

Sagade, Satchit; Scharnowski, Stefan; Westheide, Christian

Working Paper

Broker colocation and the execution costs of customer and proprietary orders

SAFE Working Paper, No. 366

Provided in Cooperation with:

Leibniz Institute for Financial Research SAFE

Suggested Citation: Sagade, Satchit; Scharnowski, Stefan; Westheide, Christian (2022) : Broker colocation and the execution costs of customer and proprietary orders, SAFE Working Paper, No. 366, Leibniz Institute for Financial Research SAFE, Frankfurt a. M., <https://doi.org/10.2139/ssrn.4289346>

This Version is available at:

<https://hdl.handle.net/10419/266687>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Satchit Sagade | Stefan Scharnowski | Christian Westheide

Broker Colocation and the Execution Costs of Customer and Proprietary Orders

SAFE Working Paper No. 366 | December 2022

Leibniz Institute for Financial Research SAFE
Sustainable Architecture for Finance in Europe

info@safe-frankfurt.de | www.safe-frankfurt.de

Electronic copy available at: <https://ssrn.com/abstract=4289326>

Broker collocation and the execution costs of customer and proprietary orders*

Satchit Sagade^{1,2}, Stefan Scharnowski³, Christian Westheide^{†2,4}

¹*Nasdaq*

²*Leibniz Institute for Financial Research SAFE*

³*University of Mannheim*

⁴*University of Vienna*

This version: December 13, 2022

Abstract

Colocation services offered by stock exchanges enable market participants to achieve execution costs for large orders that are substantially lower and less sensitive to transacting against high-frequency traders. However, these benefits manifest only for orders executed on the collocated brokers' own behalf, whereas customers' order execution costs are substantially higher. Analyses of individual order executions indicate that customer orders originating from collocated brokers are less actively monitored and achieve inferior execution quality. This suggests that brokers do not make effective use of their technology, possibly due to agency frictions or poor algorithm selection and parameter choice by customers.

Keywords: Execution Cost, Institutional Investor, Broker, High-Frequency Trading, Colocation

JEL classification: G10, G14, G15

*We thank Björn Hagströmer, Terrence Hendershott, Jan-Pieter Krahen, Phil Mackintosh, Albert Menkveld, Lorian Pelizzon, Tālis Putniņš, Stefan Rünzi, Michael Schneider, Erik Theissen, and conference and seminar participants at the 26th Annual Meeting of the German Finance Association, at ESMA, and at Goethe University Frankfurt for helpful comments and discussions, and Cyrille Guillaumie for assistance with the data. The authors gratefully acknowledge financial support from the Research Center SAFE, funded by the State of Hessen research initiative LOEWE. The authors thank ESMA for providing access to their high-frequency database. Access to Satchit Sagade was provided only while he was employed at the University of Frankfurt and the Research Center SAFE and before his employment at Nasdaq. An earlier version of this paper appeared as chapter 2 of Stefan Scharnowski's doctoral dissertation at the University of Mannheim. The analysis, interpretation, and conclusion of this study reflect solely the views of the authors. *Email addresses:* sagade@safe.uni-frankfurt.de (Satchit Sagade), scharnowski@uni-mannheim.de (Stefan Scharnowski), christian.westheide@univie.ac.at (Christian Westheide)

[†]Corresponding author: Christian Westheide, Department of Finance, Faculty of Business, Economics and Statistics, University of Vienna, Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria. Phone: +43-1-4277-37505

1 Introduction

“[B]rokers and other industry players are learning from [...] HFT strategies and in turn offering tools to their clients[...] that take advantage of the same sophisticated technology and logic. We’re all high frequency traders now.”

Credit Suisse Whitepaper ([Avramovic et al., 2017](#))

Equity market participants invest heavily in the speed of market access by employing cutting-edge hardware, relying on real-time market data feeds, and subscribing to exchanges’ colocation facilities in order to place their servers within the exchanges’ datacenters in close proximity to the matching engine. While the academic literature has explored the role of trading speed in the context of high-frequency trading (HFT), the impact of the speed of institutional investors and broker-dealers on their trading outcomes remains unexplored.¹

We investigate how the execution costs for large orders traded in a proprietary or agency capacity are affected by the use of exchanges’ colocation facilities. Trades executed by collocated exchange members in a proprietary capacity achieve about 6 basis points (henceforth bps) lower execution costs compared to similar orders executed without colocation. However, agency orders do not benefit to any extent from brokers’ colocation. After controlling for commonly used predictors of execution costs (market conditions, order characteristics, and trading strategy), agency orders executed by collocated brokers experience higher costs as compared to proprietary ones. The difference in execution costs between the two capacities is almost 5 bps across all exchange members and 9 bps within the subsample of broker-dealers

¹For the remainder of the paper, we use the terms exchange member, broker, and broker-dealer interchangeably.

who execute orders in both agency and proprietary capacities.

The magnitude of these differences in execution costs is economically large compared to average quoted spreads of 9 bps for the largest tercile of stocks, which account for the majority of trading volume in our sample. A back-of-the envelope calculation suggests that the incremental costs incurred by agency compared to proprietary orders executed by colocated brokers are economically sizable at more than 300 million euros annually.²

We also find that the execution costs of brokers using collocation for proprietary orders are less affected by the extent to which they interact with the flow of HFT firms. Trading against HFT firms is generally associated with higher execution costs, though this is driven by transactions against HFTs acting as aggressive counterparties, whereas the opposite holds for passive, liquidity providing HFT counterparties. The magnitude of these effects is around 3.5 bps and -3.0 bps for a 10 percentage point increase in the share of the aggressive and passive HFT as counterparties to the large order, respectively. Conditional on exchange members being colocated, the effect of aggressive HFT on execution costs of proprietary orders is reduced by more than half but remains positive, whereas the conditional effect of passive HFT is smaller in magnitude but remains negative. These results indicate that access to speed-enhancing facilities potentially allows large traders to reduce the sensitivity of execution outcomes to the nature and technological sophistication of their counterparties.

There are several potential explanations for the fact that, while collocation allows broker-dealers to obtain better execution outcomes for large orders, these benefits are limited to their proprietary orders. To the extent that broker-dealers provide the same algorithmic

²This number is obtained by multiplying the average trade size by the number of colocated agency trades, the difference in execution costs, the ratio of the length of a year relative to our sample period, and the ratio of the number of stocks relative to the number contained in our sample.

execution strategies to their customers that they use when trading in a principal capacity, the differences in execution outcomes indicate the presence of inefficiencies in the choice of algorithm and/or the parameters governing such algorithms employed by the customers. The use of algorithms developed by third party providers or by large buy-side investors in-house (see, for example, [Frazzini et al., 2018](#)) – possibly inferior to algorithms provided by brokers – and then deployed using an exchange member’s direct electronic access infrastructure may also help explain the results. Finally, an alternative or supplemental explanation of our results may be that the choice of execution algorithms brokers make available to customers is different to those used for their proprietary trades.

We examine two additional reasons related to the objectives the algorithms might have. First, prior literature ([Battalio et al., 2016](#); [Anand et al., 2021](#)) shows that brokers try to minimize exchange fees, likely because these fees cannot be passed on to the customer, and that such behavior leads to inferior execution outcomes. We do not find that brokers in our sample attempt to minimize exchange fees for agency compared to proprietary trades. Second, customers might care not only about average execution costs but also about their uncertainty, and such preferences might be incorporated in execution algorithms. However, we do not find the execution costs of agency orders to be more predictable than those of proprietary orders.

To better understand the channels behind our main results, we investigate the order submission behavior and the execution quality at the level of the child orders that constitute the large parent orders. In an algorithmic trading setup, one major characteristic of an algorithm is its ability to monitor orders actively with low latency. We find that, compared to colocated agency orders, the proprietary counterparts on average have higher order-to-

trade ratios. This suggests that the algorithms used to execute proprietary orders monitor the limit order book more actively – possibly in event time as opposed to calendar time (Bacidore, 2020). We furthermore find a higher degree of periodicity in colocated agency as compared to proprietary orders, which is consistent with a more prevalent use of timer-based algorithms for agency trades. Brugler (2015) and Sağlam (2020) also provide evidence of calendar-time periodicity in trading activity, which in Sağlam (2020) is associated with poor execution quality. Finally, the individual child order executions of proprietary orders experience better execution quality when measured in terms of effective spreads and price impacts in comparison to observationally similar agency orders. These results altogether suggest that the quality of child order executions for agency orders is inferior compared to that of proprietary orders.

It is unclear whether our results constitute evidence for a formal violation of brokers' best execution obligations for at least two reasons. First, in contrast to the US, brokers in the EU have significant flexibility when defining their best execution policies. Second, to the extent customers make poor choices when using the algorithms provided by the brokers, brokers are not in violation of any obligations. However, if the clients make these poor execution decisions, brokers could better educate their customers to eliminate such inefficiencies. Regardless of whether our results are interpreted as consistent with best execution, our approach of benchmarking client execution outcomes against brokers' own orders may be useful for brokers to demonstrate, and regulators to verify, best execution.

Our results are based on a pan-European equity market dataset made available by the European Securities and Markets Authority (ESMA). The data contains exchange message-level information including masked identifiers of exchange members whose activity can be

tracked across stocks and exchanges. We aggregate individual order executions to 11,724 parent orders based on unidirectional order flow of non-HFT exchange members and examine the determinants of execution costs for these parent orders using a doubly robust estimator combining panel regressions with inverse probability weights. Our results also hold for simple fixed effects models and are robust to including exchange fees in the execution costs.

A potential concern in our analysis is that our regression design does not fully control for the differences in trading motives or investment styles across agency and proprietary orders. While investors' trading motives are not directly observable in our dataset, our choice of control variables and fixed effects indirectly controls for trading motives and differences in execution strategy originating from heterogeneous trading motives. Specifically, our regressions are estimated within stock-day and also control, *inter alia*, for the time of day the execution begins. Additionally, we control for the volatility and short-term return of the stock immediately before trading, the size of the parent order, its total time of execution, its number of executed child orders, the proportion of liquidity-consuming child orders, the total trading volume in the stock, the proportion of trading executed away from the primary exchange, and the information content of the order. Thus, we consider it unlikely that unobserved trading motives would affect execution costs to any substantial extent.

Our paper contributes to the literature on brokers' order handling practices and their impact on institutional execution costs. [Barbon et al. \(2019\)](#) find evidence that brokers leak information on institutional trades to other customers, which leads to increased execution costs for those trades. [Conrad et al. \(2001\)](#) show that institutional investors' use of soft dollar brokers, i.e., those providing sell-side research in exchange for executing trades, is associated with higher transaction costs. [Anand et al. \(2011\)](#), while mainly concerned

with the performance of trading desks of institutional investors, find that broker execution performance is persistent and related to trading commissions. [Anand et al. \(2021\)](#) show that institutional execution costs are higher for brokers sending a large share of orders to affiliated alternative trading systems. [Battalio et al. \(2018\)](#) study the effect of routing orders to off-exchange high-frequency liquidity providers and find that such behavior leads to higher execution costs on institutional orders. However, apart from trading venue choices, prior literature provides little evidence on how and why execution costs differ across brokers. Our paper fills this gap by focusing on the role of technology, specifically collocation, as a differentiating factor that enables the effective use of sophisticated algorithms.

Our paper also contributes to the literature on speed dispersion. This literature finds that the effect of faster HFTs on slower non-HFTs depends on whether the former use their speed advantage to provide liquidity (as in [Ait-Sahalia and Saglam, 2017](#); [Brogaard et al., 2015](#)) or to engage in arbitrage or back-running (as in [Foucault et al., 2016](#); [Shkilko and Sokolov, 2020](#)). Our paper, while obtaining results consistent with both views, shows that broker-dealers' investment in enhancing their speed of market access can reduce the sensitivity of execution costs to trading against HFTs.

We also contribute to the literature studying the effects of HFT on the trading activity and associated costs of institutional investors.³ [Yang and Zhu \(2020\)](#) extend the two-period [Kyle \(1985\)](#) model by including a “back-runner” – a trader who infers the fundamental information by observing the order flow – alongside the informed trader. The former's presence leads to less aggressive trading by the informed trader, thereby delaying price

³More broadly, prior research shows that HFT has a positive or, at least, benign effect on common measures of market liquidity, such as bid-ask spreads, and price efficiency, such as variance ratios (e.g. [Hasbrouck and Saar, 2013](#); [Menkveld, 2013](#); [Brogaard, Hendershott, and Riordan, 2014](#)). See [Menkveld \(2016\)](#) for a survey.

discovery.⁴ [Korajczyk and Murphy \(2019\)](#) show that, in the Canadian equity market, the presence of HFT firms engaging in market-making activity leads to higher transaction costs for large institutional trades. [Van Kervel and Menkveld \(2019\)](#) find that, in the Swedish equity market, HFTs provide liquidity early-on during institutional trades but later turn to trade in the same direction. [Sağlam \(2020\)](#) observes that execution algorithms such as the Volume-Weighted Average Price (VWAP) can lead to predictable patterns in order flow that can be picked up by HFTs. [Putniņš and Barbara \(2016\)](#) find that, in the Australian equity market, the effects of HFT firms on institutional transaction costs differ in the cross-section of HFT firms, with some affecting them positively and others negatively. [Tong \(2015\)](#), in a study of the U.S. equity market, obtains evidence that institutional execution costs increase with the amount of HFT. She also observes that this effect is alleviated for some institutional investors with high levels of trading skills, as measured by historical transaction costs, though her data do not allow her to explore the determinants of institutional trading skills. [Chen and Garriott \(2020\)](#) find that HFT in the Canadian bond futures market leads to lower institutional execution costs for relatively small parent orders. Finally, [Brogaard, Hendershott, Hunt, and Ysusi \(2014\)](#) do not find any causal effect of HFT on institutional transaction costs on the London Stock Exchange. Our paper contributes to this literature by establishing differential effects of trading against aggressive and passive HFT orders, and

⁴In addition to HFTs affecting the profitability of investors' trading strategies and price discovery in the market, several studies have shown that rent seeking behaviour by HFTs negatively affects investors' information acquisition decision. [Baldauf and Mollner \(2020\)](#) model a fragmented market in which HFTs can both demand and supply liquidity. HFTs anticipate other participants' order flow and, as they become faster, information acquisition decreases. [Dugast and Foucault \(2018\)](#) show that the existence of inexpensive unprocessed but imprecise information can lead market participants to reduce their demand for processed and more accurate information, leading to a decrease in price informativeness. [Weller \(2017\)](#), in a study of algorithmic trading – a superset of HFT – in the U.S. equity market, and [Gider, Schmickler, and Westheide \(2021\)](#), in an international study on HFT, find evidence consistent with this idea.

by showing that brokers can reduce these effects through their own trading speed.

The remainder of the paper is structured as follows. [Section 2](#) describes the institutional background. [Section 3](#) presents the data and discusses the variables used. In [Section 4](#), we discuss the measurement and potential determinants of execution costs. The estimation approach is detailed in [section 5](#). [Section 6](#) contains descriptive statistics at the levels of exchange members and parent orders. In [Section 7](#), we present the parent order level results. [Section 8](#) contains analyses of order submissions and the execution quality of child orders, as well as possible motives explaining our key results. [Section 9](#) concludes.

2 Institutional Background

In this section, we discuss the institutional arrangements, the trading process, and the relationship between the sell-side (investment banks and specialized execution brokers) and the buy-side investors, as well as proprietary trading by exchange members, thus providing background information that helps understand the channels behind our results.

Large institutional orders are predominantly handled by specialized algorithmic execution desks provided by the sell-side who develop the algorithms that implement the execution strategy for an order across venues and over time with the aim of meeting the investors' objectives. These sell-side banks are exchange members and thereby provide the buy-side investors with access to the different trading venues. A smaller amount of trading volume by the sell-side on behalf of the buy-side is conducted manually ("high-touch" as opposed to the algorithmic "low-touch") and relies on the traders' individual skills rather than technological sophistication. Our data do not allow us to distinguish high-touch from low-touch executions,

though firms that invest in colocation presumably mostly trade algorithmically.

