the venue on which each child execution occurs. These data reveal the presence of six ELPs in our data set: Citadel, D.E. Shaw, Getco, Knight/Trimark, Sun Trading, and Two Sigma. As we do not observe any child executions from D.E. Shaw after January 7, 2011, we effectively have five active ELPs during our sample period.

Finally, we do not know the exact payments that our data provider received from the ELPs for our broker's child orders. However, our broker informed us that it did not incur any explicit cost when an ELP executed a child order. We understand that child orders routed by our data provider to the ELPs were to either be executed at or within the prevailing National Best Bid and Offer (NBBO) or rejected. Although we do not have any data on rejection rates, our data provider explicitly told us that its routing agreements with the ELPs specified that rejection rates cannot be high.

3.2. Summary Statistics

Table 1 contains summary statistics for the sample of 20,335 parent orders placed into our broker's VWAP algorithm between January 2011 and March 2012. The average parent-order in our sample has a market value of just over $1 million. Parent order sizes range from $50,000 to roughly $63 million. The average (median) parent-order generates 1,278 (60) child order executions. We find that on average, roughly 69% of the parent-order is filled aggressively and 31% is filled passively.[11]

In general, trading costs are increasing in a parent-order's participation rate, measured as the ratio of the parent-order's trading volume to the trading volume of the underlying stock during the parent-order's life. The average (median) sample parent-order consumed 1.8% (0.6%) of the volume traded during the parent-order's life. The average and the median parent-order in our sample had a life of three hours and thirty-eight minutes (e.g., 0.52×6.5 hours).

On average, exchanges execute 84.5% of the shares in an average parent-order, while broker operated dark pools and ELPs respectively execute 9.6% and 5.6% of the shares. The median parent-order uses ELPs to execute 2.63% of its shares and does not trade any shares in broker operated dark pools. Overall, 99% of the sample parent orders execute at least one child order on an exchange, while 61% of the parent orders have at least one child trade with an ELP. Less than half of the parent orders receive a child order execution in the broker operated dark pool.

3.3. Comparison of the data with other institutional trading data sets

In terms of order size and participation rate, our parent-order execution data are very similar to other data sets employed in the broader microstructure literature. For example, Anand et al.

[11] We define a buy (sell) child execution as a passive fill if the transaction price occurs at a price that is lower (higher) than the midpoint of the execution-time NBBO.

(2011) find the average daily participation rate of the institutions in the Ancerno institutional order database is 2.1% of the total market volume between 1999 and 2008 (1.0% in 2008). Compared to the Ancerno data, our data set has information on the exact start- and end-time of the parent-order execution (i.e., order duration), the interval volume during the execution period of the parent-order (that helps us compute participation rate), the algorithm type (i.e., VWAP), and interval return.[12] More importantly, we have information on the executed child orders, e.g., the price, quantity and execution venue of the child order, which enables us to study the impact of interacting with ELPs in the dark.

Korajczyk and Murphy (2019) and van Kervel and Menkveld (2019) also employ institutional trading data (from Canada and Sweden, respectively) to examine the impact of high-frequency trading on parent-order execution costs. Korajczyk and Murphy (2019) report an average trade size of $2.2 million with a participation rate of 2.5%. van Kervel and Menkveld (2019) examine roughly 5,000 parent orders (inferred from individual child order trades) 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.

4. Understanding ELP executions

In this section, we would like to uncover the drivers of ELP fills in detail. Due to the sensitivity of the topic, the data providers were not able to share the exact logic behind how the routing to ELPs works. In this section, we run several analyses to understand the variation in ELP fills. Intentionally, our focus will be on the visual evidence which will motivate our formal analysis in the later sections.

4.1. The life cycle of ELP executions

All else equal, the potential for information leakage is greatest the earlier one or more ELPs become aware that a parent-order is being worked. To examine when the VWAP algorithm sources ELP liquidity, we compute the fraction of shares executed by ELPs in various deciles of the parent-order.

[12] Hu et al. (2018) report that Ancerno data set has client identifiers only through 2011.

The first 10% of the shares executed are in the first decile and the last 10% of the shares executed are in the last decile. Figure 1 (top left chart) illustrates that the algorithm's use of ELPs varies throughout the life of a parent-order, as the fraction of volume executed by ELPs has a U-shaped pattern. This suggests that the VWAP algorithm sources ELP liquidity early in a parent-order's life, which potentially allows the ELPs to engage in order anticipatory trading strategies throughout most of the parent-order's life.

Figure 1 (top right chart) reveals that the algorithm's use of exchanges varies across time, with the percentage of volume executed on exchanges being highest in the middle of a parent-order's life. Figure 1 (bottom center chart) suggests the algorithm's use of the broker's dark pool is relatively constant throughout the parent-order's life. Together, these results suggest there is a negative correlation between the sourcing of exchange and ELP liquidity. Indeed, we find that the correlation between the percentage of volume executed by ELPs and exchanges in the first and tenth deciles are -50% and -54%, respectively.

We expect that all of the child orders executed by ELPs were marketable (e.g., the ELPs sold shares at the offer and bought shares at the bid). Conversely, in an effort to meet or beat the VWAP benchmark and generate liquidity rebates, we expect the algorithm to route at least some nonmarketable (passive) child orders to the exchanges. To understand the extent to which the algorithm routes child orders to ELPs after passive child orders go unexecuted on exchanges, we compute the percentage of volume executed by exchanges within each decile of the parent-order's life with regards to their aggressiveness. The results presented in Figure 2 suggest that an increase in the ELPs' share of child order fills leads to a decrease in the share of passive and aggressive executions in the exchanges. During the later stages of the parent-order execution, it is expected that the passive share would decline. Given that there is no large spike in aggressive fills in the higher deciles, it appears the algorithm is not simply canceling unexecuted child orders on exchanges and replacing them with marketable child orders that are routed to the ELPs.

4.2. Client-level instructions and ELP trades

We do not have the pre-trade instructions of the investors so we cannot exactly identify the set of clients who prohibit the broker from exposing their orders to ELPs. However, we can infer this group of investors with a simple procedure. Analyzing the ELP trades at the investor level, we find that

4 out of 146 investors have no child orders executed by an ELP. These clients have individually 280, 228, 62 and 40 parent orders consisting of 40,098, 19,997, 1,153 and 2,397 individual child order executions, respectively. These 610 parent orders are executed on 224 distinct stocks on 217 different dates, so it is almost certain that these investors have chosen not to source ELP liquidity. Further, three of these clients have also opted out of dark pool executions, implying a strategic desire to trade in lit venues. In addition, we have two additional clients who have only one child order executed by an ELP during our sample period. These clients have individually 180 and 166 parent orders consisting of 64,251 and 13,136 child executions. For the first of these clients, the lone ELP interaction occurs on the first parent-order placed by the client during our sample period. Collectively for these two clients, the ratio between ELP-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, ELP exposure of these two clients is abnormally low. Consequently, we will also label their 344 parent orders (excluding the two parent orders with ELP trades) as *Not Exposed* to ELPs during their life. Based on this classification, our sample contains 954 parent orders in the *Not Exposed* group and 19,379 parent orders in the *Exposed* group.

We now examine the differences in the venues on which child orders are executed for these two groups of parent orders. First, we study the shares of exchange and broker dealer owned dark pool executions of the *Not Exposed* group as a function of volume deciles. Figure 3 confirms our earlier conjecture about the substitution between exchange and ELP executions when we compare these plots to those in Figure 1. When a parent-order is not exposed to ELPs, we observe the percentage of shares executed on exchanges in the first 9 volume deciles are roughly the same, suggesting that child orders that would have been routed to ELPs are routed instead to the exchanges when a client prohibits ELP liquidity sourcing. Next, we look at the parent-order statistics for the *Not Exposed* and *Exposed* samples. Table 2 provides the order characteristics of each client in *Not Exposed* group and compares the average characteristics across two samples. We specifically examine the participation rate, the duration of the parent-order, the stock price volatility while the parent-order is being worked, the percentage of the parent orders' dollar volume executed in the broker owned dark pool, and the percentage of the parent orders' dollar volume executed by an ELP. On average, the participation rate (duration) of the *Exposed* parent orders is 76 basis points lower (70 minutes lower) than the orders that were not exposed to ELPs. There is no difference in volatility across the two sets of parent orders. Given that trading costs are generally increasing in participation rate, all else equal, these statistics suggest that the *Exposed* parent orders will have lower all in trading costs. In Section 7.1, we will utilize this variation of ELP exposure at the investor level to examine its impact on implementation costs.

4.3. Relative tick size and ELP trades

Considering that ELPs are also very active in the lit venues, routing relationships may be very useful for ELPs if the underlying stock has a relatively lower price. For these stocks, the queues at the best prices in the lit venues may be very large and if the ELP does not have time-priority, it will not have execution priority. Consequently, a routing relationship with the broker may allow the ELP to gain execution priority by offering price improvement to the routed order and the ELP can make almost all of the half-spread. Further, these relationships can be helpful for the ELP to manage its inventory more effectively as predicted in Aït-Sahalia and Sağlam (2017).

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 ELP trades and nominal stock price visually. We have 109 parent orders on 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 4 provides striking evidence that after the reverse split, the ratio of ELP trades drop dramatically from 37% to 4%. Given the higher spread and higher proportional take fees of low-priced stocks, ELPs can be an attractive venue to execute.