Exchange members may also trade on their own account. They may do so with the goal of profiting from taking positions in anticipation of price changes or by engaging in market-making strategies. Some proprietary trading may happen with the goal of acquiring or disposing of a position after agreeing to trade against the customer at a reference price (e.g., the value-weighted average price, VWAP). In such cases, while the trade happens to facilitate a customer trade, the sell-side tries to earn a profit by trading at prices superior to the reference price. There also is a subset of exchange members that have no customers at all but instead purely trade on their own account so as to earn a profit. Some of these firms are HFTs, others may pursue longer-horizon strategies and deliberately take on relatively large positions in individual securities.

The relationship between a buy-side investor and a sell-side firm determines who controls the execution strategy. For instance, some sell-side firms provide a suite of algorithmic execution strategies to their customers who then choose the appropriate strategy and the underlying parameters for a given order. Alternatively, the broker may have more discretion with respect to the execution strategy and make these choices on the investors' behalf.

There exist information and agency frictions between the buy-side and sell-side. For instance, brokers' willingness to continuously invest in trading speed depends on the extent to which they can recover the resulting fixed costs from end investors. This will be possible only if the buy-side has both sufficient information on differences in execution costs across brokers and is motivated to minimize execution costs. Since proficient transaction cost analysis requires large amounts of data and an investment in analytical capabilities, many members of the buy-side are unlikely to identify differences in execution quality that are economically

meaningful but, due to noise in the data, not immediately obvious. Furthermore, investment banks' services other than execution services desired by buy-side investors (e.g., access to analyst research) may mean that execution quality is not the decisive factor in choosing an execution provider due to soft dollar arrangements.⁵

If a sell-side firm has already made an investment in colocation on a given exchange, there are additional constraints that accompany the use of this colocated access. For instance, any connection to an exchange comes with limitations to its bandwidth, i.e., the amount of data sent and received via the connection has limits. Thus, an expansion of the use of a firm's colocation access may require the purchase of additional servers and expenses for additional connections to the exchange's matching engine. If it is hard for the buy-side to observe how a sell-side firm's connections to the exchange are used, there will be incentives for the sell-side to limit the intensity of the use of the connections when engaging in agency trading. Conversely, they will be used to the maximum extent for proprietary trading because the benefit thereof accrues directly to the exchange member.

These frictions drive the differences in execution costs not only between colocated and non-colocated brokers, but also between agency and proprietary orders originating from colocated brokers.

In some cases, especially those involving the largest buy-side investors, the algorithms used may be developed by the investors themselves and employed using an investment banks' technological infrastructure through a direct electronic access arrangement (DEA) setup (see, e.g., [Frazzini et al., 2018](#)). In such cases, the brokers have little to no control over the

⁵Regulations governing best execution arrangements seek to minimize the costs associated with such frictions. For instance, requirements for brokers to provide execution quality statistics to their customers help reduce the information frictions associated with access to data. However, it is unlikely that these costs are completely eliminated.

execution strategy employed by the investor. If the buy-side investors' ability to develop execution algorithms were inferior to that of the sell-side, this skill gap could also explain part of the differences in execution costs. All these factors potentially explain the differences in execution outcomes between agency and proprietary orders in cases where the investment firm providing market access is different from the firm that develops the algorithmic execution strategies.

3 Data

Our analysis is based on the proprietary high-frequency database of the European Securities and Markets Authority (ESMA). This database was designed by the regulator to study HFT activity in fragmented equity markets in the EU. The original sample comprises 98 stocks from nine countries – Belgium, Germany, Spain, France, Ireland, Italy, the Netherlands, Portugal and the United Kingdom – that were members of the EU at the time. Depending on the size of the national markets, between 5 and 16 stocks are chosen for each country via a stratified sampling approach to ensure it is representative of the EU market with respect to market value, trading volume, and level of fragmentation. The database includes all order messages and trades for each stock from the respective primary market and the three main lit multilateral trading facilities (MTFs) – Chi-X, BATS, and Turquoise – for the 22 trading days in May 2013. The data allow us to track individual exchange members' activity across all trading venues. [Bouveret et al. \(2014\)](#) provide detailed information on the sample selection approach and the database construction, as well as a first analysis of HFT activity

based on the data.⁶

3.1 Member Categorization, Trading Capacity, and Colocation

While individual members' identities remain unknown to us, ESMA groups them into three categories: HFTs, Investment Banks, and Others. Previous studies identify HFT firms either based on their primary business model or on measures of trading and quoting activity such as order-to-trade ratios or order lifetimes. [Bouveret et al. \(2014\)](#) compare these two identification approaches and find that the former provides a conservative estimate by only including proprietary trading firms and excluding HFT activity originating at investment banks. It also prevents the erroneous classification of sophisticated traders employing low latency infrastructure as HFTs, even though their business model substantially differs from that of HFTs. For example, [Avramovic et al. \(2017\)](#) argue that the quoting and trading behavior (speed of response to market events, order-to-trade ratios, etc.) of large technologically sophisticated broker-dealers may resemble that of HFTs, even though the size and holding period of their positions are vastly different. We hence rely on the HFT flag based on the former approach, which categorizes 20 exchange members as HFTs, 18 of which serve as counterparties to the institutional trades identified below.

Trading activity by individual exchange members is further split by capacity into trading on behalf of customers (agency) and trading on the firm's own account (proprietary). We rely on this flag to identify differences in execution quality between institutional trades executed by exchange members on behalf of their clients and on a proprietary basis. This flag is

⁶The dataset, however, excludes trading activity in off-exchange venues such as dark pools, systematic internalizers, and the over-the-counter market. To the best of our knowledge, statistics on institutional usage of such venues in the EU are unavailable. However, [Beason and Wahal \(2020\)](#) observe that, for a single firm providing algorithmic execution services in the US, over 77% of child orders are routed to lit venues.

generated by ESMA based on the reporting of this information by trading venues.⁷

Finally, the database also contains a binary variable indicating whether members use the colocation service offered by the different trading venues, i.e., whether members have located their equipment within the venues' data centers. Colocation, combined with a subscription to the venues' low latency market data feeds, allows members to minimize the round-trip latency between their servers and the venues' matching engines.⁸ We distinguish colocation on the primary markets from that on the MTFs. While we are able to directly observe whether an individual exchange member is colocated on the primary listing venue, we proxy for colocation on the MTFs based on whether the member is colocated on Turquoise. This is because we cannot directly observe such a flag for BATS Chi-X. We also exclude the 12 Spanish stocks from our sample as the colocation flag for the Spanish stock exchange is missing.

3.2 Identification of Institutional Orders

Large institutional parent orders are typically split into multiple smaller child orders and then executed over time and across venues. While our dataset allows us to track the message traffic (order submissions, cancellations, and transactions) of individual exchange members, we do not have information on the parent order's size. [Korajczyk and Murphy \(2019\)](#) and

⁷While [Bouveret et al. \(2014\)](#) report inconsistencies in the reporting of this information by the venues, those issues were subsequently fixed. This has been confirmed to us by ESMA. Furthermore, the reporting of riskless principal orders by some brokers as proprietary may introduce some noise in our results. However, we do not expect the magnitude of this noise to be substantial or to bias our results.

⁸We acknowledge that colocation, potentially, could also be used as a cost-efficient alternative for foreign firms to establishing an office in a country. If that were true for a subset of the exchange members in our sample, it would attenuate our results related to the differences in execution costs between orders originating from colocated versus non-colocated members. Furthermore, EU's single passport rights allows non-EU firms to trade across the European single market by using an office in one EU-member country, thus eliminating the need to colocate on multiple venues. However, the most active exchange members in our sample are in fact colocated across multiple venues.

Van Kervel and Menkveld (2019) solve this problem by stitching together a string of child order executions.

Korajczyk and Murphy (2019) define an institutional parent order as an uninterrupted sequence of one or more trades on the same side by the same exchange member if the total volume of those trades is at least CAD 100,000.⁹ Van Kervel and Menkveld (2019), on the other hand, start by first lumping together all executions on a given stock-day and then applying a directionality (the difference of buy and sell volume relative to total volume) cutoff of 90%. Both combine parent orders across days if there is a child order executed during both the last and first 30 minutes of two consecutive days. We employ the methodology of Korajczyk and Murphy (2019) to identify parent orders by applying an order size cutoff of € 100,000. Following Putniņš and Barbara (2016), we further require that a parent order is worked in the market for at least two hours.¹⁰ For our final sample of institutional orders, we additionally require an exchange member to have at least ten parent orders across all stocks. We use the capacity flag to further classify orders into those executed on behalf of a customer (agency execution) or on the exchange member's own account (proprietary execution) and construct parent orders separately for each capacity such that a single parent order consists of agency or proprietary child orders only.

⁹Korajczyk and Murphy (2019) use this cutoff based on the distribution of institutional order size in the US market as reported by Chan and Lakonishok (1995) and Cready et al. (2014).

¹⁰All parent orders identified using our approach originate from non-HFT members. This is unsurprising considering HFT firms generate their profits not by acquiring large positions, but by rapidly turning over their positions at high-frequencies, while providing liquidity and/or exploiting short-term informational asymmetries.

4 Measuring and Explaining Execution Costs

In this section, we explain how we measure the execution costs of large orders. We also describe the commonly-used determinants of execution costs that act as controls in our main analysis.

4.1 Measuring Parent Order Level Execution Costs

[Perold \(1988\)](#) defines implementation shortfall as the cost of implementing investment decisions and further decomposes it into execution costs and opportunity costs. The former relates to transactions actually executed while accounting for the fact that individual child orders may execute at different prices. Implicitly, the execution cost component includes instantaneous transaction costs such as the bid-ask spread as well as the price impact associated with the actions taken by other market participants upon detecting the presence of a large institution in the market as in [Brunnermeier and Pedersen \(2005\)](#), [Ait-Sahalia and Saglam \(2017\)](#), and [Yang and Zhu \(2020\)](#). The latter component relates to transactions that cannot be executed. As in most other studies of institutional transaction costs ([Keim and Madhavan, 1997](#); [Korajczyk and Murphy, 2019](#); [Van Kervel and Menkveld, 2019](#)), we ignore the opportunity cost of non-execution as we cannot observe the initially planned parent order size.

For parent order k executed by exchange member i in stock s on day t , the execution cost is defined as

$$\text{Execution Cost}_{istk} = \left(\frac{\text{VWAP}_{istk}}{P_{istk}^0} - 1 \right) \times Q_{istk} \quad (1)$$

where VWAP is the value-weighted average price of all child order executions, $Q \in \{-1, 1\}$ is the sign of the order which equals 1 (-1) for buy (sell) orders, and P^0 is the benchmark price. The quote midpoint immediately before the execution of the first child order serves as the benchmark price. We express the execution cost in basis points.

4.2 Determinants of Institutional Execution Costs

Our objective is to evaluate how brokers' use of colocation services affects the execution costs of their parent orders, and whether these effects differ for the brokers' proprietary versus agency executions. The use of exchanges' colocation facilities potentially allows exchange members to get better execution by strategically choosing when to supply and consume liquidity, better managing adverse selection costs, and reacting to changes in market conditions at high speeds. We construct a binary variable for each broker-stock combination based on whether the broker is colocated on any of the markets where the stock trades. A second binary flag associated with each message allows us to identify the capacity (agency or proprietary) in which the orders are submitted.

We control for several factors expected to affect institutional execution costs (see e.g. [Chan and Lakonishok, 1997](#); [Keim and Madhavan, 1997](#); [Jones and Lipson, 1999, 2001](#)). [Keim and Madhavan \(1997, 1998\)](#) argue that trading costs are driven by trade difficulty, which they proxy for using trade size, stocks' market capitalization, trader's investment style and other factors. We divide the previously established determinants of execution costs into three broad categories: order characteristics, market conditions, and execution strategy.

Several studies ([Keim and Madhavan, 1997](#); [Chan and Lakonishok, 1997](#); [Jones and Lipson, 1999](#)) find that institutional trading costs are higher for larger orders as counterparties require larger price concessions to take the other side of the trade. This is because displayed liquidity, even for the most actively traded stocks, is finite and may not fully absorb a large trade. Investors also perceive traders executing large quantities as better informed ([Easley and O'Hara, 1987](#)). To capture these effects, we include the parent order size as a control variable. Furthermore, [Chan and Lakonishok \(1995\)](#) and [Engle et al. \(2012\)](#) argue that trade size should be compared with the stock's typical trading volume when measuring execution costs. Higher overall volume in the market potentially allows the institution to better hide its trading intentions ([Kyle, 1985](#)). Hence, we also control for the total volume of all other executed orders across the primary market and the 3 MTFs in our sample during the execution of the parent order.

Previous studies (see for example [Chan and Lakonishok, 1993](#); [Saar, 2001](#); [Jones and Lipson, 2001](#)) have also shown that the execution costs of buy versus sell orders are different. For example, [Saar \(2001\)](#) predicts that institutional buy orders are more difficult to execute as they are perceived to be more informed than large sell orders. [Jones and Lipson \(2001\)](#), on the other hand, argue that buy orders tend to be cheaper to execute than sell orders. To capture any potential differences between buys and sells, we include a dummy variable for buy orders.

[Keim and Madhavan \(1996\)](#) find that establishing a long (short) position when prices are rising (falling) is more difficult due to potential information leakage before the order starts trading. We control for this pre-trade momentum by including the signed return of the stock during the 10-minute window before the first child order execution. Similarly, volatility in

the market influences the algorithms' trading strategy – for instance, how finely should a parent order be sliced or how aggressively should the order be worked in the market – as they seek to minimize the cost and risk of execution (Engle et al., 2012). Hence, we include pre-trade volatility, defined as the standard deviation of ten-millisecond quote midpoint returns during the 10-minute period before the first child order execution, as a control variable.

Previous studies have also shown that execution costs are affected by the fund manager's identity (Chan and Lakonishok, 1993, 1995) and their investment style (Keim and Madhavan, 1997; Jones and Lipson, 1999). Keim and Madhavan (1997) also observe substantial variation in trading costs within the same investment style which they attribute to trading skill, which is largely unobservable (see also Anand et al., 2011). These variables affect institutional execution costs as they are related to investors' need for immediacy, order aggressiveness, level of informational advantage, etc. While we cannot directly observe fund manager identity and investment style, we include several control variables that capture differences in trading behavior originating from heterogeneous investment styles.

To capture the degree to which an institutional investor is informed, we include the long term price impact as measured by the trade direction-signed return from the quote midpoint immediately before the parent order begins trading to the closing quote midpoint one day after the last child order execution. When determining their execution strategy, investors' urgency to trade likely influences how patient or aggressive they are while trading. For example, Keim and Madhavan (1997) find that value (technical) traders rely heavily on limit (market) orders and have lower (higher) transactions costs. Patient institutional investors may be able to lower their transaction costs by spreading their order flow over a longer time interval. Hence, we include the trade duration (measured as the number of

exchange trading hours between the first and the last child order execution) in our analysis.

Note that all of the aforementioned determinants of execution costs are directly determined by the investor's trading interest or the prevailing market conditions. In addition to these variables, there are choices the broker's order routing strategy can make to obtain the best execution outcome. In most analyses, we include variables of this kind that are expected to predict execution costs, such that the results we report as effects of trading in different capacities and with or without collocation are incremental to the choices made in these respects. Execution algorithms may attempt to trade passively by using non-marketable limit orders and thus earn the bid-ask spread and to split up the order into many child orders. Hence, we include the percentage of volume executed through aggressive orders, and the number of child order executions in a parent order. The choice of trading venue is another part of the execution strategy. To capture any differences in trading costs associated with this decision, we include the percentage of volume executed on the three MTFs. These latter characteristics of order routing strategies always have to be interpreted as a function of the, to some extent unobservable, difficulty of executing an order.¹¹ Thus, while they may be strong predictors of execution costs, their regression coefficients may not reflect causal effects. We also present results excluding these variables.