5. Rationalizing the decision to route individual child orders to ELPs

We first examine the impact of sourcing ELP liquidity on the cost of individual child order executions. The exchanges with the largest market share of trading use the make/take pricing model to arrange trades, which imposes a liquidity fee on liquidity demanding orders. Conversely, based on evidence provided in the SEC settlement documents, the ELPs do not charge liquidity fees and often offer rebates to attract marketable orders. Broker operated dark pools tend to charge low liquidity fees and/or offer rebates to attract marketable orders, but the probability of a child order execution in a broker operated dark pool is lower than that offered by ELPs. Thus, brokers seeking to minimize liquidity fees will route marketable child orders to dark pools and ELPs before routing them to exchanges.

Exchanges, dark pools, and ELPs each offer price improvement to marketable orders (i.e., execution prices within the prevailing National Best Bid and Offer (NBBO)). Brokers typically negotiate with ELPs as to the amount of price improvement that their marketable child orders will receive and the amount the ELPs will pay for the order flow. On exchanges and in dark pools, price improvement typically occurs when a marketable child order interacts with a hidden limit order whose limit price is within the NBBO. Thus, in contrast to ELP liquidity provision, price improvement is probabilistic on exchanges and in dark pools.

We use the relative effective spread, measured as the signed difference between the execution price and the midpoint of the execution-time NBBO normalized by the midpoint of the execution-time NBBO, to measure execution quality for marketable child orders. We also compute the relative effective spread net of the typical exchange liquidity fee, since some of the data provider's clients pay/receive the fees/rebates generated by their child order executions. To compute the fee-adjusted spreads, we use the fee schedules reported by the lit exchanges using the baseline values. These schedules report the default fees or rebates associated with liquidity addition or removal. We assume that each passive (aggressive) child order is subject to the fee or rebate corresponding to adding (removing) liquidity. Finally, we assume that child orders executed by ELPs and child orders executed in dark pools do not pay a liquidity fee. Table A.1 reports our fee assumptions for each venue that appears in our data set.

Table 3 illustrates the average raw and net effective spreads of the marketable child orders executed in the various venues. We also report aggregate statistics by grouping venues under "inverted exchanges," "ELPs," "make-take exchanges," "other dark pools" and the "broker's own dark pool." As a group, ELPs do not seem to offer substantially better prices for marketable orders if we rank the venues by the average nominal effective spread. However, when we account for rebates and fees, make-take exchanges become substantially costlier than ELPs. Make-take exchanges are roughly 40% more expensive than ELPs when compared in terms of net effective spread. While the inverted exchanges (e.g., exchanges that offer rebates for marketable orders) and other dark pools have lower average net effective spreads than the average ELP, for many stocks and/or in many market conditions these venues are not offering to trade at the NBBO. Overall, this univariate analysis implies that ELPs can be an attractive venue choice by allowing the broker to save from take fees for marketable orders. We next examine how the cost of the parent-order execution differs as a function of the exposure to ELPs in a multivariate setting.

6. Exposure of parent orders to ELPs and overall trading costs

6.1. Measure of parent order execution quality

Perold (1988) introduced the Implementation Shortfall (*IS*) 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 (see e.g., Anand et al., 2011, 2013). It is computed as the normalized difference between the volume-weighted average price of all child orders and the price of the asset prior to the start of the execution as proxied by the NBBO midpoint. More formally, for the ith parent order in our sample, we compute *IS* as follows:

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

where $\operatorname{sgn}(PO_i)$ is 1 if the ith parent order is a buy order and is -1 if it is a sell order, $Q_{i,j}$ is the size of the jth child order of the ith parent order, $P_{i,j}$ is the price at which the jth child order is executed, N_i is the number of child order executions for the ith parent order, and $M_{i,0}$ is the midpoint of the NBBO that was prevailing when the parent order starts being executed. Similarly, we will use *ISAdj* to denote the *IS* that is adjusted for estimates of venue fees and rebates. Note that we can decompose the *IS* into two terms where the first term is the share-weighted effective spread of the parent order's child executions normalized by the midpoint of the NBBO prevailing when the first child order trades and the second term measures the trade-weighted drift in the underlying stock price over the duration of the parent order's life:[13]

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

All else equal, the second term of this equation (which is not affected by exchange fees/rebates) can be larger when order anticipators become aware of the presence of a parent order. We use *ISAdj* to denote the implementation shortfall that has been adjusted for the assumed liquidity fee or rebate incurred by child orders. We again use the values reported in Table A.1 in the Online Appendix for our fee assumptions for each venue that appears in our data set.

The mean (median) ES is 0.68 (0.56) bps in the data set whereas the mean (median) *IS* is 3.12 (2.65) bps in the data set. Thus, *ES* (*Drift*), on average, constitutes approximately 22% (78%) of the total *IS*. We find that the average *ESadj* (*ISadj*) increases to 0.95 (3.50) bps. Percentage-wise, this is a substantial change in the effective spread but since the drift term remains unaffected from the fees or the rebates, the percentage change in *IS* is rather limited.

6.2. ELP exposure and trading costs

Execution costs can be a function of multiple trade-level and stock-level characteristics, thus, to test formally whether ELP exposure is associated with higher costs, we run the following multivariate regression at the parent order execution level with a rich set of control variables:

$$
\begin{aligned}
\textit{Trading Cost}_i &= \alpha + \beta \textit{ELP Exposure}_i + \theta_1 \textit{BODP Exposure}_i + \theta_2 \textit{Other DP Exposure}_i \qquad (3)\\
&+ \theta_3 \textit{PasExch Exposure}_i + \sum_j \delta_j \textit{Control}_{j,i} + \sum_{s=1}^{S}\gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K}\nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T}\eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i,
\end{aligned}
$$

[13] See Holden et al. (2014) for a similar decomposition.

where *Trading Cost*$_i$ is either *IS* or *ES* along with their fee adjusted versions. *ELP Exposure*$_i$, *BODP Exposure*$_i$, *Other DP Exposure*$_i$ and *PasExch Exposure*$_i$ are respectively the percentage of the ith parent order's child orders executed by an ELP, broker's own dark pool, other dark pools and passive executions in exchanges. Note that we have to exclude aggressive exchange exposure, *AggExch Exposure*$_i$, from the regression, as the sum of all venue exposure variables is equal to 1. With this regression model, $\beta \times 0.01$ will capture the additional cost incurred when 1% of a parent order is executed by ELPs instead of being filled aggressively in exchanges.

The mapping $i \overset{m}{\to} s$ is used to identify the executed stock s, the mapping $i \overset{c}{\to} k$ is used to identify the client k submitting the order, and the mapping $i \overset{d}{\to} t$ 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 parent order's participation rate, the logarithm of the arrival price, the volatility of the NBBO midpoint throughout the duration of the parent order, the duration of the parent order execution, and turnover. 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. Finally, price level and volatility may also affect the execution rates of the ELPs 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 institutional trading costs.[14] Throughout the paper, we compute standard errors by clustering at the stock and calendar day level.

Table 4 reports the regression results. Consistent with our univariate analysis, *ELP Exposure* is not a significant variable in explaining raw effective spreads. Further, the estimated coefficient on *ELP Exposure* is negative and statistically significant when the spread costs are adjusted for fees and rebates. This finding again rationalizes the routing decision to ELPs in the multivariate setting. However, when we account for the price drift in addition to spread costs in the form of *IS* and *ISAdj*, we observe that the estimated β coefficient is positive and statistically significant at the 0.01 level. Given the ELPs offer aggressive child orders lower net effective spreads, these results suggest the trade-weighted drift in the underlying stock price over the duration of the parent order's life is increasing in *ELP Exposure*. This result holds independent of whether or not the fees

[14]For example, Almgren et al. (2005) includes participation rate and volatility in analyzing the variation in *IS*. van Kervel and Menkveld (2019) uses order duration, volatility, turnover, client and stock fixed effects.

and rebates generated by the child orders are passed through to the institutional client.

To understand the exact economic impact of β, consider the following example. Suppose that in the benchmark case, 100% of a parent order execution is traded in exchanges with aggressive orders and now you consider having an ELP exposure of 5.6% (the mean value of ELP exposure in the data). In this hypothetical example, compared to the benchmark case, *IS* would increase by (0.056)(44.45)=2.49 bps. Although this number is nominally high, it would be more appropriate to gauge its economic significance considering the variance of the cost distribution. Formally, an execution with median *IS* would move to 53rd percentile of the *IS* distribution if its ELP exposure is increased from zero to 5.6% while reducing its aggressive exchange exposure by the same amount.

The economic importance can be also quantified by comparing β to the coefficient on *Participation Rate*, one of the most important drivers of execution costs documented in the literature (e.g., Almgren et al., 2005) . We observe that β is of the same order of magnitude as the *Participation Rate* coefficient. This suggests that an increase in the percentage of a parent order executed by ELPs has a similar impact on institutional trading costs as an increase in trading aggressiveness.

Our regressions also analyze the relationship between other dark pool trading activity and institutional trading costs. In contrast to the strong pattern we identified with ELP fills, we find no evidence that exposing parent orders to the broker's dark pool inflates trading costs. Similarly, the coefficient on other dark pool exposure is negative but insignificant. Overall, the different signs of 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; 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.

6.3. Robustness Tests

GETCO accounts for approximately 70% of the trades (60% of the dollar volume) that are executed by ELPs. To examine whether the results presented in Table 4 are driven by exposure of parent orders to GETCO, we decompose our measure of ELP trading activity into two components: the percentage of child orders executed by GETCO and the percentage of child orders executed by

ELPs other than GETCO and re-estimate equation 3. We present the results in Table A.2 in the Online Appendix. As is the case when all ELP child executions are aggregated, we find that exposure to GETCO and exposure to the other ELPs is associated with inflated trading costs for the parent order. This suggests child order executions by GETCO are not driving our results.

IS uses the NBBO midpoint prior to the execution of the first child order as a benchmark price to compute the execution cost measure. One popular ex-post benchmark for the average execution price is the volume-weighted average price (VWAP) during the parent-order's trading interval. Madhavan (2002) argues that VWAP slippage may not be a reliable proxy of execution quality, as early aggressive trading would significantly affect the realized VWAP. Nevertheless, we examine the robustness of our findings using VWAP slippage as a measure of overall trading costs. Table A.3 in the Online Appendix reports the main regression results using VWAP slippage and its fee-adjusted version. In both cases, the estimated coefficient on ELP Exposure is highly significant implying the robustness of our results with respect to a different proxy.

7. Evidence from client aversion to ELPs

Our findings in the previous section point to a strong positive correlation between proxies of ELP trading activity and parent order execution costs which may not directly imply a causal relationship. There may be an omitted variable that is correlated with execution costs and ELP 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 parent order executions from six institutional investors who have no more than one child order executed by an ELP during our sample period. We assume that these investors chose not to directly source ELP liquidity for their parent orders. Second, we examine the trading costs incurred by a particular institutional client before and after she seemingly stops allowing the broker to route her child orders to ELPs.

7.1. Clients without ELP exposure

Recall that Table 2 provides the order characteristics of the distinct six clients who have zero (clients 1 through 4) or abnormally small exposure to ELPs (clients 5 and 6). First, we observe that each client has quite different order characteristics. For example, client 6 has a large participation rate and relatively short duration, which may lead to higher execution costs at the parent order level. Conversely, client 2 has a relatively small participation rate but longer duration, which can lead to lower trading costs. In contrast to clients 4, 5 and 6, who execute between 7% and 15% of their child orders in our data provider's dark pool, clients 1, 2, and 3 have no exposure to the broker's dark pool. One takeaway from this heterogeneity is that there does not seem to be a common pattern across these six investors aside from the fact that they avoid ELPs.[15]

The clients choosing not to expose their orders to ELPs, may be following different trading strategies with regards to which stock to buy or sell and when to trade. Ultimately, we would like to compare the execution costs of two investors who follow identical trading strategies but only differ in their exposure to ELPs. We employ an exact matching procedure to obtain such a sample of treated and control groups. For each parent order i in the *Not Exposed* 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 five 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 364 (out of 954) parent orders using this exact procedure. We will label these groups of parent orders, *NoELP* and *ELP*, respectively. This exact matching procedure addresses the potential timing and stock-selection ability of the clients in the *Not Exposed* group.

Next, we compare the characteristics of the *NoELP* and *ELP* matched samples. Table 5 provides the detailed comparison. We specifically check BODP and PasExch exposure, participation rate, logarithm of the arrival price, volatility, and quoted spread. We observe that the differences in logged price, volatility, and quoted spread are not statistically significant implying the success of the exact matching for these characteristics. We observe that parent order executions in the *NoELP* group have lower BODP exposure, and larger participation rate, which would typically

[15] Although the investor identities are masked, the description field regarding the parent order identifier may signal the true identity of the investor. For three of these six clients, we observe that texts like "BNP," "SUSQ" and "FID" appear suggesting that these orders may be coming from BNP Paribas, Susquehanna International Group and Fidelity Investments, respectively. For the remaining clients, we observe "AMC," "AMMN" and "TXP" appearing in their parent order identifiers. Overall, this group of investors does not seem to be part of the broader ELP group. We thank Andy Puckett for suggesting this check.

lead to higher execution costs.

We now test whether parent orders in the *NoELP* group have lower implementation shortfall compared to the ones in the ELP group by running the following multivariate regression:

$$\textit{Trading Cost}_i = \alpha + \beta \textit{Not Exposed to ELPs}_i + \sum_j \delta_j \textit{Control}_{j,i} + \epsilon_i, \tag{4}$$

where the cost measure is the raw and net *IS* and the control variables include BODP and PasExch exposure, participation rate, and order duration.

Table 6 reports the regression results. In both specifications, the coefficient on *Not Exposed to ELPs* is negative and highly significant. The estimates point to an additional cost of approximately 12 bps in the parent order executions with ELP trading relationship. Overall, this evidence provides a causal interpretation to the observed positive correlation between trading costs and ELP exposure. Our results are robust to different matching procedures (e.g., propensity score matching).

7.2. The client that stops exposing its orders to ELPs

Large institutional clients like Fidelity and Vanguard can experiment with pre-trade instructions to maximize their execution performance. These clients have the ability to statistically examine the relative performance of parent orders that are exposed to ELPs. For each client in our sample, we evaluate the time series of child executions to determine whether the client can be inferred to initiate or stop exposing child orders to ELPs. None of the clients in our analysis avoid ELPs in the early part of the sample period and source ELP liquidity in the latter part of the sample period. We do, however, find one institutional investor who seemingly allows its orders to source ELP liquidity in the early part of our sample period, but refrains in the latter part - none of the client's 40,000+ child orders generated by more than 100 parent orders executed after March 18, 2011 are routed to ELPs. This sharp switch suggests the client has prohibited the broker from routing the child orders to the ELPs. Furthermore, the client's child orders execute in the broker's and in other dark pools throughout the sample period. The fact that the client appears to stop sourcing ELP liquidity is consistent with what one would expect if the client detected that parent orders with ELP exposure had higher trading costs than those that were not exposed to ELPs. In this section, we test this hypothesis.

Between January 2011 and March 2011 this client has 40 parent orders (5,239 child order trades) that have sporadic exposure to ELPs. On March 18, 2011, in his last parent order execution during this period, there is a child order executed by an ELP suggesting that the client consents for ELP exposure. During this period, 67.5% of the client's parent order executions have ELP exposure and 4.2% of the trades are executed in the broker's own dark pool. The client's next parent order execution after this period occurs on April 6, 2011 and starting from this execution, the client's child orders are not routed to ELPs in 111 distinct parent orders consisting of 40,479 child order trades. The client still continues to trade in the dark pools as 4.1% of the client's orders are executed in the broker's own dark pool in this no-ELP period. These statistics strongly imply that the client switched its instructions on ELP routing after March 18, 2011.

Given the switching decision, one may wonder whether the client changes his overall trading strategy. We compare various parent order and stock-level statistics between the pre- and post-switch period and report the statistics in Table A.4 in the Online Appendix. We specifically check BODP and PasExch exposure, participation rate, logarithm of the arrival price, volatility, and quoted spread. Since volatility and quoted spread can change significantly, we check whether these statistics for the client remain stable over time by computing the z-scores of his executed stock volatility and quoted spread by normalizing with the daily means and standard deviations of the statistics. We find that apart from the ELP exposure, all of the parent order and stock-level characteristics are very similar suggesting that the client does not seem to follow a different trading strategy in the post-switch period.

Given that the same client switches from active ELP exposure to zero ELP exposure, we can design a difference-in-differences framework to formally test the impact of ELP exposure on execution costs. Let *Switched* take a value of 1 for parent order executions generated by the switching client's orders and zero otherwise. *Post* is an indicator variable that takes a value of 1 for executions after March 18, 2011. In the control group, we can use the parent order executions in the Exposed group that have ELP exposure throughout the data period. We then run the following diff-in-diff regression using the same set of control variables and stock, day and client fixed effects as in equation 3:

$$Trading\ Cost_i = \alpha + \beta Post_i \times Switched_i + \kappa Post_i + \upsilon Switched_i + \sum_j \delta_j Control_{j,i} \tag{5}$$
$$+ \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i,$$

Table 7 reports the regression results using raw and net implementation shortfall as the cost metrics. We find that β is negative and statistically significant for both measures implying that after the switch, the client realized lower trading costs. Given that the client seems to be following the same trading strategy except disallowing ELP exposure, these findings provide further causal support for the cost increase in parent order executions with exposure to ELPs.

8. How do ELPs increase transaction costs?

In this section, we conduct four additional analyses to further explore whether the evidence is consistent with the claim that ELPs use their knowledge that an institutional order is being worked to engage in order anticipatory trading strategies to the detriment of the institutional order.

8.1. Early versus late ELP exposure

In Figure 1, we observe the U-shaped pattern for ELP fills. Given the positive relationship with execution costs and ELP exposure, we expect that early exposure to ELPs would lead to higher execution costs. This finding would be consistent with the information leakage channel. To investigate this explanation formally, we decompose our ELP trading activity proxy into three components: the percentage of child executions by ELPs in the first decile (*ELP Exposure D1*), the tenth decile (*ELP Exposure D10*), and in deciles 2 through 9 (*ELP Exposure D2_D9*). By construction, the sum of these variables equals ELP Exposure. To examine this hypothesis, we re-run our main specification (see equation 3) with the three decomposed ELP trading variables replacing *ELP*

Exposure:

$$\begin{aligned}
\textit{Trading Cost}_i &= \alpha + \beta_{\text{Early}} \textit{ELP Exposure D1}_i + \beta_{\text{Late}} \textit{ELP Exposure D10}_i + \beta_{\text{Middle}} \textit{ELP Exposure D2_D9}_i \\
&+ \theta_1 \textit{BODP Exposure}_i + \theta_2 \textit{Other DP Exposure}_i + \theta_3 \textit{PasExch Exposure}_i + \sum_j \delta_j \textit{Control}_{j,i} \\
&+ \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i \qquad (6)
\end{aligned}$$

Table 8 reports the regression results and verifies our conjecture. Using a Wald test, we confirm that β_{Early} is statistically greater than β_{Late} for each specification. Consistent with our hypothesis, we also observe that $\beta_{\text{Middle}} > \beta_{\text{Late}}$ with statistical significance. Overall, these findings point to higher information leakage with earlier ELP fills.