Finally, one potential channel through which collocation likely helps institutions is by allowing them to mimic HFT strategies and to minimize the sensitivity of their trading costs to trading against HFT. To test for this possibility, we include three measures capturing the interactions of HFT activity and institutional investors' trades: the fraction of the order exe-

¹¹E.g., *ceteris paribus*, splitting up a large order into many small pieces may help reduce execution costs. However, this decision may be made when executing the parent order at low cost is particularly difficult.

cuted against all HFTs, aggressive (liquidity-taking) HFTs, and passive (liquidity-providing) HFTs. Separating HFT into liquidity taking and liquidity supplying HFT allows us to capture the impact on execution costs associated with different HFT strategies. For example, liquidity supplying (consuming) HFTs are more likely to trade against (with) “the wind” (Van Kervel and Menkveld, 2019). As is the case for choices the broker has with respect to the way of executing a parent order, interaction with HFTs results endogenously from the trading process. Thus, coefficients estimated in the regressions should be only considered indicative of the mechanisms affecting execution costs.

5 Estimation Approach

We use order level regressions to examine the determinants of institutional execution costs and particularly whether the use of colocation services allows members to lower the execution costs for their trades executed in an agency and a proprietary capacity. We employ the following panel regression setup for this purpose:

$$\begin{aligned}
 \text{Execution Cost}_{istk} = & \alpha + \beta_1 \text{TradeChar}_{istk} + \beta_2 \text{Agency}_{istk} + \beta_3 \text{Colo}_{is} \\
 & + \beta_4 \text{Colo}_{is} \times \text{Agency}_{istk} + \beta_5 \text{HFT}_{istk} + \beta_6 \text{ColoHFT}_{istk} \\
 & + \beta_7 \text{Agency}_{istk} \times \text{HFT}_{istk} + \beta_8 \text{Agency}_{istk} \times \text{ColoHFT}_{istk} \\
 & + \mu_{st} + \eta_k + \nu_i + \varepsilon_{istk}
 \end{aligned}$$

where for parent order k on day t in stock s by exchange member i , *TradeChar* is a vector of the order and execution characteristics defined in Subsection 4.2, *Agency* indicates agency

orders, *Colo* indicates whether an exchange member is colocated on the stock's primary listing venue or Turquoise¹², *HFT* is a vector containing the fraction of the parent order executed against all, passive, or aggressive HFT counterparties, and *ColoHFT* is the fraction of the parent order executed against an HFT on a venue where the exchange member is colocated. Stock-day fixed effects μ_{st} capture any unobservables affecting the execution cost for a stock on the day the order starts trading. For example, trend-chasers, momentum (Jegadeesh and Titman, 1993) and contrarian traders (De Bondt and Thaler, 1985), trade stocks based on their past performance. Quarter-hour fixed effects η_k capture systematic differences in execution costs depending on when during the trading day the execution begins. This includes differences in execution costs likely arising due to distinct patterns in intraday price formation.¹³ For example, Madhavan et al. (1997) find that intraday patterns in liquidity and volatility are consistent with a decreasing level of information asymmetry and increasing inventory costs through the day. Another example is the potential impact of the increased trading activity around the US market open. Exchange-member fixed effects ν_i capture the impact of broker skill and other unobserved effects associated with the individual brokers, for example differences between high-touch versus low-touch brokers. We estimate the above model with and without exchange-member fixed effects to separately understand the differences in execution costs by trading capacity and collocation within individual brokers. In all models, standard errors are clustered by exchange member and stock. We winsorize all variables, except the dummy variables and those expressed in percentages, at the 99.5% level.

¹²As discussed above, we use collocation on Turquoise as a proxy for collocation on the MTFs.

¹³In untabulated results, we examine the distribution of orders throughout the trading session and do not find any large differences in when agency and proprietary trades begin.

A potential concern in our analysis is that our regression design does not fully control for the differences in trading motives or investment styles across agency and proprietary orders. While investors' trading motives are not directly observable in our dataset, our choice of control variables and fixed effects indirectly controls for trading motives and differences in execution strategy originating from heterogeneous trading motives. Furthermore, if trades are observationally indistinguishable in our data, they are likely also hard to distinguish by other market participants, whose response ultimately, apart from observable execution characteristics, determines execution costs. It appears unlikely that unobserved trading motives would affect execution costs.

A second possible concern in our regression design could be that the functional form of the dependence of the execution costs on the observable trade characteristics might differ from that assumed in our regression specification, which would pose a problem in the presence of an imbalance in the covariates. For example, larger orders are more difficult to execute and are generally associated with higher execution costs. However, the exact relationship, whilst apparently concave (Keim and Madhavan, 1996, 1998), is unknown. While the square root law is widely accepted in the industry (see, for example, Torre, 1997), Almgren et al. (2005) and Zarinelli et al. (2015) obtain evidence against it.

We address this concern by using inverse propensity score weighting (IPW) in our regressions. This estimator combines weighting and regression adjustment and has the doubly-robust property (Robins and Rotnitzky, 1995; Imbens and Wooldridge, 2009), meaning that it is consistent even if one of the models for the propensity score or the outcome regression is misspecified. While this approach has occasionally been used in the finance and economics literature in the case of binary treatments (see e.g. Bird, 2018; Van Biesebroeck et al., 2015),

our context entails multiple treatments: the exchange members' colocation status, and the trading capacity (agency vs. proprietary).^{14,15}

We thus follow the literature started by Imbens (2000), who extends the binary treatment propensity score methodology to multiple treatments. In our case, there are two binary treatment variables forming four treatment categories with no natural ordering: (\bar{C}, P) , (C, P) , (\bar{C}, A) , (C, A) , where C and \bar{C} denotes colocation and non-colocation and A and P denotes agency and proprietary trades, respectively. In the first stage, we estimate the propensity score for the four treatment categories using a multinomial logit model, where we include as independent variables the same vector of market, trade, and execution characteristics described above. Similar to Uysal (2015), the inverse of the predicted propensity score for each of the four respective treatment categories is then used to weigh the observations in the second stage regression where we again control for market, trade, and execution characteristics as well as stock-day, quarter-hour, and, in some of our analyses, exchange member fixed effects. Since this is a two-step estimation approach with no known analytical expression for the standard errors, we estimate these using bootstrapping with 1,000 iterations.

¹⁴Another commonly used doubly-robust approach employing propensity scores uses matching. While matching is possible even in the multiple treatment case (Linden et al., 2016), finding good matches becomes increasingly difficult in finite samples.

¹⁵We report results from unweighted panel regressions in the internet appendix A. The results do not qualitatively differ from those obtained using the doubly-robust estimator.

6 Descriptive Statistics

6.1 Stock Characteristics

[Table 1](#) shows the summary statistics for the 83 out of 86 stocks in our final sample which have parent orders. We also report the summary statistics for each market capitalization tercile. The stocks have an average market capitalization of €8.75 billion, daily continuous trading volume of €29 million, and a quoted bid-ask spread of 20 bps. Trading activity is highly fragmented across the four main venues with the primary listing exchange responsible for 62% of the market share and the three largest MTFs contributing the remaining 38% for the average stock. However, there is substantial cross-sectional variation across all these dimensions as evidenced by the difference in mean (and median) values across the three size groups. Stocks of larger firms have higher trading volumes, lower bid-ask spreads, and their trading activity is more fragmented. Finally, according to Refinitiv data, the four lit venues in our data jointly comprise more than 95% of the total lit trading.

Insert [Table 1](#) about here

6.2 Exchange Member Characteristics

We identify 11,724 parent orders executed by 139 exchange members in 83 stocks from 8 countries. [Table 2](#) provides exchange-member level summary statistics separately for all exchange members in our sample (Panel A) and for those who trade both in a proprietary and agency capacity (Panel B).

Insert [Table 2](#) about here

The average exchange member executes 84 parent orders in about 14 stocks from 3 countries totaling €48.1 million during our sample period. There is substantial heterogeneity across exchange-members: some are active in only three (one) stocks (country), whereas others are active in 54 (eight) stocks (countries). 36 exchange-members engage in both agency and proprietary trading and 60 (43) members specialize in agency (proprietary) trading. However, the dual capacity members are much more active. The average exchange member trading in a dual capacity trades more than 29 stocks from 5 countries totaling €138.3 million during our sample period. This translates to such exchange-members accounting for more than 70% of all the parent orders both in terms of order count and total trading volume.

35% of all exchange members and two-thirds of those acting in a dual capacity execute some fraction of their institutional orders on the three MTFs. This may seem surprising considering the best execution obligations imposed on brokers by EU regulations, in particular by the Markets in Financial Instruments Directive (MiFID).¹⁶ However, broker-dealers enjoy substantial flexibility while ensuring compliance with these rules. For example, MiFID allows them to define best execution in terms of price, costs, speed, likelihood of execution and settlement, size, nature or any other consideration relevant to order execution. At the same time, brokers can apply these factors differently to different clients, orders, instruments, and venues. It is also not mandatory that brokers are connected to multiple trading venues.¹⁷

The mean and median number of primary markets in which the average exchange member is colocated is 0.6 and 0, respectively, indicating that the majority of members in our sample

¹⁶Contrary to the US, where the obligation to provide best execution rests largely with trading venues, in the EU, brokers are obligated to provide best execution to their clients.

¹⁷See [CESR \(2007\)](#) for further details.

choose not to collocate. In fact, only 30% of the exchange members use any collocation services. The vast majority of these members choose to collocate on one of the primary listing exchanges as opposed to on the MTFs. Conditional on being collocated anywhere, exchange-members are on average collocated on 2.2 venues. Exchange-members trading in a dual capacity are more likely to collocate on the primary markets and the MTFs. They are also collocated on more trading venues compared to exchange-members trading exclusively in an agency or proprietary capacity. Specifically, 56% and 25% of all exchange-members trading in dual capacity are collocated on the primary markets and MTFs, respectively. Unconditionally, such members on average collocate on 1.3 primary markets and, conditional on being collocated anywhere, they are collocated on 2.8 venues.

6.3 Parent Order Characteristics

[Table 3](#) provides summary statistics at the parent order level separately for the complete set of orders and for those executed in an agency and proprietary capacity. The table also reports the mean differences between agency and proprietary orders, the associated t-statistics, and the standardized mean difference as a measure of the economic magnitude of the differences.

Insert [Table 3](#) about here

Panel A shows that, for approximately 53% of all parent orders, the executing broker is collocated on the stock's primary exchange or on the MTFs. The difference between these figures for agency and proprietary orders (51% versus 55%), while statistically significant, is economically small.

Panel B shows that there are roughly equal numbers of buy and sell orders. The average

trade size is 570 thousand Euro and *Information*, the average long-term price impact of trades, amounts to -9 bps suggesting that, on average, a parent order is uninformed. These numbers do not substantially differ between trades executed in the different capacities as indicated by the insignificant differences. The execution lasts on average about 5 exchange trading hours and, while the difference between agency and proprietary trades is statistically significant, its magnitude is rather small.

Panel C describes the parent order execution characteristics. There is a substantial difference between the two capacities in their choice to supply or consume liquidity: 44% of an average agency trade consumes liquidity, whereas this percentage is only 39% of proprietary trades. The difference concerning the choice of trading venue is even larger: only 17% of the average agency order executes on the MTFs, whereas the corresponding percentage for proprietary orders is 39%.

Panel D reports the execution costs of the average parent order, the effective half spread paid (earned) by the aggressive (passive) executions that are part of an average parent order, and finally the one-minute price impact associated with every aggressively and passively executed child order, respectively. The execution cost, our main variable of interest, is -4.46 bps on average. The negative average may result from the fact that the average parent order in our sample is uninformed, as evidenced by the negative value of the *Information* variable, and traded predominantly passively, as evidenced by the *Aggressiveness* variable. For instance, [Keim and Madhavan \(1998\)](#) find that value (technical) traders who trade passively (aggressively) end up with negative (positive) execution costs. Furthermore, based on the regression coefficients estimated later in this paper, the expected execution cost for an agency order with zero information and an even split between liquidity-consuming and

-providing executions is positive. However, average execution costs are significantly smaller for proprietary orders than for agency orders. The difference of more than 4 bps, while not large relative to the variation in execution cost between individual trades, is economically large and statistically significant. The average aggressive (passive) trade incurs (earns) a half-spread of 4.97 bps (3.62 bps). The average trade when executed aggressively incurs a lower cost for proprietary trades vis-à-vis agency trades (4.75 bps vs 5.20 bps). At the same time, the average passively executed proprietary trade earns a larger half spread than the average passively executed agency trade (3.74 bps vs 3.49 bps). The better child order execution outcomes measured in terms of effective spreads also extended to the average price impact. The average price impact of an aggressively executed order is higher for agency trades as compared to proprietary trades (3.75 bps vs 3.25 bps). The corresponding numbers for passive executions are nearly identical (-4.68 bps vs -4.67 bps). The difference in effective spread and price impact between agency and proprietary trades provides the first evidence of the execution performance of underlying algorithms. In [Section 8](#), we test whether differences in these (and other) variables hold using the empirical model described in [Section 5](#).

Finally, Panel E reports the HFT interaction statistics. Agency orders trade about 30% of their volume against HFT firms, whereas the number for proprietary orders is about 33%. This difference is significant and consistent with proprietary orders being executed to a larger extent on MTFs, where HFTs are generally more active than on the primary listing exchanges. The average parent order is approximately twice as likely to interact with aggressive HFTs as compared to passive HFTs. However, agency orders execute to a slightly lesser extent against aggressive as opposed to passive HFTs (65% vs 68%). These statistics

are broadly consistent with the usage of aggressive orders in Panel C. For instance, roughly two-thirds of the average parent order is executed passively and roughly two-thirds of the average parent order's HFT interactions are against aggressive HFT. Similarly, the average agency order is more aggressively executed compared to the average proprietary order and is also more likely to interact with passive HFT counterparties. These statistics suggest that the algorithmic decision to supply or consume liquidity does not substantially influence the extent to which the parent order interacts with HFT firms.

To conclude, agency and proprietary executions are largely similar in terms of their order characteristics, duration, and the use of exchanges' colocation services. However, they differ with respect to their choice of trading venue, level of aggressiveness, and interactions with HFT firms. We overcome this imbalance via the doubly robust approach of controlling for all these variables in our analyses in combination with inverse propensity weighting. [Table 4](#) in the internet appendix demonstrates the success at eliminating meaningful covariate imbalances across all four combinations of capacity and colocation.

6.4 IPW balancing performance

[Table 4](#) provides evidence for the covariate imbalance in the original data and the success IPW achieves in reducing this imbalance. It shows means of the predictors for each of the four combinations of capacity and colocation, as well as comparisons between each of the six pairs of combinations. Whereas the significance tests are highly dependent on sample size, the degree of potentially economically important difference can be inferred from the standardized differences ($\frac{\Delta}{\sigma}$). We observe in Panel A that there are economically meaningful

differences for all pairs, and that these differences are not limited to variables that can be assumed to reflect differences between execution algorithms (last three rows). IPW succeeds at reducing the differences to magnitudes that mostly appear negligible. Among 60 comparisons, only five differences exceed 0.1 standard deviations. Three of these differences result from aggregate trading volume in the stock during execution (*log* Market Trading) in the weighted sample being somewhat smaller for colocated proprietary trades, whereas the other two reflect differences in the degree of order splitting and the use of MTFs. Note that covariate imbalances do not have to be problematic *per se* because we employ regression adjustments in all of our analysis; yet, the reduction in imbalances that IPW achieves limits the sensitivity of results to the exact specification of the regression model.

Insert Table 4 about here

7 Parent Order Results

This section analyzes the determinants of execution costs of large orders. Our main focus is on understanding how execution costs differ by the capacity in which the exchange member trades an order and the use of colocation. Additionally, we consider the relation between execution costs and interactions with HFTs.

7.1 Baseline Analysis

Execution costs are positively related to trade difficulty (Keim and Madhavan, 1997). For example, orders that are large (relative to the daily volume of a stock) are more difficult to execute. A market participant acquiring a large position may prefer to trade patiently

over a long period while relying on limit orders (to avoid crossing the bid-ask spread) and minimizing executions on MTFs to avoid being back-run by HFTs. However, if a market participant is impatient, possibly because it possesses short-term information, it may have to trade aggressively using marketable orders while accessing all sources of liquidity. We examine this tradeoff by regressing execution costs on the predictors described in [Subsection 4.2](#). [Table 5](#) contains the baseline results.