8.2. Child order analysis of ELP executions

If ELPs anticipate institutional orders with high accuracy, we would expect that buy (sell) child orders executed by ELPs are abnormally higher when compared to recent transactions right before. This would be consistent with a typical "front-running" theory where ELPs buy (sell) the stock, push the price up (down) and then execute an institution's buy (sell) child order. To test this potential channel, we run an analysis at the child order level and study whether trading with ELPs are costlier compared to other venue transactions. Formally, we define *PreCost* at the child order level:

$$\textit{PreCost}_k = \text{sgn}\left(CO_i\right) \frac{P_k - VWAP_{(t_k - 30s, t_k)}}{VWAP_{(t_k - 30s, t_k)}}, \qquad (7)$$

where sgn (CO_k) is 1 if the kth child order is a buy order and is -1 if it is a sell order, $VWAP_{(t_k-30s,t_k)}$ is the volume-weighted average price realized over the 30 seconds prior to the kth child order trade using TAQ data set.

Using 2.6 million child orders, we now run a multivariate regression at the child order level and test formally whether each ELP child order trade is associated with higher costs, as measured by

PreCost after controlling for child order size, stock, client and date dummies:

$$PreCost_k = \alpha + \beta ELP_k + \theta_1 BODP_k + \theta_2 OtherDP_k + \theta_3 PasExch_k + \delta Log\ (TradeSize)_k \tag{8}$$
$$+ \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(k)=s\}} + \sum_{j=1}^{K} \nu_j \mathbb{I}_{\{c(k)=j\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(k)=t\}} + \epsilon_k,$$

where ELP_k is a binary variable taking a value of 1 if the kth child order is executed by an ELP. $BODP_k$, *PasExch*, and $OtherDP_k$ are defined analogously. $Log\ (TradeSize)_k$ is the logarithm of the dollar size of the child order. In the full specification, we include stock ($k \overset{m}{\to} s$), client ($k \overset{c}{\to} j$), and day ($k \overset{d}{\to} t$) with the appropriate mapping functions. With this regression model, since the corresponding binary variable for $AggExch_k$ is omitted from the model, β captures the relative cost incurred if the child order is executed by an ELP instead of being filled aggressively in an exchange.

Table 9 reports the regression results. In all specifications with and without fixed effects, the coefficient on *ELP* is statistically significant implying that child orders are sold (bought) at relatively higher (lower) prices by the ELPs when compared to comparable aggressive fills in exchanges. The coefficient is also higher than that of dark pools both economically and statistically, so ELP trades do not perform well compared to dark pools either. Note that combined with our analysis in Section 5, our findings imply that the mid-quote of the asset is abnormally rising (falling) right before ELP transactions aligning with the order anticipation hypothesis.

8.3. Cost of successive fills

One important characteristic of ELP fills is the lack of anonymity due to direct routing of the order. Thus, one would expect that successive child order trades in the ELPs 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 ELP would be much larger. Figure 5 illustrates a visual evidence for this conjecture using three different venue types: ELP, dark pool and a lit venue. We chose GETCO, BODP and BATS, respectively, for these venue types as these markets have roughly similar number of trades in terms of order of magnitude (GETCO: 108 thousand, BODP: 166 thousand, and BATS: 157 thousand). We observe that ELP 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.[16]

8.4. Competition between ELPs

Hendershott and Madhavan (2015) write that "revealing trading intentions to many potential counterparties can lead to costly information leakage." Recall that 61.1% of all parent orders have at least one child order executed by an ELP. We compute the breakdown of this statistic with regards to distinct ELPs that provide liquidity to a parent order and find that 25.3% of the parent orders interact with one ELP, 21.4% trade with two ELPs, 11.7% source liquidity from three ELPs, 2.4% transact with four ELPs, and 0.25% trade with five of the six sample ELPs. None of our parent orders interact with each of the six sample ELPs. If ELPs engage in order anticipatory trading strategies, one might expect transactions costs to be higher when a parent order is exposed to multiple ELPs.

Conversely, Bessembinder et al. (2016) 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 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 ELPs in the data, if the broker routes orders to multiple ELPs, this theory would suggest a potential decrease in execution costs compared to the case where there is a single ELP taking advantage of the order flow information.

To examine the impact that exposing parent orders to multiple ELPs has on trading costs, we

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

run the following regression:

$$\begin{aligned}
&\mathit{Trading\ Cost}_i = \alpha + \lambda_1 \mathit{More\ than\ 2\ ELPs}_i + \lambda_2 \mathit{More\ than\ 3\ ELPs}_i + \beta \mathit{ELP\ Exposure}_i \\
&\quad + \theta_1 \mathit{BODP\ Exposure}_i + \theta_2 \mathit{Other\ DP\ Exposure}_i + \theta_3 \mathit{PasExch\ Exposure}_i + \sum_j \delta_j \mathit{Control}_{j,i} \\
&\quad + \sum_{s=1}^{S} \gamma_s \mathbb{I}_{\{m(i)=s\}} + \sum_{k=1}^{K} \nu_k \mathbb{I}_{\{c(i)=k\}} + \sum_{t=1}^{T} \eta_t \mathbb{I}_{\{d(i)=t\}} + \epsilon_i, \qquad (9)
\end{aligned}$$

where the cost measure is *IS* and *ISAdj* and the control variables include the same set of variables used in equation 3. We also include the binary variables, *More than 2 (3) ELPs* which take a value of one if the parent order is exposed to more than 2 (3) distinct ELPs, to examine the incremental impact of routing to multiple ELPs.

Table 10 reports the regression results. We find that the coefficients on *More than 2 ELPs* and *More than 3 ELPs* are positive and significant. These findings suggest that after controlling for the ratio of ELP child order executions, having additional competing ELPs does not decrease execution costs. These findings are instead more consistent with higher information leakage when the orders are routed to multiple ELPs.

9. Conclusion

This paper examines the costs of routing relationships between a broker and ELPs from the perspective of an institutional investor. Typically, brokers do not pass the fees and rebates associated with executing parent orders directly through to their institutional clients. This creates an incentive for brokers to reduce the fees and/or increase the rebates generated when working institutional orders. One common way for brokers to reduce fees is to route marketable child orders to ELPs, who provide child order executions at or within the prevailing NBBO and, unlike many exchanges, do not charge liquidity fees. A potential cost of this routing strategy is that ELPs become aware that a large parent order is being worked in the market place and can use this information to engage in profitable order anticipatory trading strategies to the detriment of the parent order. We use parent orders executed by a single VWAP algorithm to examine whether the sourcing of ELP liquidity ultimately leads to inflated trading costs.

We first present evidence consistent with the idea that routing marketable child orders to ELPs reduces net effective spreads relative to the major exchanges that charge positive liquidity fees. These results, however, ignore the potential impact that ELPs can have on the underlying stock price over the life of the order. We next examine how exposing parent orders to ELPs effects the implementation shortfall for parent orders. We obtain robust evidence suggesting parent orders that source liquidity from ELPs incur higher transactions costs than those that avoid ELPs. This finding is robust to controlling for venue fees and rebates along with other control variables including stock, day and investor fixed-effects. In the most conservative case, our results suggest that if 1% of a parent order is routed to ELPs instead of the stock exchanges, the implementation shortfall for the order will increase by roughly 12%.

We establish a causal relationship by examining the parent order executions from a set of investors who strategically opt out of the routing agreement with the ELPs. We exploit this variation in the data set to test the causal relationship between ELP exposure and execution costs. Further, we also exploit the fact that one of our sample institutions seemingly ceases to allow our data provider to source ELP liquidity in the latter part of our sample period. Using both of these identification strategies, we find statistically significant differences in execution costs between ELP-exposed and ELP-excluded parent orders. Investigating the underlying mechanism for the cost increase in detail, we examine the cost of exposing the child orders to ELPs in different stages of a parent order's life and find that early exposure is substantially costlier. We also present evidence that exposing a parent order to multiple ELPs leads increases the parent order's implementation shortfall. Each of these findings is consistent with the information leakage channel.

More generally, our results suggest that there is a cost to sourcing liquidity directly from ELPs. When an ELP interacts with a single child order on an exchange, it does not know where the order originated. The results of van Kervel and Menkveld (2019) suggest that ELPs can detect the presence of an institutional order (with error) by interacting with multiple child orders on lit exchanges. When an institutional broker sources liquidity from an ELP with whom it has an ongoing relationship, the ELP immediately becomes aware that a parent order is being worked. This is because the ELP can see the broker from which the order originated. Thus, while the relationship allows the broker to obtain low cost liquidity from ELPs for individual child orders, it also allows ELPs to more quickly ascertain that a large institutional order is being worked and to use that knowledge to earn trading profits at the expense of the institutional order.

References

Aït-Sahalia, Y., Sağlam, M., 2017. High frequency market making: Optimal quoting. Unpublished working paper. University of Cincinnati.

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.

Anand, A., Samadi, M., Sokobin, J.S., Venkataraman, K., 2019. Institutional order handling and broker-affiliated trading venues. Available at SSRN 3338129 .

Bacidore, J., Otero, H., Vasa, A., 2011. Does smart routing matter? The Journal of Trading 6, 32–37.

Barbon, A., Di Maggio, M., Franzoni, F., Landier, A., 2019. Brokers and order flow leakage: Evidence from fire sales. The Journal of Finance 74, 2707–2749.

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.

Baxter, B.P., 2017. The securities black market: Dark pool trading and the need for a more expensive regulation ats-n. Vand. L. Rev. 70, 311.

Bertsimas, D., Lo, A.W., 1998. Optimal control of execution costs. Journal of Financial Markets 1, 1–50.

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.

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

Buti, S., Rindi, B., Werner, I.M., 2011. Diving into dark pools. Charles A. Dice Center Working Paper .

Comerton-Forde, C., Putniņš, T.J., 2015. Dark trading and price discovery. Journal of Financial Economics 118, 70–92.

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

Foley, S., Putniņš, T.J., 2016. Should we be afraid of the dark? dark trading and market quality. Journal of Financial Economics 122, 456–481.

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.

Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity. The Journal of Finance 66, 1–33.

Hendershott, T., Madhavan, A., 2015. Click or call? auction versus search in the over-the-counter market. The Journal of Finance 70, 419–447.

Holden, C.W., Jacobsen, S., Subrahmanyam, A., et al., 2014. The empirical analysis of liquidity. Foundations and Trends® in Finance 8, 263–365.

Hu, G., Jo, K.M., Wang, Y.A., Xie, J., 2018. Institutional trading and abel noser data. Journal of Corporate Finance 52, 143–167.

Johnson, B., 2010. Algorithmic Trading & DMA: An introduction to direct access trading strategies. volume 200. 4Myeloma Press London.

Keim, D.B., Madhavan, A., 1998. The cost of institutional equity trades. Financial Analysts Journal 54, 50–69.

van Kervel, V., Menkveld, A.J., 2019. High-frequency trading around large institutional orders. The Journal of Finance 74, 1091–1137.

Korajczyk, R.A., Murphy, D., 2019. High-frequency market making to large institutional trades. The Review of Financial Studies 32, 1034–1067.

Madhavan, A.N., 2002. Vwap strategies. Trading 2002, 32–39.

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.

Sağlam, M., Moallemi, C.C., Sotiropoulos, M.G., 2019. Short-term trading skill: An analysis of investor heterogeneity and execution quality. Journal of Financial Markets 42, 1–28.

SEC, 2010. Concept release on equity market structure. Federal Register 75, 3594–3614.

Sofianos, G., 2007. Dark pools and algorithmic trading. Algorithmic Trading Handbook, 2nd ed.(The Trade, London) , 59–72.

Sofianos, G., Xiang, J., Yousefi, A., 2011. Smart order routing: All-in shortfall and optimal order placement. Goldman Sachs, Equity Executions Strats, Street Smart 42, 14.

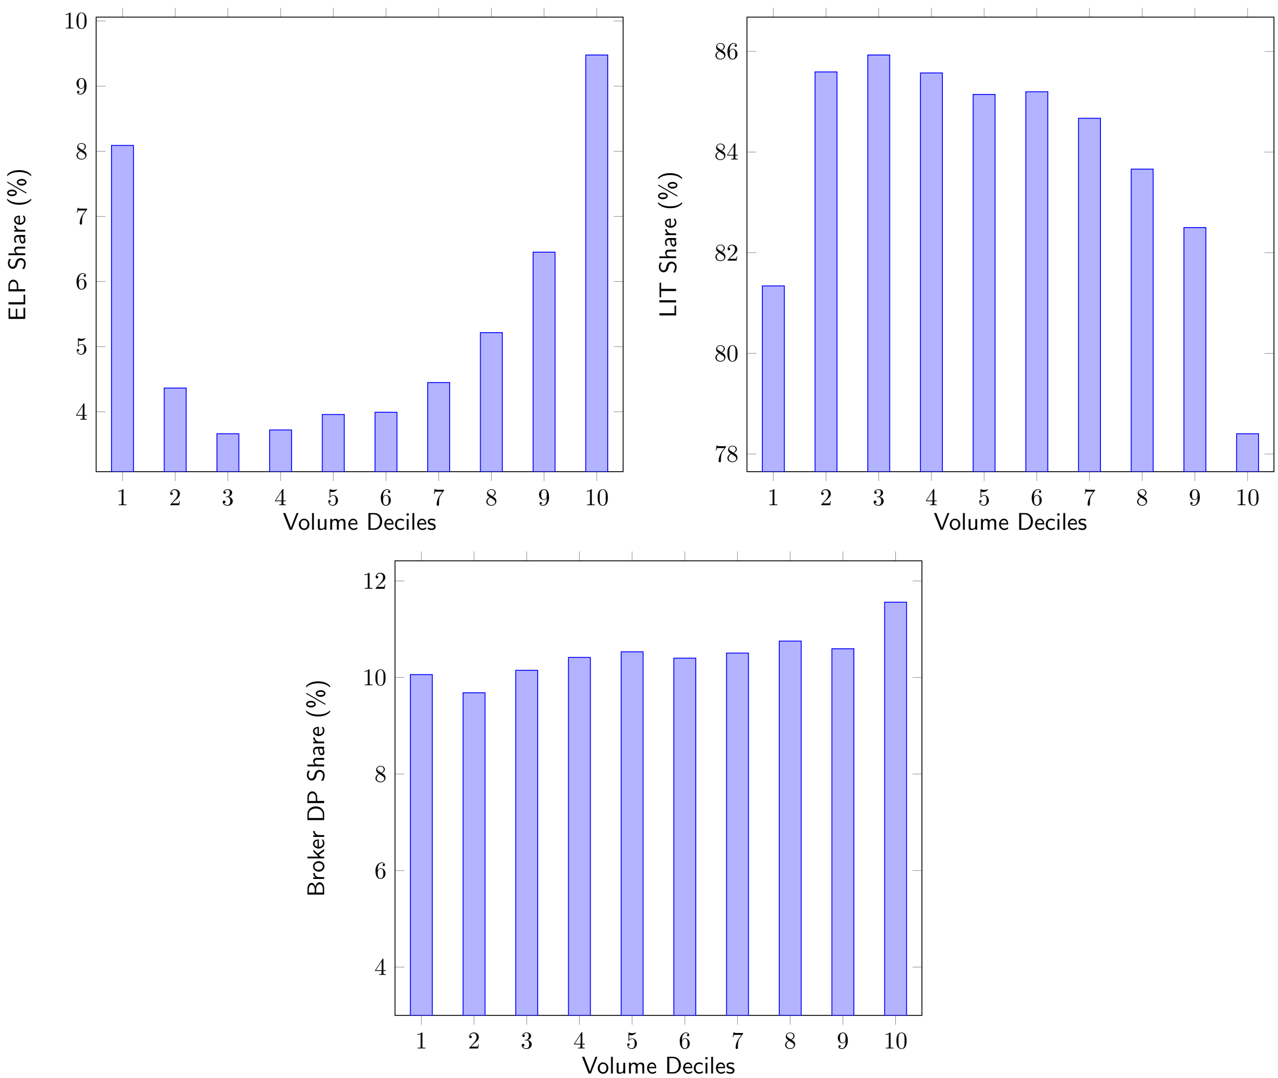

Figure 1: Share of ELPs, exchanges and the broker's dark pool in executed volume deciles.

Notes: We plot the percentages of ELP (top left), lit-venue (top right) and the broker's dark pool (bottom) executions in various deciles of executed volume, e.g., ELPs' share of each 1000 shares for a parent order consisting of 10,000 shares.

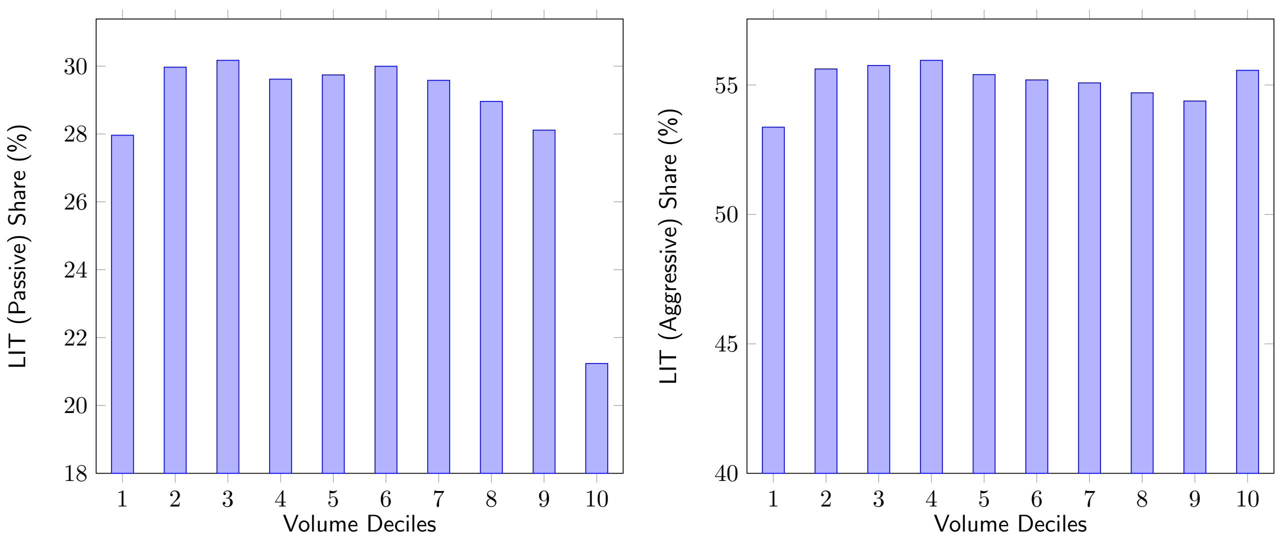

Figure 2: Share of passive and aggressive fills in executed volume deciles.
Notes: We plot the ratios of passive (left) versus aggressive (right) lit-venue executions in various deciles of executed volume.