Insert [Table 5](#) about here

As expected, transaction costs increase with parent order size. Depending on the regression specification, a 0.1 unit increase in log trade size is associated with between 0.26 and 0.33 bps higher execution costs. As trade sizes vary substantially and can be very large in some cases, this effect is economically significant. The total trading volume in the market, *Market Trading*, is negatively related to execution costs, consistent with the notion that a higher trading volume introduces more noise in the signal generated by the parent order executions. A 0.1 unit increase in log total volume during the parent order's lifetime is associated with 0.34 bps lower execution costs.

The execution costs for buys are around 2 bps higher than for sells, though the effect is statistically insignificant.

Pre-trade volatility, measured as the standard deviation of ten-millisecond quote midpoint returns over the ten minutes before the first child order execution, predicts a 17 bps difference in parent order execution costs for every 1 bps change in volatility, however the impact is statistically insignificant. Note that the mean and standard deviation of volatility are both less than one-tenth of one basis point. The stock return over the ten minutes before the

trade has no explanatory power across all specifications.

The permanent information content of the trade is, as expected, positively associated with the execution cost. A 10 bps increase in *Information* is associated with approximately 1.5 bps higher execution costs. The time it takes to complete a parent order is negatively associated with transaction costs and the association is statistically significant. Specifically, working a parent order for one additional hour is associated with a 1.7 bps lower execution cost. This is because, keeping parent order size constant, duration is inversely related to the institution's participation rate and, consequently, negatively associated with transaction costs. However, this effect becomes somewhat smaller and statistically weaker once we control for the stock's total trading volume (*Market Trading*) during the parent order's life time, which together with the child order size determines the participation rate. Trading impatiently by using aggressive (i.e., marketable) orders is associated with higher transaction costs. A trade that is executed with a 10 percentage points higher share of liquidity taking orders is associated with a 3.1 bps higher execution cost.

The number of child order executions is also significant in explaining parent order transaction costs: a 0.1 unit increase in log child order executions is associated with a 0.41 bps increase in execution costs, though this effect decreases slightly once we control for the fraction of a parent order traded on MTFs, where individual order executions are generally smaller. Considering that the size of individual child order executions, and therefore their number, is endogenously related to the liquidity available in the limit order book, even in the presence of stock-day fixed effects, the direction of any causality between the number of child orders and execution costs is unclear. The situation is similar when interpreting the effect of the volume executed on MTFs because the choice of venue partially depends on the

liquidity available in the primary market. The coefficients suggest that shifting 10% of the parent order away from the primary market and towards the MTFs is associated with a 0.7 bps increase in execution costs.

7.2 Trading Capacity, Colocation, and Interaction with HFT

In [Table 6](#), we consider the effects of exchange members' execution capacity, whether the exchange member executing the parent order is colocated on a venue where the stock trades, and the fraction of a parent order that executes against liquidity providing and consuming HFTs. We report the results from panel regressions with stock-day and quarter-hour fixed effects, and the variables included in column (4) of [Table 5](#) as untabulated controls. We center all interacted continuous variables at zero.

Insert Table 6 about here

We find that, after including all baseline control variables, orders executed in an agency capacity are substantially more expensive than proprietary orders. The difference in execution costs is economically meaningful at about 4.5 bps, or more than 50% of the quoted spread of a typical large cap stock in our sample. Column (2) shows that additionally including a colocation dummy does not unconditionally explain any variation in execution costs. In columns (3) and (4) we add an interaction term that identifies the effects of colocation on the execution costs of agency versus proprietary orders. The main effect of colocation, now representing the effect for proprietary orders, turns significant and shows an execution cost reduction of around 6 bps, whereas colocation for agency trades is associated with 7.2 bps higher execution costs compared to proprietary orders, suggesting that agency orders do not

benefit from exchange members' subscription to colocation facilities.

We next include trading volume against HFTs and find that an additional 10% of an order traded against HFTs is associated with an increase in execution costs by 0.9 bps. When splitting HFT into liquidity taking versus liquidity providing executions, we find that they have opposite effects on transaction costs: trading an additional 10% of an order against aggressive HFT is associated with around 3.7 bps higher execution costs, whereas the effects of trading the same amount against passive HFT is around -3.0 bps.

We then include the fraction of trading volume against HFT on a venue where the exchange member is colocated and find that both the adverse effects of aggressive, and the beneficial effects of passive, HFT are reduced when the exchange member is colocated. The adverse effect of aggressive HFT is reduced by around 50% whereas the beneficial effect of passive HFT is smaller by more than 80%. This indicates that execution costs depend on the type of counterparty to a lesser extent for colocated market participants, allowing them to obtain execution costs that are less sensitive to the nature of the counterparty. The net impact of HFT on orders originating from colocated exchange members nevertheless remains positive due to the higher magnitude of the effect of aggressive HFT on execution costs and due to parent orders in our sample being twice as likely to trade against aggressive HFTs than passive HFTs.

Finally, we interact the fraction of trading against HFT on a venue where the exchange member is colocated with the agency dummy to examine whether agency orders not only generally fail to benefit from colocation but possibly also see a smaller effect of colocation on the costs when facing HFTs. However, the coefficients suggest no significant effects when trading against HFTs.

7.3 Within Exchange Member Variation in Execution Costs

Despite all the covariates and fixed effects we employed in the previous analyses, it might be that unobservable differences across exchange members drive our results. This includes differences in the type of colocation facilities, such as access speed and capacity, chosen by individual members as well as other technological and connectivity differences. For example, as of June 2022, Deutsche Börse allows exchange members to configure their colocated infrastructure based on power requirements, type of market data, and other connectivity options.¹⁸ Finally, differences in trading skills, specifically the quality and choice of order splitting and routing algorithms, might also explain part of our results.

To address these concerns, we conduct an alternative analysis in Table 7 that, in addition to stock-day and quarter-hour fixed effects, includes exchange member-fixed effects.¹⁹ While the inclusion of the additional fixed effects eliminates the variation in colocation on Turquoise, we retain the colocation variable that captures exchange members' decision to collocate on individual primary markets, which we denote $Colo_{PM}$. This is because brokers are not colocated on all primary venues. The identification of the colocation effect now stems only from those exchange members that execute parent orders on primary markets where they are colocated and others where they are not colocated. Similarly, the identification of the differences between execution costs for agency and proprietary orders now comes from exchange members executing orders in both capacities. Employing exchange-member fixed effects in combination with our control variables and IPW effectively allows us to estimate

¹⁸Further details on Deutsche Boerse's colocation service are available at <https://www.xetra.com/xetra-en/technology/co-location-services>.

¹⁹We do not report the baseline results with exchange-member fixed effects as these are qualitatively similar to the results in Table 5.

the impact of colocation for agency versus proprietary orders for similar orders originating at the same exchange-member.²⁰

Insert Table 7 about here

Similar to the previous analyses, parent orders executed in an agency capacity incur higher execution costs. Compared to the earlier results, the size of the coefficient increases to 9.2 bps. This finding implies that, while agency executions across all member firms incur higher execution costs, execution costs of those firms that specialize in agency or proprietary executions differ less compared to the differences between agency and proprietary orders originating from exchange members that trade in a dual capacity.

After including the colocation dummy and its interaction with the capacity dummy, the results are again similar to those obtained without using exchange member-fixed effects: proprietary trades benefit from the use of colocation, whereas this is not the case for agency trades. Specifically, the effect of colocation on the execution costs of proprietary orders is -6.3 bps. However, when these colocated brokers execute agency orders, they have significantly higher execution costs with a coefficient of 7.8 bps. The effect of the agency dummy variable is statistically significant at 5 bps, which suggests that agency orders originating from dual-capacity exchange members have higher execution costs compared to proprietary orders executed by the same firm even when the broker is not colocated.

The unconditional effects of parent orders interacting with HFT firms remain qualitatively unchanged. In aggregate, trading against HFTs is associated with higher execution costs. The volume transacted against aggressive HFT orders is associated with higher execution

²⁰The reduction in effective sample size necessarily reduces the estimations' statistical power.

costs whereas the volume against passive HFT orders has the opposite effect. Compared to the results without exchange-member fixed effects, the coefficients for the interaction of colocation with aggressive and passive HFTs retain their signs whereas the size and statistical significance changes somewhat. The benefit enjoyed by a non-located broker from trading against passive HFTs is somewhat larger, whereas the interaction effect between colocation and passive HFT is larger and now significant. The two effects cancel out suggesting that sophisticated traders do not systematically experience large differences in execution costs based on the kind of counterparties they trade with.

In conclusion, the main results do not appear to be sensitive to the chosen regression specifications. Colocation helps proprietary traders reduce their transaction costs and avoid losses to aggressive HFTs, but these benefits are limited to proprietary executions.

7.4 Analysis of Individual Member Fixed Effects

While our previous results document that, on average, the benefits of colocation in the form of lower execution costs are limited to proprietary executions, they leave unanswered the question whether this phenomenon is common or driven by a small number of outliers among the brokers. To answer this question, we perform regressions similar to the above, except that we include a larger set of dummy variables to study the cross-sectional distribution of within-broker differences in execution costs between agency and proprietary trades. In particular, in addition to the baseline control variables and stock-day and quarter-hour fixed effects as in the regressions in [Table 5](#), we include dummy variables at the broker-capacity-colocation level. There are 18 brokers trading in both capacities on venues where they are

colocated. For each of these brokers, we compute the difference between the fixed effects of agency and proprietary trades, respectively, for trades executed on a venue where the broker is colocated.

Insert Figure 1 about here

Figure 1 shows the estimates and error bars of this difference for the individual brokers. Execution costs are higher for agency orders for 15 of the 18 brokers and for 5 of them the difference is statistically significant at the 10% level. At the same time, none of the 3 differences in the opposite direction are statistically significant. A non-parametric Wilcoxon signed-rank tests provides strong evidence that agency orders do not obtain the same benefit from collocation as do proprietary orders, with a p-value of 0.0021. In summary, the higher execution costs of agency compared to proprietary trades executed by the same colocated broker can be observed for most dual capacity brokers in our sample. Overall, these results show that our main findings are not driven by a small subset of these brokers but the result of likely similar choices made by brokers and their institutional customers.

8 Potential channels behind the key results

The results reported in Section 7 raise the question as to *how* and *why* agency trades, in particular those executed via colocated brokers, obtain execution costs that are higher than comparable trades executed by exchange members in a proprietary capacity. We here do not aim to establish causal effects of any particular characteristics of the order submission strategies on parent order execution costs. This would require reengineering the logic of

the order execution algorithms, a task that appears infeasible, considering the fact that we observe data on a broker level and each broker likely employs multiple algorithms.²¹ Instead, we first consider the algorithms' ability to monitor the order books. Next, we analyze the execution quality of the individual child orders and how these measures relate to the brokers' colocation status and the capacity in which they execute orders. Then, we consider the possibility that brokers route their customer orders with the objective of minimizing exchange fees. Finally, we ask whether differences in execution cost risk might rationalize the differences in average execution costs observed.

8.1 Order Book Monitoring

Subscribing to exchanges' colocation facilities should allow brokers to better monitor the order book at high frequencies and employ order execution algorithms that quickly respond to changing market conditions. For instance, it can allow brokers to rapidly revise a passive child order after submission if the order becomes stale due to price changes, thereby resulting in improved fill ratios and/or reduced price impact costs. In this subsection, we provide evidence indicative of the improved monitoring ability by analyzing order flow periodicity and order-to-trade ratios (OTR) for orders submitted in both capacities and with or without the use of colocation technology.

[Bacidore \(2020\)](#) argues that *“one tremendous avoidable source of latency stems from algo provider’s decision to use ‘timer-based’ instead of ‘event-based’ processing.[...] A timer-based algorithm ‘wakes up’ periodically, e.g., every 500 milliseconds, samples the state of*

²¹For a detailed analysis of the execution algorithms offered by a single broker, see [Beason and Wahal \(2020\)](#).

the market, and then takes action before beginning another 500-millisecond slumber[...].”

The relatively higher latency in the absence of collocation may be a reason for brokers to employ timer-based strategies. [Figure 2](#) shows the the frequencies of orders submitted during different seconds within a minute separately for the four combinations of collocation/no collocation and agency/proprietary orders. The plots show deviations of the logarithm of the number of order submissions from their mean. For all four groups, the largest abnormal number of order submissions occurs in the first second of a minute, and the effect is larger for agency compared to proprietary orders, and for non-located compared to located orders. While we cannot rule out the possibility that (some of) the excess order volume may be in response to other market participants’ high activity at those times, the plot does suggest the existence of algorithms using clock-time scheduling. There are also spikes at the other 10-second intervals within the minute for both agency and proprietary located orders. These spikes are more pronounced for agency orders, which is indicative of agency algorithms being less responsive to market events. The lower periodicity observed for non-located orders is likely due to the presence of high-touch trading in this category, which introduces noise in the observed periodicity. In untabulated tests, we find that the observed differences at second 0 hold after employing stock-fixed effects or stock- and broker-fixed effects, whereas the coefficients for those at seconds 10, 20, 30, 40, and 50, are of the expected sign though not statistically significant.

Insert [Figure 2](#) about here

Altogether, the figure and the statistical tests imply that there is more periodicity in the order submission of agency versus proprietary located orders. The presence of such

predictable patterns in order submissions in calendar time, which has also been observed in earlier studies (Brugler, 2015; Sağlam, 2020), might contribute to poor execution performance, as such patterns can be anticipated by HFTs engaging in back-running behavior. For instance, Sağlam (2020) finds that the use of time- and volume-based execution algorithms leads to predictable patterns in order flow, which are correlated with higher execution costs. Conversely, a reduction in predictability is associated with lower execution costs. While the above analysis reveals the presence of such predictable patterns for parent orders executed with and without collocation in both capacities, the disproportionately higher predictability of colocated agency orders may help explain the higher execution costs for such orders.

Next, we analyze the differences in the OTR – defined as the ratio of the number of order submissions to the number of child order executions – by capacity and use of collocation. Specifically, we consider whether the OTR systematically differs by the type of parent orders by regressing parent orders’ OTR on the same predictors as in the previous analyses. OTRs are generally used to track algorithmic and high-frequency trading activity in electronic markets (see, for example, Hendershott et al., 2011). The results in Table 8 show that there is no statistically significant difference between agency and proprietary orders in the absence of collocation.²² The results without broker-fixed effects in column (1) show that proprietary colocated orders feature a significantly higher OTR than non-colocated ones, which is consistent with collocation enabling brokers to more successfully manage their orders’ positions in the limit order book. However, this effect is absent for colocated agency orders, with a coefficient that, albeit only weakly statistically significant, nearly offsets the

²²The table is based on observations for 80 of the 83 sample stocks. Due to only intermittent access to the data, we will update the table with results for the complete sample in the next revision.

positive coefficient for collocation. The within-broker differences in OTR between colocated and non-colocated proprietary orders reported in column (2) are positive, though smaller than in column (1) and statistically insignificant. This reduced effect is consistent with the idea that brokers that colocate anywhere employ relatively sophisticated algorithms across trading venues regardless of their collocation. This would allow them to actively manage their orders to some degree even where they are not colocated. However, OTR is significantly lower for colocated agency orders, which, somewhat surprisingly, results in colocated agency orders having a lower OTR than non-colocated agency orders. However, the latter difference (untabulated) is not statistically significant.

Insert Table 8 about here

Altogether, this subsection provides evidence suggesting that the algorithms colocated brokers use for their proprietary orders materially differ from those used for customer execution. The latter feature higher periodicity and, in particular, a lower OTR, which is indicative of a less active management strategy of child orders, and likely contributes to higher parent order execution costs.

8.2 Child Order Execution Quality

Having established some differences between colocated agency and proprietary orders with respect to their child order submission behavior, in this subsection, we turn to their execution quality. When executing a child order, an institutional investor would like to accomplish two goals: first, low execution costs of individual child orders, and second, to minimize the extent of adverse price changes that would make subsequent child order executions more expensive.

As a measure of how well an order performs with respect to the first goal, we compute parent order level averages of the effective spreads for the individual child executions that make up a given parent order. We disaggregate the analysis into spread costs for aggressive versus passive order execution.