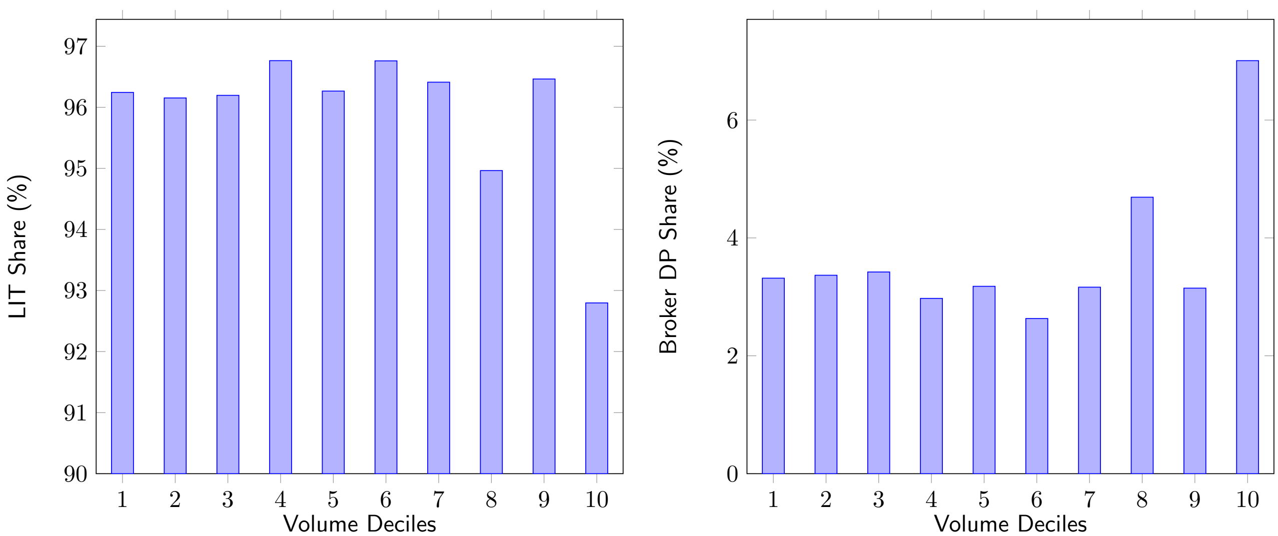

Figure 3: Share of exchanges and the broker's dark pool in executed volume deciles when ELPs are not sourced.

Notes: We plot the ratio of lit exchanges (left) and the broker's dark pool (bottom) executions in various deciles of executed volume in the *Not Exposed* group.

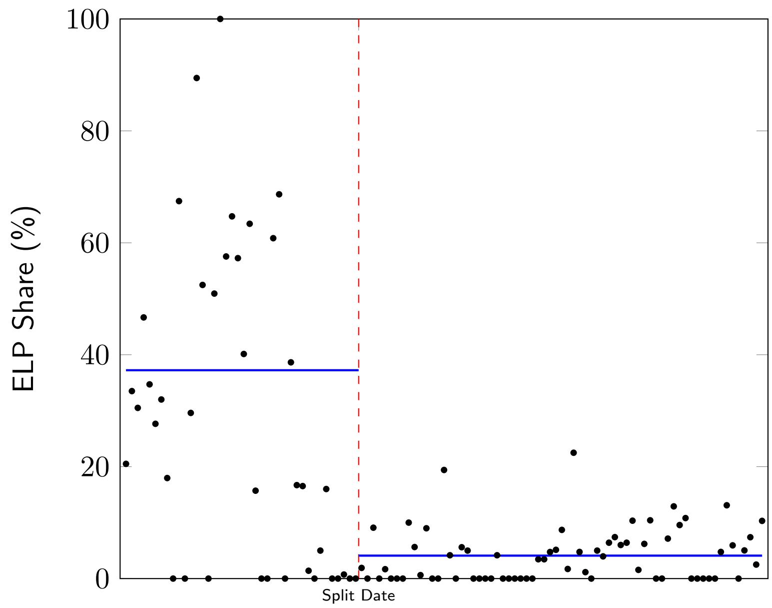

Figure 4: Shares of ELP fills on Citigroup Inc. (NYSE:C) before and after the 1-for-10 reverse stock-split on May 9, 2011. *Notes:* We plot the average share of ELP fills before and after the reverse stock split.

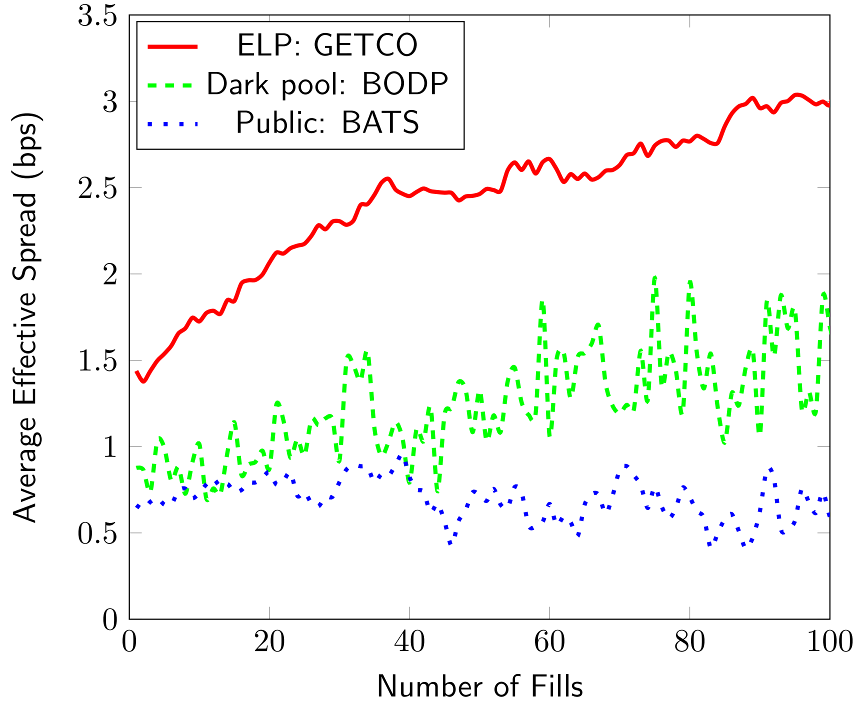

Figure 5: Effective spread of successive fills in different venues.
Notes: Fixing the execution venue, we plot the mean effective spread of the nth child order fill where n is shown in the x-axis.

Table 1: Summary statistics

Notes: This table presents descriptive statistics for the 20,335 parent order executions placed with a large institutional broker between January 2011 and March 2012. There are 9,865 buy orders and 10,479 sell orders in 498 of the S&P 500 stocks. Participation rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked.

	Average	Min.	25th Pctl.	Median	75th Pctl.	Max.
Value of parent order ($ millions)	1.015	0.050	0.131	0.343	1.001	62.864
# of child order executions	127.8	5	26	60	148	4,533
% of parent order executed w/ aggressive child executions (trades)	69.31	0.00	60.00	70.00	79.17	100.00
% of parent order executed w/ passive child executions (trades)	30.69	0.00	20.83	30.00	40.00	100.00
Participation rate	0.018	0.000	0.002	0.006	0.0019	0.521
Parent order duration (% of trading day)	0.52	0.03	0.16	0.52	0.90	1.00
% of parent order executed on an exchange (trades)	84.51	0.00	80.00	91.38	97.22	100.00
% of parent order executed in a broker operated dark pool (trades)	9.57	0.00	0.00	0.00	10.00	100.00
% of parent order executed by an ELP (trades)	5.60	0.00	0.00	2.63	7.69	100.00
% of parent orders with an execution by an exchange	99.27	0.00	100.00	100.00	100.00	100.00
% of parent orders with an execution in a broker operated dark pool	48.27	0.00	0.00	0.00	100.00	100.00
% of parent orders with an execution by an ELP	61.10	0.00	0.00	100.00	100.00	100.00

Table 2: Descriptive statistics for parent orders that avoid ELPs

Notes: Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. BODP Exposure is the percentage of a parent order's child executions that occur in the broker's own dark pool. ELP Exposure is the percentage of a parent order's number of trades with ELPs. The table below presents averages for parent orders placed by each of the six clients that seemingly avoid sourcing ELP liquidity, the entire sample of parent orders placed by clients that appear to avoid sourcing ELP liquidity, and the sample of parent orders placed by clients that allow their parent orders to interact with ELPs.