Panel A of [Table 9](#) contains the results. Note that the number of observations in this table is smaller than in the previous ones because a parent order is included only if it contains at least one aggressive (passive) child order execution. Column (1) shows that colocation reduces the spread paid by liquidity consuming orders, and this effect does not significantly differ between proprietary and agency orders. However, agency orders pay wider spreads regardless of the use of colocation, suggesting that execution algorithms used for agency trades may generally be inferior to those used for proprietary trades. This cost disadvantage of 0.4 bps amounts to more than half of the benefit colocation provides. However, it makes up only a small part of the execution cost disadvantage agency parent orders suffer. When we consider within-broker variation in column (2), we do not find statistically significant differences between trades in stocks listed on an exchange where a broker is colocated and where it is not, and, though the coefficient estimate is of the expected sign, it is of substantially smaller magnitude than in the analysis without broker fixed effects. This finding is consistent with that in the previous subsection that brokers' execution capabilities may to a significant extent translate across exchanges, and that the use of colocation anywhere goes hand-in-hand with higher technological skills. The interaction effect of colocation with agency orders is also statistically insignificant, though the magnitude and sign suggest that any advantage of using colocation almost exactly cancels out for agency orders. Agency orders pay wider spreads regardless of colocation.

Insert Table 9 about here

Turning to effective spreads earned in passive order executions, the negative coefficient on colocation in column (3) indicates that colocation is associated with wider spreads earned by the liquidity provider. The effect size is smaller than for aggressive orders, suggesting that speed and/or sophistication is more important for liquidity taking orders. We do not observe any significant differences between proprietary and agency orders. While the results obtained when we consider within-broker variation in column (4) are not statistically significant, agency orders generally obtain inferior execution, which, in this case, means they earn smaller spreads. This result suggests that dual capacity brokers manage agency limit orders worse than comparable proprietary limit orders, whereas this finding is not observable in the results exploiting differences across brokers. Compared to the analysis of aggressive orders, we observe that, in relative terms, colocation remains an economically more important predictor of effective spreads. While the effect size is somewhat smaller than in column (3), it does appear that speed, in addition to general technological sophistication, helps brokers earn wider spreads through facilitating timely responses to changing market conditions. The interaction effect of colocation and agency is statistically insignificant, though the point estimate amounts to a reduced benefit of employing colocation by more than half.

We now turn to the second goal institutional investors have with respect to minimizing their execution costs, i.e., avoiding price impact as a result of information leakage. Whereas in classic microstructure models the price impact is a reflection of the information content of a trade, investors splitting their trade among multiple child order executions desire to minimize their price impact so as to obtain beneficial trade prices in subsequent child trades.

Thus, for aggressive orders, a low price impact points to a good execution outcome. For passive orders, a price movement against the direction of the order is usually interpreted as adverse selection and therefore undesirable. However, assuming that a trade happened at a certain price, an institutional trader would prefer experiencing a large price impact to obtain better prices for subsequent child orders. Yet, a limit order trader who does not manage the order effectively may get an execution when the order is stale, i.e., the price at which the child order trades may be disadvantageous as it would have been possible to trade at a better price. In that situation, while the price impact incurred may be large, the prices of subsequent child trades may not necessarily improve because the price would have moved anyway. Thus, for passive trade executions, it is unclear whether a large price impact should be interpreted as beneficial or harmful.

We measure price impact as the quote midpoint change from immediately before to either 1 or 60 seconds after a trade, signed by the direction of the institutional order. Similar to the effective spreads, we report the results separately for aggressive and passive orders.

Panels B and C of [Table 9](#) show the results for price impacts at a one-second and one-minute horizon, respectively. For aggressive orders, the results show that colocation enables significantly smaller price impacts, with an effect size of about 0.7 bps within one second, which grows to 1.0 bps within one minute. Non-located agency orders are associated with statistically insignificant immediate and near-zero one-minute additional price impacts. However, the benefit to using colocation vanishes for agency orders. While the effect is slightly smaller and not statistically significant within one second, it grows to almost 1 bp after one minute, offsetting the advantage implied by the coefficient for colocation. The within-broker analysis yields similar results to those for effective spreads: the effects of colo-

cation and its interaction with agency are small, whereas agency orders in general, regardless of collocation, result in price impacts that are higher by about 0.8 bps, within one second, though the coefficient shrinks to about 0.4 bps and becomes statistically insignificant for the one-minute horizon.

Columns (3) and (4) report the price impact results for passive orders. We note that a negative coefficient indicates a larger price impact from the institution's perspective or, put differently, a larger price move against the institution's trade direction. Column (3) shows that passive agency orders incur a larger price impact, with an effect size that rises to about 0.7 bps within a minute. As explained above, this could be an indication of trade execution at stale prices, and thus reflect poor execution quality. Alternatively, it suggests lower information leakage, potentially enabling successive trades at better prices. While we consider the former channel more likely as the results in [Subsection 8.1](#) indicate that it is unlikely that agency orders are better monitored than proprietary ones, our estimations do not allow us to discriminate between the two interpretations. Proprietary colocated orders enjoy a price impact that is larger by about 0.8 bps after one second and 0.9 bps after one minute, whereas the positive coefficients for the interaction effect of collocation and agency show that the collocation effect nearly cancels out for agency orders. Based on the findings of the previous subsection and for spreads, it is likely that the effect of collocation for proprietary orders is reflective of the avoidance of information leakage. By contrast, agency orders do not enjoy the same benefit from collocation, and appear to be more exposed to stale executions, though the coefficient estimates alone do not firmly lead to a unique interpretation.

The within-broker analysis in column (4) shows that here, significant differences between non-colocated proprietary and agency orders do not exist. Collocation by itself does not

statistically significantly predict price impact, though the effect size at a one-minute horizon is about half that observed in column (3). However, for colocated agency orders, the price impact is smaller, from the parent order's perspective, by 0.8 bps after one second and 0.5 bps after one minute. While the latter coefficient is not statistically significant, it is still sizable, when considering that this effect size accumulates over multiple child orders. To sum up, the effects of colocation are smaller within than across brokers, consistent with the previous results in this subsection.

8.3 Exchange fees

Prior literature has shown that brokers may prefer routing customers' orders in a way that minimizes fees (Battalio et al., 2016; Anand et al., 2021) at the expense of the orders' execution quality. While, as is the case for the other order routing characteristics considered in this section, we cannot test whether such behavior is the cause of our main results, we can test whether there are differences in the fees dependent on the principal versus agency capacity and the existence of colocation. To do so we conduct analyses similar to those for the execution costs, with the exchange fees as the dependent variable instead. However, here we omit in our main specification the predictors that are already a function of the algorithm used rather than of a parent order's objective: the aggressiveness, the number of child orders, and the use of MTFs. This is because, if brokers attempt to minimize exchange fees, much of this minimization could be captured by these variables, whereas we are interested in whether brokers minimize exchange fees for a given trading interest. For exchanges offering fee schedules dependent on the members' total trading volume, we assume the fee/rebate

corresponding to the highest volume threshold. The baseline results, not including capacity and colocation, are reported in the internet appendix [Table C1](#). [Table 10](#) provides the results for our variables of interest, columns (1) to (3) and (4) to (6) showing the results without and with broker fixed effects, respectively.

We find that, for the specification without broker-fixed effects, exchange fees for agency orders in aggregate do not significantly differ from those for comparable principal orders. However, when conditioning the results on colocation, we observe that non-located agency orders incur higher exchange fees (by about 0.03 bps) than comparable principal orders. Fees for principal located orders are higher than those for non-located ones, though when customer orders are executed by a located broker, fees are lower than those for comparably principal orders. Thus, among non-located orders, brokers pay higher fees for agency than for principal orders, whereas the converse hold for located orders. It appears that brokers achieve lower exchange fees precisely for that type of parent order that we observe to obtain the poorest execution quality relative to similar proprietary orders. When including exchange-member fixed effects, we again find that agency orders incur higher exchange fees. However, located principal orders do not significantly differ in their trading fees from non-located ones. The interaction effect of colocation and agency orders roughly cancels out the agency main effect. Thus, within broker, located agency orders again incur lower expenses than non-located ones, and they are on a par with non-located principal orders.

Insert [Table 10](#) about here

To sum up, we do not find evidence that brokers attempt to generally minimize exchange

fees when trading for customers as compared to on their own account. By contrast, non-colocated agency orders incur the highest exchange fees.

8.4 Execution cost risk

If institutional investors are risk-averse with respect to execution costs and proprietary traders possibly are not, the differences in average execution costs that we observe could result from the optimization of different targets. To test for this possibility, we analyze whether the magnitude of the residuals from the execution cost regressions differs dependent on capacity and colocation. For reasons analogous to those outlined in the analysis of exchange fees above, our focus here is on the results that exclude observable characteristics of the execution algorithm (aggressiveness, number of child orders, use of MTFs). The regression models we employ are, with the exception of the dependent variables, identical to those used when explaining execution costs. We present, in [Table 11](#), results for two different measures of the variation in unexplained execution costs: the absolute values of the residuals (models (1) and (2)), and the squared residuals (models (3) and (4)). The results across exchange members (models (1) and (3)) show significant effects neither for agency trades nor for the interaction of agency trades with colocation, thus providing no evidence for seeking to reduce variation in execution costs of agency orders. Colocation is associated with a larger variation in outcomes (by about 3 bps in absolute value), suggesting that algorithms used by colocated brokers pay less regard to the variation in outcomes. The within-exchange member analyses (models (2) and (4)) show significantly larger variation in outcomes for non-colocated agency orders, rejecting the notion that the variation in costs of agency orders

is minimized. The coefficients for colocation and its interaction with agency trading are not statistically significant, and their size implies that, holding the colocation status constant, costs of agency orders vary to a larger extent than those of proprietary orders. Thus, there is no evidence for the idea that an attempt at minimizing variation of trading cost could explain our main results.

8.5 Summary of potential channels

In conclusion, the results on effective spreads and price impact contribute to our understanding of the results on parent order execution costs. Across brokers, in aggregate we observe poorer execution quality for agency orders and, in particular, colocated agency orders. Most of the results of the within-broker analyses are quantitatively and statistically weaker, but they tend to point toward poorer execution quality for agency orders irrespective of colocation. The OTR results, indicating a less active management of child orders, complement the evidence described above and help explain the incremental parent order execution costs not captured by measures of immediate child order execution quality.

We find no evidence for the notion that our main results are driven by brokers' attempts at minimizing trading fees when trading for their customers, nor for the idea that the results stem from minimizing agency trades' variation in unexpected execution costs.

9 Conclusion

Speed dispersion across market participants is of importance because recent empirical and theoretical research has shown that HFTs can make large institutional trades more expensive,

which, in turn, can lead to reduced price discovery and allocative efficiency of capital in the economy. Subscribing to exchanges' colocation facilities allows exchange members to submit orders and access the exchanges' market data feeds at low latencies. Based on a pan-European highly granular dataset comprising the main equity markets, we study how the execution costs of exchange members attempting to acquire or dispose of a large position in a proprietary or agency capacity vary with their use of the exchanges' colocation facilities.

Exchange members who colocate obtain lower execution costs for their proprietary orders but not for the orders of their buy-side customers. These results hold not only across brokers but also for the subset of brokers who execute large orders in a dual capacity. Furthermore, while institutional investors' execution costs are, on average, sensitive to trading against HFT firms, the magnitude of this effect is significantly smaller for colocated exchange members trading in a proprietary capacity. Investigations into the underlying trading patterns reveal that the above differences between customer and proprietary orders executed via colocated brokers coincide with differences in the monitoring and the execution quality of individual child orders. By contrast, there is no evidence that they result from attempts at minimizing exchange trading fees or the variation in execution costs to a larger extent for customer than for proprietary orders.

We conjecture that the lack of benefit from colocation for customer orders is due to customers making inferior choices when selecting the most appropriate algorithm for a particular order or due to brokers using different algorithms for their proprietary and customer orders. While regulations do not require benchmarking the execution costs of orders originating from buy-side customers to the brokers' own executions, this comparison could be considered a natural approach to evaluating best execution.

References

- Ait-Sahalia, Y. and M. Saglam (2017). High Frequency Market Making: Optimal Quoting. *Working Paper*.
- Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Direct Estimation of Equity Market Impact. *Working Paper*.
- Anand, A., P. Irvine, A. Puckett, and K. Venkataraman (2011). Performance of Institutional Trading Desks: An Analysis of Persistence in Trading Costs. *Review of Financial Studies* 25(2), 557–598.
- Anand, A., M. Samadi, J. Sokobin, and K. Venkataraman (2021). Institutional Order Handling and Broker-Affiliated Trading Venues. *The Review of Financial Studies* 34(7), 3364–3402.
- Avramovic, A., V. Lin, and M. Krishnan (2017). We’re All High Frequency Traders Now. *Credit Suisse Trading Strategy*.
- Bacidore, J. (2020). Does Your Algorithm Contain A Ticking Time(r) Bomb? <https://web.archive.org/web/20201022213341/https://www.bacidore.com/post/does-your-algorithm-contain-a-ticking-time-r-bomb>.
- Baldauf, M. and J. Mollner (2020). High-frequency trading and market performance. *Journal of Finance* 75(3), 1495–1526.
- Barbon, A., M. Di Maggio, F. Franzoni, and A. Landier (2019). Brokers and Order Flow Leakage: Evidence from Fire Sales. *Journal of Finance* 74(6), 2707–2749.
- Battalio, R., S. A. Corwin, and R. Jennings (2016). Can brokers have it all? on the relation between make-take fees and limit order execution quality. *The Journal of Finance* 71(5), 2193–2238.
- Battalio, R. H., B. C. Hatch, and M. Sağlam (2018). The cost of exposing large institutional orders to electronic liquidity provider. *Working Paper*.
- Beason, T. and S. Wahal (2020). The Anatomy of Trading Algorithms. *SSRN Working Paper*.
- Bird, A. (2018). Taxation and Executive Compensation: Evidence from Stock Options. *Journal of Financial Economics* 127(2), 285–302.
- Bouveret, A., C. Guillaumie, C. Aparicio Roqueiro, C. Winkler, and S. Nauhaus (2014). High-frequency Trading Activity in EU Equity Markets. *ESMA Economic Report* (2014–1).
- Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan (2015). Trading fast and slow: Colocation and liquidity. *Review of Financial Studies* 28(12), 3407–3443.

- Brogaard, J., T. Hendershott, S. Hunt, and C. Ysusi (2014). High-Frequency Trading and the Execution Costs of Institutional Investors. *Financial Review* 49(2), 345–369.
- Brogaard, J., T. Hendershott, and R. Riordan (2014). High-frequency Trading and Price Discovery. *Review of Financial Studies* 27(8), 2267–2306.
- Brugler, J. (2015). Into the light: dark pool trading and intraday market quality on the primary exchange. *Working Paper*.
- Brunnermeier, M. K. and L. H. Pedersen (2005). Predatory Trading. *Journal of Finance* 60(4), 1825–1863.
- CESR (2007). Best Execution Under MiFID: Questions and Answers.
- Chan, L. K. and J. Lakonishok (1993). Institutional trades and intraday stock price behavior. *Journal of Financial Economics* 33(2), 173–199.
- Chan, L. K. and J. Lakonishok (1997). Institutional Equity Trading Costs: NYSE versus Nasdaq. *Journal of Finance* 52(2), 713–735.
- Chan, L. K. C. and J. Lakonishok (1995). The Behavior of Stock Prices Around Institutional Trades. *Journal of Finance* 50(4), 1147–1174.
- Chen, M. and C. Garriott (2020). High-frequency trading and institutional trading costs. *Journal of Empirical Finance* 56, 74–93.
- Conrad, J., K. Johnson, and S. Wahal (2001). Institutional Trading and Soft Dollars. *Journal of Finance* 56(1), 397–416.
- Cready, W., A. Kumas, and M. Subasi (2014). Are Trade Size-Based Inferences About Traders Reliable? Evidence from Institutional Earnings-Related Trading. *Journal of Accounting Research* 52(4), 877–909.
- De Bondt, W. F. M. and R. Thaler (1985). Does the Stock Market Overreact? *Journal of Finance* 40(3), 793–805.
- Dugast, J. and T. Foucault (2018). Data Abundance and Asset Price Informativeness. *Journal of Financial Economics* 130(2), 367–391.
- Easley, D. and M. O’Hara (1987). Price, trade size, and information in securities markets. *Journal of Financial Economics* 19(1), 69–90.
- Engle, R., R. Ferstenberg, and J. Russell (2012). Measuring and Modeling Execution Cost and Risk. *Journal of Portfolio Management* 38(2), 14–28.
- Foucault, T., J. Hombert, and I. Roşu (2016). News Trading and Speed. *Journal of Finance* 71(1), 335–382.