	Client 1	Client 2	Client 3	Client 4	Client 5	Client 6	Not Exposed	Exposed	Diff
# of parent orders	280	228	62	40	179	165	954	19,379	18,427
Participation rate (%)	1.82	0.84	0.95	1.02	2.03	7.50	2.52	1.76	0.76^{*}
Parent order duration	0.31	0.70	0.04	0.48	0.18	0.25	0.36	0.53	-0.18^{***}
Volatility (%)	1.81	1.38	1.84	1.89	2.53	1.40	1.78	1.49	0.28
BODP Exposure	0.00	0.00	0.00	12.46	9.22	5.90	3.27	9.87	-6.60^{***}
ELP Exposure	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.87	-5.87^{***}

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

Table 3: Relative and fee-adjusted net effective spreads for marketable child orders by venue

Notes: Table **??** reports our fee assumptions for each venue that appears in our data set.

Venue	Type	# of Child	*Effective Spread*		*Net Effective Spread*	
		Executions	Average (bps)	Std. Error	Average (bps)	Std. Error
BYX Stock Exchange: BYX	Inverted Exch.	2,776	2.15	0.04	1.94	0.04
Nasdaq BX Stock Exchange: BX	Inverted Exch.	2,668	2.55	0.05	2.33	0.05
Citadel: CDRG	ELP	15,199	1.88	0.01	1.88	0.01
Knight Securities: TRIM	ELP	5,386	1.90	0.02	1.90	0.02
Knight Securities: NITE	ELP	1,214	2.03	0.05	2.03	0.05
Getco: GFLO	ELP	97,415	2.29	0.01	2.29	0.01
D.E. Shaw: SHAW	ELP	555	2.61	0.12	2.61	0.12
Two Sigma: SOHO	ELP	1,969	2.81	0.05	2.81	0.05
Sun Trading: FSOM	ELP	724	3.03	0.10	3.03	0.10
Broker-owned dark pool: BODP	Broker's D.P.	93,745	4.33	0.02	4.33	0.02
BIDS ATS: BIDS	Other D.P.	400	0.85	0.05	0.85	0.05
Level ATS: EBXL	Other D.P.	6,798	2.29	0.03	2.29	0.03
EDGA Stock Exchange: EDGA	Low Fee Exch.	9,049	2.16	0.02	3.18	0.03
New York Stock Exchange: NYSE	Make/Take Exch.	230,185	2.11	0.01	2.91	0.01
Nasdaq Stock Exchange: Nasdaq	Make/Take Exch.	614,059	2.13	0.00	3.05	0.00
Archipelago Stock Exchange: ARCX	Make/Take Exch.	289,624	2.14	0.00	3.07	0.00
EDGX Stock Exchange: EDGX	Make/Take Exch.	65,885	2.23	0.01	3.20	0.10
BZX Stock Exchange: BZX	Make/Take Exch.	94,401	2.28	0.01	3.50	0.01
Archipelago Stock Exchange: ARCA	Make/Take Exch.	6,701	3.04	0.03	4.60	0.05
ELPs		122,462	2.24	0.01	2.24	0.01
Make/Take Exchanges		1,300,855	2.15	0.00	3.08	0.00
Inverted Exchanges		5,444	2.34	0.03	2.13	0.03
Other Dark Pools		7,419	2.21	0.03	2.21	0.03
Broker's Dark Pool		93,745	4.33	0.02	4.33	0.02

Table 4: Multivariate analysis of the relationship between trading costs and ELP exposure

Notes: For a given child order execution, effective spread is measured as the difference between the trade price and the execution-time NBBO midpoint divided by the execution-time NBBO midpoint. Net effective spread is the effective spread net of the assumed fee or rebate charged by the executing venue. Implementation Shortfall is computed as the normalized difference between the volume-weighted average child order execution price and the NBBO midpoint prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. ELP Exposure is the percentage of a parent order's trades with ELPs. BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in dark pools that our data provider does not own). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:			
	Effective Spread	Net Effective Spread	Implementation Shortfall	Net Implementation Shortfall
ELP Exposure	0.14	−0.45**	44.45***	43.85***
	(0.26)	(0.23)	(8.56)	(8.56)
BODP Exposure	−1.48***	−2.21***	−7.81*	−8.53**
	(0.21)	(0.22)	(4.13)	(4.13)
Other DP Exposure	−0.40	−1.36	−39.20	−39.78
	(0.64)	(0.85)	(29.06)	(29.06)
PasExch Exposure	−3.29***	−4.23***	6.37	5.35
	(0.17)	(0.18)	(5.24)	(5.24)
Participation Rate	1.99***	2.28***	52.09***	52.33***
	(0.58)	(0.62)	(19.85)	(19.85)
Log (Price)	−0.94***	−1.35***	4.14	3.78
	(0.31)	(0.33)	(4.79)	(4.79)
Volatility	9.20	6.00	56.37	56.99
	(12.07)	(15.58)	(314.67)	(314.74)
Duration	0.05	0.09	−3.99	−3.98
	(0.08)	(0.09)	(8.68)	(8.68)
Turnover	−0.002	−0.003	0.37	0.37
	(0.003)	(0.004)	(0.38)	(0.38)
N	20,335	20,335	20,335	20,335
Adjusted R^2	0.13	0.19	0.10	0.10

Table 5: Summary statistics for parent orders submitted by clients who avoid ELPs and their matches

Notes: ELP Exposure is the percentage of a parent order's trades with ELPs. BODP Exposure is the percentage of a parent order's child executions that occur in the broker's own dark pool. PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Quoted Spread is the time-weighted percentage bid-ask spread over the parent order's life. Standard errors are adjusted by double-clustering on stock and day.

Statistic	No ELP	ELP	Difference	p-value
# of parent orders	364	364		
ELP Exposure (%)	0.00	6.31	−6.31	<0.01
BODP Exposure (%)	2.99	10.11	−7.12	<0.01
PasExch Exposure (%)	37.0	31.2	5.80	<0.01
Participation Rate (%)	2.41	1.22	1.19	<0.01
Log (Price)	3.934	3.935	−0.001	0.98
Volatility (%)	1.51	1.60	0.09	0.28
Quoted Spread (bps)	4.10	4.20	−0.10	0.50

Table 6: Trading costs of comparable parent orders with and without exposure to ELPs.

Notes: Implementation Shortfall is computed as the normalized difference between the average child order execution price and the NBBO midpoint prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. Not Exposed to ELPs is an indicator variable that is set equal to one if the client submitting the parent order does not source ELP liquidity. BODP Exposure is the percentage of a parent order's child executions that occur in the broker's own dark pool. PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Duration is the fraction of the trading day that the parent order is worked. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:	
	Implementation Shortfall	Net Implementation Shortfall
Not Exposed to ELPs	−12.22**	−12.07**
	(5.69)	(5.69)
BODP Exposure	−19.44	−20.23
	(19.85)	(19.86)
PasExch Exposure	24.33	22.65
	(18.97)	(18.97)
Participation Rate	157.50*	156.57*
	(83.13)	(83.16)
Duration	2.85	2.81
	(7.84)	(7.84)
N	728	728
Adjusted R^2	0.01	0.01

Table 7: Difference-in-differences regression of trading costs incurred by the client who stops sourcing ELP liquidity.

Notes: All executions generated by parent orders in the *Exposed* group of orders are included in this analysis. Implementation Shortfall is computed as the normalized difference between the average child order execution price and the price of the asset prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. Post takes a value of 1 for executions after March 18, 2011. Switched takes a value of 1 for executions generated by the switching client's orders. BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in a dark pool that is not owned by the data provider). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:	
	Implementation Shortfall	Net Implementation Shortfall
Post × Switched	-32.27^{**}	-32.23^{**}
	(15.77)	(15.77)
Post	49.23^{***}	49.27^{***}
	(11.86)	(11.86)
Switched	-1.44	-1.38
	(10.97)	(10.97)
BODP Exposure	-13.56^{***}	-14.19^{***}
	(3.85)	(3.85)
Other DP Exposure	-58.50^{**}	-58.94^{**}
	(29.16)	(29.16)
PasExch Exposure	-0.09	-1.01
	(4.69)	(4.69)
Participation Rate	46.13^{***}	46.47^{***}
	(15.96)	(15.96)
Log (Price)	1.45	1.15
	(4.34)	(4.34)
Volatility	67.60	68.25
	(252.00)	(252.00)
Duration	-7.60^{*}	-7.56^{*}
	(3.95)	(3.95)
Turnover	0.49^{*}	0.49^{*}
	(0.29)	(0.29)
N	19,386	19,386
Adjusted R^2	0.10	0.10

Table 8: The cost of early versus late exposure to ELPs

Notes: Implementation Shortfall is computed as the normalized difference between the average child order execution price and the price of the asset prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. ELP Exposure D1 (D10) is the percentage of child executions by ELPs in the first (last) ten percent of the parent order's executed volume. ELP Exposure D2_D9 is defined analogously. BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in a dark pool that is not owned by the data provider). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:	
	Implementation Shortfall	Net Implementation Shortfall
ELP Exposure D1	55.84***	55.21***
	(9.45)	(9.45)
ELP Exposure D10	26.58***	25.93***
	(9.14)	(9.15)
ELP Exposure D2_D9	33.04***	32.40***
	(9.70)	(9.71)
BODP Exposure	−7.20	−7.97
	(4.44)	(4.44)
Other DP Exposure	−43.33	−43.97
	(31.77)	(31.78)
PasExch Exposure	8.26	7.18
	(5.49)	(5.49)
Participation Rate	59.98***	60.23***
	(19.86)	(19.87)
Log (Price)	5.22	4.85
	(4.98)	(4.97)
Volatility	155.34	155.93
	(326.57)	(326.63)
Duration	−4.07	−4.06
	(7.66)	(7.66)
Turnover	0.50	0.50
	(0.32)	(0.32)
N	20,335	20,335
Adjusted R^2	0.11	0.11

Table 9: Multivariate analysis of the relationship between child order trading costs and ELPs

Notes: For a given child order execution, PreCost is measured as the normalized difference between the trade price and volume-weighted average price realized over the 30 seconds prior to the child order execution. ELP is a binary variable taking a value of 1 if the child order is executed by an ELP. BODP (OtherDP) is a binary variable taking a value of 1 if the child order is executed by the broker's own (other) dark pool. PasExch is a binary variable taking a value of 1 if the child order is executed passively by an exchange. Log (TradeSize) is the logarithm of the dollar size of the child order. Standard errors are given in parentheses and are adjusted by double-clustering on stock and day.