- Frazzini, A., I. Ronen, and T. J. Moskowitz (2018). Trading costs. *SSRN Working Paper*.
- Gider, J., S. Schmickler, and C. Westheide (2021). High-Frequency Trading and Price Informativeness. *Working Paper*.
- Hasbrouck, J. and G. Saar (2013). Low-latency Trading. *Journal of Financial Markets* 16(4), 646–679.
- Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *Journal of Finance* 66(1), 1–33.
- Imbens, G. W. (2000). The Role of the Propensity Score in Estimating Dose-Response Functions. *Biometrika* 87(3), 706–710.
- Imbens, G. W. and J. M. Wooldridge (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature* 47(1), 5–86.
- Jegadeesh, N. and S. Titman (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *Journal of Finance* 48(1), 65–91.
- Jones, C. M. and M. L. Lipson (1999). Execution Costs of Institutional Equity Orders. *Journal of Financial Intermediation* 8(3), 123–140.
- Jones, C. M. and M. L. Lipson (2001). Sixteenths: Direct Evidence on Institutional Execution Costs. *Journal of Financial Economics* 59(2), 253–278.
- Keim, D. B. and A. Madhavan (1996). The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects. *Review of Financial Studies* 9(1), 1–36.
- Keim, D. B. and A. Madhavan (1997). Transactions Costs and Investment Style: An Interexchange Analysis of Institutional Equity Trades. *Journal of Financial Economics* 46(3), 265–292.
- Keim, D. B. and A. Madhavan (1998). The Cost of Institutional Equity Trades. *Financial Analysts Journal* 54(4), 50–69.
- Korajczyk, R. A. and D. Murphy (2019). High-Frequency Market Making to Large Institutional Trades. *Review of Financial Studies* 32(3), 1034–1067.
- Kyle, A. S. (1985). Continuous Auctions and Insider Trading. *Econometrica* 53(6), 1315–1335.
- Linden, A., S. D. Uysal, A. Ryan, and J. L. Adams (2016). Estimating Causal Effects for Multivalued Treatments: A Comparison of Approaches. *Statistics in Medicine* 35(4), 534–552.
- Madhavan, A., M. Richardson, and M. Roomans (1997). Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks. *Review of Financial Studies* 10(4), 1035–1064.

- Menkveld, A. J. (2013). High-frequency Trading and the New Market Makers. *Journal of Financial Markets* 16(4), 712–740.
- Menkveld, A. J. (2016). The Economics of High-Frequency Trading: Taking Stock. *Annual Review of Financial Economics* 8(1).
- Perold, A. F. (1988). The Implementation Shortfall: Paper versus Reality. *Journal of Portfolio Management* 14(3), 4.
- Putniņš, T. J. and J. Barbara (2016). Heterogeneity in the Effects of Algorithmic and High-Frequency Traders on Institutional Transaction Costs. *Working Paper*.
- Robins, J. M. and A. Rotnitzky (1995). Semiparametric Efficiency in Multivariate Regression Models with Missing Data. *Journal of the American Statistical Association* 90(429), 122–129.
- Saar, G. (2001). Price Impact Asymmetry of Block Trades: An Institutional Trading Explanation. *Review of Financial Studies* 14(4), 1153–1181.
- Sağlam, M. (2020). Order anticipation around predictable trades. *Financial Management* 49(1), 33–67.
- Shkilko, A. and K. Sokolov (2020). Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity, and Trading Costs. *Journal of Finance* 75(6), 2899–2927.
- Tong, L. (2015). A Blessing or a Curse? The Impact of High Frequency Trading on Institutional Investors. *Working Paper*.
- Torre, N. (1997). Barra market impact model handbook. *BARRA Inc., Berkeley* 208.
- Uysal, S. D. (2015). Doubly Robust Estimation of Causal Effects with Multivalued Treatments: An Application to the Returns to Schooling. *Journal of Applied Econometrics* 30(5), 763–786.
- Van Biesebroeck, J., E. Yu, and S. Chen (2015). The Impact of Trade Promotion Services on Canadian Exporter Performance. *Canadian Journal of Economics/Revue canadienne d'économique* 48(4), 1481–1512.
- Van Kervel, V. and A. J. Menkveld (2019). High-frequency trading around large institutional orders. *Journal of Finance* 74(3), 1091–1137.
- Weller, B. M. (2017). Does Algorithmic Trading Reduce Information Acquisition? *Review of Financial Studies* 31(6), 2184–2226.
- Yang, L. and H. Zhu (2020). Back-running: Seeking and hiding fundamental information in order flows. *Review of Financial Studies* 33(4), 1484–1533.
- Zarinelli, E., M. Treccani, J. D. Farmer, and F. Lillo (2015). Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity* 1(02), 1550004.

Figure 1: Difference between Agency and Proprietary Trades of Colocated Exchange Members

This graph shows the distribution of within-exchange member differences in execution costs in bps between agency and proprietary trades that a given exchange member executes on a trading venue where they are colocated. The estimates are based on regressions that include exchange member-capacity-colocation fixed effects in addition to baseline control variables and stock-day and intraday fixed effects. Standard errors are clustered by exchange member and stock and estimated using bootstrapping with 1,000 iterations. The ranges indicate 90% confidence intervals.




Figure 2: Periodicity in Order Submissions

This graph shows the distribution of child order submissions within a minute at the exchange member-stock-capacity level. The number of submissions is given in deviations of the logarithm of the number of order submissions at a given second from their overall mean. The horizontal axis indicates the second within a minute. The data is split according to whether an exchange member is colocated at an exchange where a given stock can be traded.

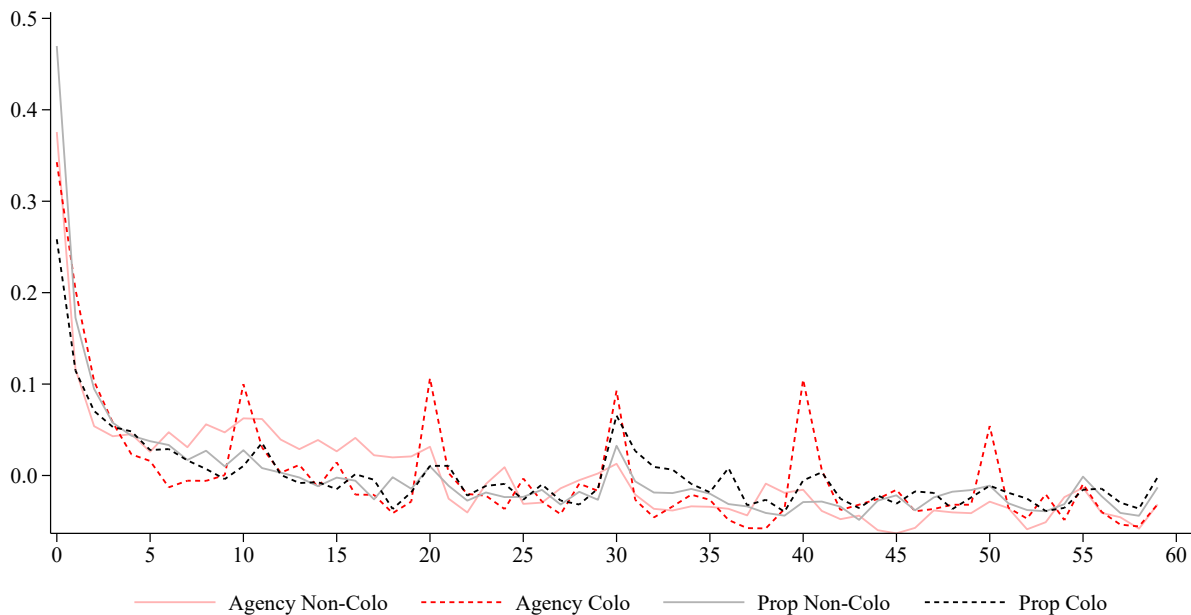


Table 1: Stock Summary Statistics

This table shows summary statistics on the 83 sample stocks. The stocks are sorted into terciles based on their average market capitalization during the sample period. *Market Cap.* is the average market capitalization in 1bn Euro. *Trading Value* is the average daily lit trading value across all exchanges in 1mn Euro. *Share_{PM}* and *Share_{MTF}* are the market shares of the stock's primary market and of the three MTFs combined, respectively, in percent. Spread is the average relative bid-ask spread at the primary market in percent.

Size	Market Cap.		Trading Value		Share _{PM}		Share _{MTF}		Spread	
	Mean	P50	Mean	P50	Mean	P50	Mean	P50	Mean	P50
Small	0.83	0.81	2.25	1.57	77.5	78.9	22.5	21.1	0.30	0.20
Medium	3.33	3.04	11.86	7.39	69.8	71.7	30.2	28.3	0.20	0.14
Large	22.57	9.72	74.53	36.44	60.6	57.4	39.4	42.6	0.09	0.08
Total	8.75	2.94	29.00	7.29	62.3	72.4	37.7	27.6	0.20	0.13

Table 2: Exchange Member Summary Statistics

This table shows exchange member level summary statistics. *Active Stocks* is the number of different stocks that the exchange member is actively trading. *Active Countries* is the number of different countries where those stocks are listed. *Total Volume* is the total trading value of parent orders in 1mn Euro. *# Orders* is the number of parent orders. *Active on MTF* is a binary variable set to one if the exchange member trades on any MTF, scaled by 100. *Agency* is the number of agency trades relative to all trades of an exchange member, in percent. *Colo*, *Colo_{PM}*, and *Colo_{TQ}* are binary variables set to one if the exchange member is colocated on any venue, any primary market, or Turquoise, respectively, scaled by 100. *Colo # PMs* is the number of different primary markets at which the exchange member is colocated. *Colo #|Colo* is the number of different venues at which the exchange member is colocated conditional on being colocated at any venue.

	Mean	P5	P50	P95
Panel A: All Exchange Members (N = 139)				
Active Stocks	14.2	3.0	7.0	54.0
Active Countries	3.2	1.0	2.0	8.0
Total Volume	48.1	2.8	10.7	288.8
# Orders	84.3	11.0	29.0	497.0
Active on MTF	35.3	0.0	0.0	100.0
Agency	56.5	0.0	68.6	100.0
Colo	29.5	0.0	0.0	100.0
Colo _{PM}	28.8	0.0	0.0	100.0
Colo _{TQ}	8.6	0.0	0.0	100.0
Colo # PMs	0.6	0.0	0.0	4.0
Colo # Colo	2.2	1.0	1.0	6.0
Panel B: Members Trading in Both Capacities (N = 36)				
Active Stocks	29.4	6.0	23.5	68.0
Active Countries	5.3	1.0	6.0	8.0
Total Volume	138.3	9.5	78.1	489.8
# Orders	229.5	26.0	134.5	755.0
Active on MTF	66.7	0.0	100.0	100.0
Agency	51.5	20.3	49.8	86.0
Colo	55.6	0.0	100.0	100.0
Colo _{PM}	55.6	0.0	100.0	100.0
Colo _{TQ}	25.0	0.0	0.0	100.0
Colo # PMs	1.3	0.0	1.0	4.0
Colo # Colo	2.8	1.0	2.5	5.5

Table 3: Parent Order Summary Statistics

This table shows parent order level summary statistics. *Colo* is a binary variables set to one if the exchange member is colocated on the stock's primary market or the MTFs, scaled by 100. *Buy* is a binary variable for buy trades, scaled by 100. *Order Size* is the size of the parent order in 1k Euro. *Information* is the signed return from the quote midpoint before the first to the closing price one day after the last child order execution in bps. *Duration* is the trading-hours adjusted time the parent order is worked in the market. *Aggressiveness* is the volume of executions of marketable orders relative to the size of the parent order, in percent. *MTF* is the percentage of trade volume executed on an MTF. *Execution Costs* are the execution costs in bps. *Effective Half Spread* and *Price Impact* are the average effective half-spread and 1 minute price impact of all aggressive (liquidity-taking) and passive (liquidity providing) child order executions for a given parent order in bps. *Volume HFT* is the fraction of trade volume with an HFT counterparty in percent. *Volume agr. HFT* is the fraction of trading volume where the counterparty is a liquidity-taking HFT in percent. The last columns show the results of t-tests for differences between agency and proprietary orders. $\frac{\Delta}{\sigma}$ is the difference standardized by the standard deviation of the overall sample.

	All Orders				Agency Orders				Proprietary Orders				Difference		$\frac{\Delta}{\sigma}$	
	Mean	P5	P50	P95	Mean	P5	P50	P95	Mean	P5	P50	P95	Δ	t-stat		
Panel A: Colocation and Capacity																
Colo	52.95	0.00	100.00	100.00	51.35	0.00	100.00	100.00	54.47	0.00	100.00	100.00	3.12	(3.38)***	0.06	
Agency	48.65	0.00	0.00	100.00	100.00	100.00	100.00	100.00	0.00	0.00	0.00	0.00				
Panel B: Order Characteristics																
Buy	50.03	0.00	100.00	100.00	49.98	0.00	0.00	100.00	50.07	0.00	100.00	100.00	0.08	(0.09)	0.00	
Trade Size	570.7	109.0	265.3	1,992.6	560.3	109.5	265.4	2,016.5	580.5	108.5	264.9	1,976.7	20.14	(1.15)	0.02	
Information	-9.41	-330.08	-7.32	312.08	-6.71	-333.87	-3.41	311.50	-11.97	-322.22	-10.67	312.20	-5.27	(-1.38)	-0.03	
Duration	4.95	2.13	4.30	8.68	5.09	2.15	4.45	9.05	4.82	2.13	4.15	8.49	-0.27	(-5.47)***	-0.10	
Panel C: Execution Characteristics																
Aggressiveness	41.70	0.00	37.58	97.78	44.05	1.23	40.00	99.86	39.46	0.00	35.10	94.89	-4.59	(-8.28)***	-0.15	
Child Order Executions	72.87	4.00	35.00	273.00	65.72	3.00	30.00	258.00	79.65	4.00	40.00	286.50	13.92	(6.92)***	0.13	
MTF	28.39	0.00	3.63	100.00	17.20	0.00	0.00	100.00	38.99	0.00	28.40	100.00	21.79	(33.71)***	0.59	
Panel D: Execution Costs																
Execution Costs	-4.46	-112.78	-2.60	102.69	-2.22	-110.35	-0.47	108.74	-6.59	-115.96	-4.56	94.91	-4.38	(-3.51)***	-0.06	
Effective Half Spread ^{aggr.}	4.97	0.63	3.30	14.46	5.20	0.69	3.39	14.77	4.75	0.60	3.16	14.09	-0.45	(-3.18)***	-0.06	
Effective Half Spread ^{pass.}	-3.62	-10.50	-2.48	0.74	-3.49	-9.72	-2.45	1.06	-3.74	-11.15	-2.50	0.49	-0.25	(-1.83)	-0.04	
Price Impact _{1min} ^{aggr.}	3.49	-4.52	1.97	15.78	3.75	-4.39	2.08	16.44	3.25	-4.65	1.89	15.07	-0.51	(-2.88)***	-0.06	
Price Impact _{1min} ^{pass.}	-4.67	-15.99	-3.30	3.52	-4.68	-16.35	-3.22	3.95	-4.67	-15.16	-3.35	2.94	0.01	(0.05)	0.00	
Panel E: HFT Interactions																
Volume HFT	31.41	2.18	30.21	64.96	29.50	1.38	27.73	62.49	33.22	3.52	32.24	66.35	3.72	(10.58)***	0.19	
Volume agr. HFT	20.83	0.00	17.87	53.65	19.07	0.00	15.60	51.73	22.49	0.00	20.10	55.11	3.42	(10.61)***	0.19	
		N = 11,724					N = 5,704					N = 6,020				

Table 4: IPW Performance

This table shows parent order level averages of and differences between agency (A) and proprietary (P) orders originating from exchange members that are (C) or are not (\bar{C}) collocated on a venue where the stock can be traded. In panel A, each observations is equally weighted. In panel B, the observations are weighted by the inverse probability of being in their observed category of collocation and agency, where the probabilities are estimated in a first stage using a multinomial logit model with the same independent variables as in Table 5. t-statistics of the differences are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively. $\frac{\Delta}{\sigma}$ is the difference standardized by the standard deviation of the respective variable in the unweighted sample.

	A \bar{C}	P \bar{C}	AC	PC	P \bar{C} - A \bar{C}			AC - A \bar{C}			PC - A \bar{C}			AC - P \bar{C}			PC - P \bar{C}			AC - PC		
					Δ	t	$\frac{\Delta}{\sigma}$	Δ	t	$\frac{\Delta}{\sigma}$	Δ	t	$\frac{\Delta}{\sigma}$	Δ	t	$\frac{\Delta}{\sigma}$	Δ	t	$\frac{\Delta}{\sigma}$	Δ	t	$\frac{\Delta}{\sigma}$
Panel A: Unweighted Sample																						
log Order Size	12.62	12.78	12.78	12.63	0.16	(6.39)***	0.18	0.16	(6.69)***	0.18	0.01	(0.51)	0.01	-0.00	(-0.03)	-0.00	-0.15	(-6.21)***	-0.16	0.15	(6.54)***	0.16
log Market Trading	15.86	16.01	15.39	15.00	0.16	(3.07)***	0.08	-0.47	(-9.43)***	-0.25	-0.86	(-17.63)***	-0.45	-0.63	(-12.57)***	-0.33	-1.01	(-20.87)***	-0.53	0.39	(8.07)***	0.20
Buy	48.40	50.09	51.49	50.05	1.69	(1.26)	0.03	3.09	(2.33)**	0.06	1.65	(1.28)	0.03	1.39	(1.05)	0.03	-0.05	(-0.04)	-0.00	1.44	(1.13)	0.03
Volatility t_{-10min}	0.09	0.09	0.08	0.08	0.01	(2.66)***	0.07	-0.01	(-2.59)***	-0.07	-0.00	(-0.61)	-0.02	-0.01	(-5.12)***	-0.14	-0.01	(-3.32)***	-0.09	-0.00	(-2.07)**	-0.05
Return t_{-10min}	-2.08	-1.71	0.67	0.70	0.37	(0.51)	0.01	2.75	(3.86)***	0.10	2.77	(4.08)***	0.10	2.37	(3.28)***	0.09	2.40	(3.46)***	0.09	-0.03	(-0.04)	-0.00
Information	-21.04	-13.42	6.87	-10.76	7.62	(1.34)	0.04	27.91	(5.07)***	0.14	10.27	(1.90)*	0.05	20.29	(3.76)***	0.10	2.66	(0.50)	0.01	17.63	(3.44)***	0.09
Duration	5.11	5.12	5.07	4.57	0.01	(0.08)	0.00	-0.04	(-0.54)	-0.01	-0.54	(-8.08)***	-0.20	-0.05	(-0.62)	-0.02	-0.55	(-8.16)***	-0.20	0.50	(7.39)***	0.19
Aggressiveness	45.62	39.60	42.57	39.35	-6.01	(-7.30)***	-0.20	-3.05	(-3.72)***	-0.10	-6.27	(-8.05)***	-0.21	2.97	(3.74)***	0.10	-0.25	(-0.34)	-0.01	3.22	(4.31)***	0.11
log Child Exec.	2.99	3.61	3.74	3.62	0.61	(16.10)***	0.47	0.75	(22.71)***	0.58	0.63	(19.50)***	0.48	0.14	(3.97)***	0.11	0.02	(0.53)	0.01	0.12	(4.19)***	0.09
MTF	17.09	41.82	17.31	36.62	24.73	(23.28)***	0.67	0.22	(0.26)	0.01	19.53	(22.41)***	0.53	-24.51	(-24.86)***	-0.67	-5.20	(-5.15)***	-0.14	-19.31	(-24.86)***	-0.52
Panel B: Weighted Sample																						
log Order Size	12.76	12.71	12.71	12.71	-0.05	(-1.70)*	-0.05	-0.04	(-1.41)	-0.05	-0.05	(-1.67)*	-0.05	0.01	(0.31)	0.01	0.00	(0.09)	0.00	0.01	(0.23)	0.01
log Market Trading	15.67	15.24	15.53	15.57	-0.43	(-2.60)***	-0.22	-0.15	(-2.05)**	-0.08	-0.11	(-1.63)	-0.06	0.29	(1.74)*	0.15	0.33	(2.02)**	0.17	-0.04	(-0.66)	-0.02
Buy	50.58	49.53	49.68	51.85	-1.05	(-0.66)	-0.02	-0.90	(-0.58)	-0.02	1.26	(0.87)	0.03	0.15	(0.09)	0.00	2.31	(1.52)	0.05	-2.17	(-1.48)	-0.04
Volatility t_{-10min}	0.09	0.09	0.09	0.09	0.00	(0.25)	0.01	-0.00	(-0.52)	-0.02	-0.00	(-0.19)	-0.01	-0.00	(-0.67)	-0.03	-0.00	(-0.40)	-0.01	-0.00	(-0.35)	-0.01
Return t_{-10min}	-0.16	-1.26	-0.77	-0.71	-1.10	(-1.20)	-0.04	-0.61	(-0.70)	-0.02	-0.55	(-0.72)	-0.02	0.49	(0.50)	0.02	0.55	(0.63)	0.02	-0.06	(-0.07)	-0.00
Information	-9.33	-6.70	-11.76	-11.44	2.63	(0.39)	0.01	-2.43	(-0.37)	-0.01	-2.11	(-0.33)	-0.01	-5.06	(-0.75)	-0.02	-4.74	(-0.72)	-0.02	-0.32	(-0.05)	-0.00
Duration	5.16	4.98	4.95	5.03	-0.17	(-2.14)**	-0.06	-0.21	(-2.58)***	-0.08	-0.12	(-1.50)	-0.05	-0.03	(-0.40)	-0.01	0.05	(0.59)	0.02	-0.08	(-0.99)	-0.03
Aggressiveness	41.66	41.59	39.37	41.17	-0.07	(-0.08)	-0.00	-2.29	(-2.48)**	-0.08	-0.49	(-0.57)	-0.02	-2.22	(-2.37)**	-0.07	-0.42	(-0.47)	-0.01	-1.81	(-2.02)**	-0.06
log Child Exec.	3.52	3.49	3.57	3.43	-0.03	(-0.72)	-0.02	0.05	(1.33)	0.04	-0.09	(-2.22)**	-0.07	0.09	(2.04)**	0.07	-0.06	(-1.36)	-0.04	0.14	(3.89)***	0.11
MTF	29.24	26.64	31.11	28.68	-2.60	(-2.03)**	-0.07	1.87	(1.25)	0.05	-0.56	(-0.47)	-0.02	4.47	(3.44)***	0.12	2.04	(2.20)**	0.06	2.43	(2.00)**	0.07

Table 5: Baseline Execution Cost Analysis

This table shows parent order level regressions with stock-day fixed effects in addition to quarter-hour intraday fixed effects for the beginning of the order's execution. The dependent variable is the execution cost in basis points. *Order Size* is the parent order size in Euro. *Market Trading* is the total trading volume of all other trades during the execution of the parent order. *Buy* is a binary variable equal to one for buy orders. *Volatility_{t-10min}* is the standard deviation of 10ms quote midpoint returns and *Return_{t-10min}* is the trade direction signed quote midpoint return in the 10min interval before the first execution of a child order, in bps, respectively. Since the previous two variables are based on less than 10min of observations for early trades, a binary variable for trades starting before 9:10h is included in all models but not reported. *Information* is the permanent price impact as measured by the trade direction signed return from the quote midpoint immediately before the parent order to the closing quote midpoint one day after the last child order execution, in bps. *Duration* is the time difference between the first and last execution of a child order, adjusted for exchange trading hours. *Aggressiveness* is the volume of executions of marketable orders relative to the size of the parent order, in percent. *Child Order Executions* is the number of individual child order executions. *MTF* is the total volume executed on MTFs relative to the size of the parent order, in percent. The observations are weighted by the inverse probability of being in their observed category of colocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the same independent variables while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)
<i>log Order Size</i>	2.597*** (2.63)	2.663*** (2.71)	3.330*** (3.21)
<i>log Market Trading</i>		-3.370* (-1.67)	-3.441* (-1.71)
Buy	1.870 (1.29)	1.936 (1.34)	1.985 (1.37)
Volatility _{t-10min}	17.310 (1.35)	17.035 (1.33)	17.081 (1.33)
Return _{t-10min}	0.005 (0.17)	0.006 (0.19)	0.006 (0.18)
Information	0.146*** (27.65)	0.145*** (27.57)	0.145*** (27.59)
Duration	-1.675*** (-4.72)	-1.033* (-1.88)	-0.975* (-1.78)
Aggressiveness	0.305*** (12.32)	0.305*** (12.36)	0.310*** (12.45)
<i>log Child Order Executions</i>	4.114*** (5.73)	4.142*** (5.77)	3.280*** (4.06)
MTF			0.065*** (2.97)
Stock-Day FE	✓	✓	✓
Intraday FE	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓
Exchange Member FE	—	—	—
N = 11,724			

Table 6: Exchange Member Colocation and HFT

This table shows parent order level regressions with stock-day and intraday fixed effects. The dependent variable is the execution cost in basis points. *Agency* is a binary variable equal to one for agency and zero for proprietary orders. *Colo* is a binary variable equal to one if the exchange member is collocated on the primary market or on the MTF. *HFT* is the fraction of the parent order executed against an HFT counterparty in percent. *HFT^{Aggr.}* and *HFT^{Pass.}* is the fraction of the parent order executed against an HFT counterparty where the HFT takes or provides liquidity, respectively, in percent. *ColoHFT^{Aggr.}* and *ColoHFT^{Pass.}* is the fraction of the parent order executed against a liquidity taking or providing HFT, respectively, on a venue where the exchange member is collocated, in percent. The HFT variables are centered at zero. The remaining variables capturing order and execution characteristics are defined as in Table 5 and all included in all models. The observations are weighted by the inverse probability of being in their observed category of collocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the independent variables as in Table 5 while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	4.465*** (2.87)	4.365*** (2.80)	0.555 (0.26)	0.682 (0.32)	1.029 (0.49)	3.654 (0.96)
Colo		-2.301 (-1.45)	-5.978*** (-2.71)	-5.960*** (-2.71)	-5.659*** (-2.58)	-5.857*** (-2.62)
Colo×Agency			7.223** (2.48)	7.165** (2.47)	6.289** (2.16)	6.024** (2.02)
HFT				0.088** (1.97)		
HFT ^{Aggr.}					0.365*** (6.65)	0.451*** (5.77)
HFT ^{Pass.}					-0.304*** (-4.15)	-0.382*** (-3.87)
ColoHFT ^{Aggr.}						-0.211** (-2.01)
ColoHFT ^{Pass.}						0.334*** (2.85)
Agency×HFT ^{Aggr.}						-0.090 (-0.88)
Agency×HFT ^{Pass.}						-0.056 (-0.48)
Agency×ColoHFT ^{Aggr.}						0.254 (1.64)
Agency×ColoHFT ^{Pass.}						-0.189 (-1.07)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	—	—	—	—	—	—
N = 11,724						

Table 7: Exchange Member Colocation and HFT with Exchange Member FE

This table shows parent order level regressions with stock-day, intraday, and additionally exchange member fixed effects. The dependent variable in all models is the execution cost in basis points. The independent variables are defined as in Table 5 and Table 6. As being collocated on the MTF does not vary within exchange members, only collocation at the stock's primary market is considered. Consequently, trading against HFTs has been split by trading venue type. The HFT variables are centered at zero. The observations are weighted by the inverse probability of being in their observed category of collocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the independent variables as in Table 5 while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	9.217*** (4.06)	9.215*** (4.06)	4.991* (1.70)	5.111* (1.74)	5.180* (1.77)	4.529 (1.52)
Colo _{PM}		-2.937 (-0.90)	-6.258* (-1.74)	-6.126* (-1.70)	-6.182* (-1.71)	-6.879* (-1.91)
Colo _{PM} × Agency			7.834** (2.22)	7.763** (2.20)	7.704** (2.18)	7.659** (2.11)
HFT				0.091** (2.10)		
HFT _{Aggr.}					0.351*** (6.34)	0.352*** (4.75)
HFT _{Pass.}					-0.297*** (-4.28)	-0.353*** (-3.99)
ColoHFT _{PM} ^{Aggr.}						-0.110 (-0.91)
ColoHFT _{PM} ^{Pass.}						0.407*** (2.62)
Agency × HFT _{PM} ^{Aggr.}						0.038 (0.36)
Agency × HFT _{PM} ^{Pass.}						0.009 (0.07)
Agency × HFT _{MTF} ^{Aggr.}						0.115 (0.93)
Agency × HFT _{MTF} ^{Pass.}						-0.031 (-0.21)
Agency × ColoHFT _{PM} ^{Aggr.}						0.021 (0.13)
Agency × ColoHFT _{PM} ^{Pass.}						-0.361* (-1.73)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	✓	✓	✓	✓	✓	✓
N = 11,724						

Table 8: Order to Trade Ratios

This table shows parent order level regressions with stock-day and intraday fixed effects. The dependent variable is the number of child order submissions relative to the number of child order executions. The independent variables are defined as in Table 5 and Table 6. The model in column (2) additionally includes exchange member fixed effects. As being colocated on the MTF does not vary within exchange members, only colocation at the stock's primary market is considered in this model. The observations are weighted by the inverse probability of being in their observed category of colocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the independent variables as in Table 5 while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)
Agency	-1.425 (-0.42)	-0.424 (-0.16)
Colo	10.722** (2.12)	3.883 (0.99)
Colo×Agency	-9.679* (-1.80)	-10.481*** (-2.60)
Order Controls	✓	✓
Execution Controls	✓	✓
Stock-Day FE	✓	✓
Intraday FE	✓	✓
Inv. Prob. Weighting	✓	✓
Exchange Member FE	–	✓
$N = 10,411$		

Table 9: Child-order Execution Quality

This table shows parent order level regressions with stock-day fixed effects in addition to quarter-hour intraday fixed effects for the beginning of the order's execution. Columns (2) and (4) additionally include exchange member fixed effects. Each panel shows a separate set of regression results where the dependent variables are the effective half spread and the price impact after 1 second and 1 minute, respectively. All dependent variables are parent order level averages of the respective measures across all aggressive (liquidity-taking) or passive (liquidity-providing) child order executions of a given parent order. All three measures are expressed in basis points where positive (negative) values indicate a loss (gain) for the institution. The independent variables are defined as in Table 5 and Table 6. The observations are weighted by the inverse probability of being in their observed category of colocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the independent variables as in Table 5 while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)
	Aggr.	Aggr.	Pass.	Pass.
Panel A: Effective Half Spread				
Agency	0.421*** (3.09)	0.385* (1.92)	-0.009 (-0.07)	0.261 (1.49)
Colo	-0.739*** (-5.62)	-0.147 (-0.62)	-0.487*** (-3.82)	-0.302 (-1.40)
Colo×Agency	0.045 (0.26)	0.176 (0.72)	0.084 (0.50)	0.170 (0.87)
Panel B: Price Impact at 1 Second				
Agency	0.335* (1.67)	0.768*** (2.80)	-0.331* (-1.93)	0.010 (0.05)
Colo	-0.694*** (-3.80)	0.290 (0.88)	-0.753*** (-5.09)	-0.035 (-0.15)
Colo×Agency	0.351 (1.45)	-0.096 (-0.30)	0.616*** (3.05)	0.788*** (3.28)
Panel C: Price Impact at 1 Minute				
Agency	-0.091 (-0.31)	0.390 (1.00)	-0.741*** (-2.91)	0.117 (0.38)
Colo	-1.009*** (-3.35)	0.084 (0.18)	-0.853*** (-3.79)	-0.385 (-1.05)
Colo×Agency	0.930** (2.34)	0.096 (0.19)	0.873*** (2.70)	0.517 (1.38)
Order Controls	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓
Exchange Member FE	—	✓	—	✓
	N = 10,926		N = 10,969	