	Dependent variable:			
	PreCost (bps)			
	(1)	(2)	(3)	(4)
ELP	3.68***	3.66***	3.54***	3.53***
	(0.12)	(0.12)	(0.11)	(0.11)
BODP	2.64***	2.65***	2.59***	2.62***
	(0.14)	(0.14)	(0.12)	(0.12)
OtherDP	3.22***	3.24***	2.30***	2.35***
	(0.30)	(0.30)	(0.27)	(0.27)
PasExch	−0.57***	−0.58***	−0.62***	−0.64***
	(0.03)	(0.03)	(0.03)	(0.03)
Log (TradeSize)		−0.07**		−0.13***
		(0.03)		(0.01)
Stock FE	No	No	Yes	Yes
Date FE	No	No	Yes	Yes
Client FE	No	No	Yes	Yes
N	2,588,910	2,588,910	2,588,910	2,588,910
Adjusted R^2	0.01	0.01	0.02	0.02

Table 10: The cost of exposing to multiple ELPs.

Notes: Implementation Shortfall is computed as the normalized difference between the average child order execution price and the price of the asset prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. More than 2 ELPs (More than 3 ELPs) is an indicator variable that equals one if the parent order sourced liquidity from at least two different ELPs (three different ELPs). BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in a dark pool that is not owned by the data provider). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:	
	Implementation Shortfall	Net Implementation Shortfall
More than 2 ELPs	5.51***	5.51***
	(1.41)	(1.41)
More than 3 ELPs	5.93**	5.95**
	(2.77)	(2.77)
ELP Exposure	33.30***	32.69***
	(8.26)	(8.26)
BODP Exposure	−5.24	−5.96
	(4.06)	(4.06)
Other DP Exposure	−32.15	−32.73
	(28.54)	(28.54)
PasExch Exposure	7.26	6.24
	(5.21)	(5.21)
Participation Rate	48.34**	48.58**
	(19.96)	(19.96)
Log (Price)	4.84	4.48
	(4.81)	(4.82)
Volatility	43.25	43.85
	(312.86)	(312.92)
Duration	−6.01	−6.00
	(8.58)	(8.58)
Turnover	0.37	0.37
	(0.37)	(0.37)
N	20,335	20,335
Adjusted R^2	0.10	0.10

Table A.1: Fee assumptions (per share) for child order executions in each venue.

Notes: Negative numbers refer to rebates.

Venue	Type	Liquidity Providing Trades	Liquidity Demanding Trades
Archipelago Stock Exchange: ARCA	Make/Take Exch.	−\$0.0025	\$0.0030
Archipelago Stock Exchange: ARCX	Make/Take Exch.	−\$0.0025	\$0.0030
BZX Stock Exchange: BZX	Make/Take Exch.	−\$0.0020	\$0.0030
EDGX Stock Exchange: EDGX	Make/Take Exch.	−\$0.0020	\$0.0030
Nasdaq Stock Exchange: Nasdaq	Make/Take Exch.	−\$0.0020	\$0.0030
New York Stock Exchange: NYSE	Make/Take Exch.	−\$0.0014	\$0.0027
EDGA Stock Exchange: EDGA	Low Fee Exch.	\$0.0003	\$0.0003
Virtu Americas: GFLO	ELP	n.a.	\$0.0000
Citadel: CDRG	ELP	n.a.	\$0.0000
Trimark: TRIM	ELP	n.a.	\$0.0000
Knight Securities: NITE	ELP	n.a.	\$0.0000
D.E. Shaw: SHAW	ELP	n.a.	\$0.0000
Two Sigma Securities: SOHO	ELP	n.a.	\$0.0000
Sun Trading: FSOM	ELP	n.a.	\$0.0000
BIDS ATS: BIDS	Other D.P.	\$0.0000	\$0.0000
Level ATS: EBXL	Other D.P.	\$0.0000	\$0.0000
Broker-owned dark pool: BODP	Broker's D.P.	\$0.0000	\$0.0000
Nasdaq BX Stock Exchange: BX	Inverted Exch.	\$0.0020	−\$0.0006
BYX Stock Exchange: BYX	Inverted Exch.	\$0.0018	−\$0.0008

Table A.2: Multivariate relationship between implantation shortfall and exposure to GETCO and the other ELPs.

Notes: Implementation Shortfall is computed as the normalized difference between the average child order execution price and the price of the asset prior to the start of the execution. Net Implementation Shortfall is adjusted for assumed fees and rebates. GETCO Exposure is the percentage of child executions by GETCO. Other ELP Exposure is the percentage of child executions by ELPs other than GETCO. BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in a dark pool that is not owned by the data provider). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:	
	Implementation Shortfall	Net Implementation Shortfall
GETCO Exposure	34.92***	34.31***
	(8.92)	(8.92)
Other ELP Exposure	89.67***	89.10***
	(15.61)	(15.60)
BODP Exposure	−7.26*	−7.98*
	(4.14)	(4.13)
Other DP Exposure	−37.35	−37.93
	(29.05)	(29.05)
PasExch Exposure	6.82	5.81
	(5.26)	(5.25)
Participation Rate	54.51***	54.75***
	(19.88)	(19.88)
Log (Price)	4.15	3.78
	(4.76)	(4.76)
Volatility	50.49	51.11
	(314.23)	(314.30)
Duration	−4.08	−4.07
	(8.69)	(8.69)
Turnover	0.37	0.37
	(0.38)	(0.38)
N	20,335	20,335
Adjusted R^2	0.10	0.10

Table A.3: Multivariate analysis of the relationship between VWAP slippage and ELP exposure.

Notes: VWAP Slippage is computed as the normalized difference between the average child order execution price and the volume-weighted average price observed in the market during the parent order's lifetime. Net VWAP Slippage is adjusted for assumed fees and rebates. ELP Exposure is the percentage of a parent order's trades with ELPs. BODP Exposure (Other DP Exposure) is the percentage of a parent order's child executions that occur in the broker's own dark pool (in a dark pool that is not owned by the data provider). PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Duration is the fraction of the trading day that the parent order is worked. Turnover is the ratio of the number of shares traded during the life of the parent order to the outstanding number of shares in thousands. Each regression includes 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:	
	VWAP Slippage	Net VWAP Slippage
ELP Exposure	4.70***	4.10***
	(1.27)	(1.28)
BODP Exposure	−0.96*	−1.68***
	(0.54)	(0.53)
Other DP Exposure	9.90	9.32
	(7.69)	(7.71)
PasExch Exposure	−1.85**	−2.87***
	(0.78)	(0.79)
Participation Rate	3.15*	3.40**
	(1.68)	(1.69)
Log (Price)	0.46	0.09
	(1.05)	(1.08)
Volatility	142.92*	143.56*
	(82.81)	(82.79)
Duration	−0.28	−0.27
	(0.78)	(0.78)
Turnover	0.01	0.01
	(0.06)	(0.06)
N	20,335	20,335
Adjusted R^2	0.08	0.08

Table A.4: Summary statistics for parent orders submitted by the switching client before and after the switch date

Notes: ELP Exposure is the percentage of a parent order's trades with ELPs. BODP Exposure is the percentage of a parent order's child executions that occur in the broker's own dark pool. PasExch Exposure is the percentage of a parent order's child executions that provide liquidity on a lit stock exchange. Participation Rate is measured as the ratio of the parent order's trading volume to the overall trading volume of the underlying stock over the period of time that the parent order is being worked. Log (Price) is the log of the midpoint of the NBBO prevailing when the parent order is placed with the broker. Volatility is measured as the volatility of the midpoint of the NBBO over the parent order's life. Quoted Spread is the time-weighted percentage bid-ask spread over the parent order's life. We compute the z-scores by normalizing the raw variables with their daily means and standard deviations. Standard errors are adjusted by double-clustering on stock and day.

Statistic	Pre	Post	Difference	p-value
# of parent orders	40	111		
ELP Exposure (%)	6.19	0.00	6.19	<0.01
BODP Exposure (%)	4.22	4.13	0.09	0.95
PasExch Exposure	0.39	0.36	0.03	0.21
Participation Rate (%)	2.05	2.05	0.00	0.99
Log (Price)	3.853	3.813	0.039	0.70
Volatility (z-score)	-0.13	-0.13	0.00	0.99
Quoted Spread (z-score)	-0.03	0.06	-0.09	0.55