Table 10: Relative Fees and Exchange Member Colocation

This table shows parent order level regressions with stock-day and intraday fixed effects. The dependent variables in all models are the exchange fees in basis points. The independent variables are defined as in Table 5 and Table 6. The corresponding baseline model including the control variables can be found in Table C1 in the internet appendix. The models in columns (4)-(6) additionally include exchange member fixed effects. As being collocated on the MTF does not vary within exchange members, only collocation at the stock's primary market is considered in these models. The observations are weighted by the inverse probability of being in their observed category of collocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the independent variables as in the baseline model in Table C1 while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	0.029*** (3.21)	0.030*** (3.38)	0.063*** (5.41)	0.067*** (6.62)	0.067*** (6.58)	0.094*** (5.65)
Colo		0.028*** (3.56)	0.060*** (4.38)		-0.011 (-0.60)	0.011 (0.58)
Colo×Agency			-0.063*** (-3.80)			-0.054*** (-2.98)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	-	-	-	-	-	-
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	-	-	-	✓	✓	✓
$N = 11,724$						

Table 11: Execution Cost Risk

This table shows parent order level regressions with stock-day and intraday fixed effects. The dependent variables are the absolute residuals (in columns (1) and (2)) or squared residuals (in columns (3) and (4)) of the execution cost regressions shown in model (3) of Table 6 (for columns (1) and (3)) and Table 7 (for columns (2) and (4)), respectively. The independent variables are defined as in Table 5 and Table 6. The models in columns (2) and (4) additionally include exchange member fixed effects. As being colocated on the MTF does not vary within exchange members, only colocation at the stock's primary market is considered in these models. The observations are weighted by the inverse probability of being in their observed category of colocation and capacity, where the probabilities are estimated in a first stage using a multinomial logit model with the same independent variables except for colocation and capacity while the standard errors are estimated using bootstrapping with 1,000 iterations. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	Abs. Residuals		Sq. Residuals	
	(1)	(2)	(3)	(4)
<i>log</i> Order Size	1.056** (2.56)	0.726* (1.68)	99.068 (1.52)	81.114 (1.23)
<i>log</i> Market Trading	7.003*** (3.64)	6.300*** (3.26)	990.406*** (3.10)	825.692*** (2.78)
Buy	-0.032 (-0.04)	-0.273 (-0.36)	-85.322 (-0.62)	-98.286 (-0.75)
Volatility t_{-10min}	6.705 (1.09)	5.324 (0.88)	310.106 (0.27)	329.209 (0.31)
Return t_{-10min}	0.027* (1.78)	0.024 (1.62)	3.321 (1.20)	3.524 (1.36)
Information	-0.002 (-0.60)	0.000 (-0.13)	-0.388 (-0.75)	-0.184 (-0.38)
Duration	0.483 (1.14)	0.602 (1.45)	82.168 (1.14)	108.089 (1.62)
Agency	1.088 (1.05)	3.037** (2.06)	149.992 (0.91)	446.405** (2.01)
Colo	3.191*** (2.95)	-1.113 (-0.40)	442.918** (2.56)	-226.732 (-0.57)
Colo \times Agency	-1.450 (-1.00)	-2.447 (-1.24)	-222.289 (-0.95)	-397.380 (-1.25)
Stock-Day FE	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓
Exchange Member FE	—	✓	—	✓
$N = 11,724$				

Internet Appendix

In this internet appendix, we first provide evidence that our key results remain robust to not using IPW in [Appendix A](#). In [Appendix B](#), we show that the results are not meaningfully affected by exchange fees. [Appendix C](#) shows the baseline results for exchange fees. [Appendix D](#) shows that the main results hold if we exclude predictors that are directly affected by the execution algorithm.

A Unweighted estimates

First, we estimate the panel regression for parent order execution costs without employing IPW. [Table A1](#) reports the baseline results. The coefficients for log order size, Information, parent order duration, aggressiveness, log child order executions, and MTF trading are consistent in magnitude and statistical significance with the main specification employing IPW. Additionally, the market trading variable that is weakly significant in the main specification becomes insignificant in the estimation without IPW. Finally, the trade direction variable and historical volatility become weakly significant in the specification without IPW. [Table A2](#) reports the relationship between execution costs on the one hand and trading capacity, colocation, and HFT interactions on the other hand. Again, the results remain very similar to the estimations with IPW. After controlling for the baseline variables, agency orders have a higher execution cost than proprietary orders, colocated proprietary orders have lower execution costs than colocated agency orders, and aggressive (passive) HFT activity is positively (negatively) associated with order execution costs. The sign, magnitude, and statistical significance of the coefficients is also similar between the two specifications.

Table A3 contains the results after including exchange-member fixed effects. While the overall conclusions remain unchanged, we do observe some differences compared to the IPW results. Specifically, the coefficient for the execution costs of agency orders, while positive, is no longer weakly significant after controlling for colocation and HFT. At the same time, the coefficient of the colocation dummy and its interaction with the agency dummy in specifications (3) to (6) is larger in absolute terms and gains in statistical significance. Finally, the relationship between execution costs and total HFT as well as its decomposition into aggressive and passive strategies remain largely unchanged.

Figure A1 shows that the broker-level difference in fixed effects between agency and proprietary orders is similar to the output with IPW. Overall, the above findings indicate that our key results are robust to the weighting scheme employed in the estimation process.

Insert Table A1 about here

Insert Table A2 about here

Insert Table A3 about here

Insert Figure A1 about here

B Execution costs incorporating exchange fees

Next, we estimate all the main results by computing execution costs after including exchange fees and subtracting any rebates. For exchanges offering fee schedules dependent on the members' total trading volume, we assume the fee/rebate corresponding to the highest volume threshold. As some of the competing MTFs operate a maker-taker fee structure,

offering rebates to liquidity providers, the sum of fees incurred can be zero or even negative in some cases. We report the coefficients from the estimation with and without exchange member-fixed effects. The results are reported in [Table B1](#), [Table B2](#), and [Table B3](#). While we do observe small differences in the statistical significance of some coefficients, the overall conclusions concerning the relationship between execution costs and trader capacity, collocation, and HFT remain robust to including fees/rebates.

Insert [Table B1](#) about here

Insert [Table B2](#) about here

Insert [Table B3](#) about here

C Baseline model for exchange fees

[Table C1](#) shows the baseline results for exchange fees. Most predictors of execution costs do not predict exchange fees. Exchange fees decrease in order size and increase in execution duration, two variables that are generally positively correlated with each other.

Insert [Table C1](#) about here

D Robustness to excluding routing characteristics

In [Tables D1](#) and [D2](#), we repeat the main analyses contained in [Subsections 7.2](#) and [7.3](#) while excluding predictors reflecting characteristics of the execution algorithm. Thus, results reported in [Table D2](#) measure the total impact of capacity and collocation on execution costs

rather than the direct effect after conditioning on certain characteristics of the algorithm.²³ The results without exchange member fixed effects (models (1) to (3)) are qualitatively unchanged compared to those reported in [Subsection 7.2](#), with slightly larger effect sizes. The within-exchange member analysis (models (4) to (6)) provides results that also do not qualitatively differ from those reported in [Subsection 7.3](#), again with some minor differences in effect size estimates. The results of the baseline analysis shown in [Table D1](#) also do not qualitatively differ from those in [Table 5](#). The only quantitatively apparent difference is in the coefficient for order size, which can be explained by the fact that the excluded variables, and in particular the number of child orders, are correlated with the parent order size.

Insert [Table D1](#) about here

Insert [Table D2](#) about here

²³We use the term “algorithm” in a broad sense, including choices high-touch traders may make when executing a large order.

Figure A1: Difference between Agency and Proprietary Trades of Colocated Exchange Members without IPW

This graph shows the distribution of within-exchange member differences in execution costs in bps between agency and proprietary trades that a given exchange member executes on a trading venue where they are colocated. The estimates are based on regressions that include exchange member-capacity-colocation fixed effects in addition to baseline control variables and stock-day and intraday fixed effects. Standard errors are clustered by exchange member and stock. The ranges indicate 90% confidence intervals.




Table A1: Baseline Execution Cost Analysis without IPW

This table shows parent order level regressions for the execution cost in basis points similar to Table 5, but without inverse probability weighting. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)
<i>log</i> Order Size	2.646*** (2.90)	2.671*** (2.92)	3.379*** (3.56)
<i>log</i> Market Trading		-1.226 (-0.61)	-1.280 (-0.64)
Buy	2.133* (1.68)	2.140* (1.69)	2.173* (1.71)
Volatility $t-10\text{min}$	19.553* (1.77)	19.645* (1.78)	19.733* (1.78)
Return $t-10\text{min}$	0.000 (0.00)	0.000 (0.00)	0.000 (-0.01)
Information	0.142*** (31.40)	0.142*** (31.39)	0.142*** (31.37)
Duration	-1.304*** (-4.18)	-1.065** (-2.03)	-1.001* (-1.91)
Aggressiveness	0.305*** (14.11)	0.305*** (14.10)	0.309*** (14.28)
<i>log</i> Child Order Executions	4.238*** (6.25)	4.243*** (6.26)	3.396*** (4.58)
MTF			0.060*** (3.12)
Stock-Day FE	✓	✓	✓
Intraday FE	✓	✓	✓
Inv. Prob. Weighting	-	-	-
Exchange Member FE	-	-	-
$N = 11,724$			

Table A2: Exchange Member Colocation and HFT without IPW

This table shows parent order level regressions for the execution cost in basis points similar to Table 6, but without inverse probability weighting. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	4.593*** (3.07)	4.477*** (2.98)	-0.304 (-0.15)	-0.234 (-0.11)	0.377 (0.18)	3.765 (1.08)
Colo		-1.840 (-1.29)	-6.212*** (-3.11)	-6.195*** (-3.11)	-5.819*** (-2.94)	-5.309*** (-2.70)
Colo×Agency			8.840*** (3.21)	8.826*** (3.21)	7.793*** (2.83)	7.085** (2.55)
HFT				0.081** (2.09)		
HFT ^{Aggr.}					0.347*** (7.02)	0.468*** (6.54)
HFT ^{Pass.}					-0.296*** (-4.70)	-0.334*** (-3.72)
ColoHFT ^{Aggr.}						-0.258*** (-2.72)
ColoHFT ^{Pass.}						0.280*** (2.63)
Agency×HFT ^{Aggr.}						-0.122 (-1.33)
Agency×HFT ^{Pass.}						-0.044 (-0.41)
Agency×ColoHFT ^{Aggr.}						0.244* (1.68)
Agency×ColoHFT ^{Pass.}						-0.224 (-1.36)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	—	—	—	—	—	—
Exchange Member FE	—	—	—	—	—	—
N = 11,724						

Table A3: Exchange Member Colocation and HFT with Exchange Member FE without IPW

This table shows parent order level regressions for the execution cost in basis points similar to Table 7, but without inverse probability weighting. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	8.668*** (4.12)	8.543*** (4.06)	3.607 (1.37)	3.756 (1.43)	3.843 (1.46)	3.489 (1.30)
Colo _{PM}		-4.806 (-1.63)	-8.305*** (-2.60)	-8.150** (-2.55)	-8.259*** (-2.59)	-8.417*** (-2.64)
Colo _{PM} × Agency			8.865*** (2.66)	8.764*** (2.63)	8.674*** (2.60)	8.515** (2.48)
HFT				0.095** (2.44)		
HFT ^{Aggr.}					0.343*** (6.93)	0.362*** (5.28)
HFT ^{Pass.}					-0.274*** (-4.39)	-0.274*** (-3.20)
ColoHFT _{PM} ^{Aggr.}						-0.123 (-1.14)
ColoHFT _{PM} ^{Pass.}						0.219 (1.55)
Agency × HFT _{PM} ^{Aggr.}						-0.026 (-0.28)
Agency × HFT _{PM} ^{Pass.}						-0.025 (-0.22)
Agency × HFT _{MTF} ^{Aggr.}						0.088 (0.75)
Agency × HFT _{MTF} ^{Pass.}						-0.047 (-0.33)
Agency × ColoHFT _{PM} ^{Aggr.}						0.133 (0.82)
Agency × ColoHFT _{PM} ^{Pass.}						-0.231 (-1.15)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	—	—	—	—	—	—
Exchange Member FE	✓	✓	✓	✓	✓	✓

$N = 11,724$

Table B1: Baseline Execution Cost Analysis after Exchange Fees

This table shows parent order level regressions similar to Table 5, but for the execution cost after adding exchange fees and subtracting rebates in basis points. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)
<i>log</i> Order Size	2.281** (2.25)	2.378** (2.36)	2.975*** (2.79)
<i>log</i> Market Trading		-4.906** (-2.27)	-4.970** (-2.30)
Buy	1.925 (1.30)	2.021 (1.37)	2.064 (1.40)
Volatility $t-10\text{min}$	15.370 (1.10)	14.968 (1.07)	15.010 (1.07)
Return $t-10\text{min}$	0.002 (0.04)	0.003 (0.07)	0.002 (0.07)
Information	0.154*** (25.15)	0.154*** (25.13)	0.154*** (25.13)
Duration	-1.857*** (-4.92)	-0.922 (-1.58)	-0.870 (-1.49)
Aggressiveness	0.303*** (11.94)	0.304*** (12.00)	0.309*** (12.06)
<i>log</i> Child Order Executions	4.261*** (5.77)	4.302*** (5.83)	3.530*** (4.23)
MTF			0.059*** (2.59)
Stock-Day FE	✓	✓	✓
Intraday FE	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓
Exchange Member FE	-	-	-
$N = 11,724$			

Table B2: Exchange Member Colocation and HFT after Exchange Fees

This table shows parent order level regressions similar to Table 6, but for the execution cost after adding exchange fees and subtracting rebates in basis points. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	4.600*** (2.88)	4.496*** (2.81)	0.578 (0.27)	0.692 (0.32)	1.044 (0.48)	4.081 (1.02)
Colo		-2.420 (-1.48)	-6.200*** (-2.70)	-6.184*** (-2.69)	-5.879** (-2.57)	-6.124*** (-2.61)
Colo×Agency			7.426** (2.45)	7.375** (2.44)	6.485** (2.14)	6.209** (1.99)
HFT				0.079* (1.73)		
HFT ^{Aggr.}					0.360*** (6.47)	0.448*** (5.65)
HFT ^{Pass.}					-0.319*** (-4.25)	-0.400*** (-3.95)
ColoHFT ^{Aggr.}						-0.202* (-1.87)
ColoHFT ^{Pass.}						0.351*** (2.88)
Agency×HFT ^{Aggr.}						-0.104 (-0.99)
Agency×HFT ^{Pass.}						-0.067 (-0.55)
Agency×ColoHFT ^{Aggr.}						0.260 (1.64)
Agency×ColoHFT ^{Pass.}						-0.189 (-1.04)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	-	-	-	-	-	-
N = 11,724						

Table B3: Exchange Member Colocation and HFT with Exchange Member FE after Exchange Fees

This table shows parent order level regressions similar to Table 7, but for the execution cost after adding exchange fees and subtracting rebates in basis points. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%–level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	9.234*** (3.99)	9.232*** (3.98)	4.378 (1.45)	4.486 (1.48)	4.555 (1.51)	3.988 (1.30)
Colo _{PM}		−3.296 (−0.99)	−7.112* (−1.91)	−6.993* (−1.88)	−7.051* (−1.89)	−7.689** (−2.06)
Colo _{PM} ×Agency			9.002** (2.45)	8.938** (2.43)	8.878** (2.41)	8.742** (2.31)
HFT				0.082* (1.84)		
HFT ^{Aggr.}					0.348*** (6.18)	0.356*** (4.72)
HFT ^{Pass.}					−0.316*** (−4.42)	−0.366*** (−4.00)
ColoHFT ^{Aggr.} _{PM}						−0.108 (−0.87)
ColoHFT ^{Pass.} _{PM}						0.416*** (2.60)
Agency×HFT ^{Aggr.} _{PM}						0.011 (0.10)
Agency×HFT ^{Pass.} _{PM}						−0.021 (−0.16)
Agency×HFT ^{Aggr.} _{MTF}						0.100 (0.79)
Agency×HFT ^{Pass.} _{MTF}						−0.033 (−0.22)
Agency×ColoHFT ^{Aggr.} _{PM}						0.052 (0.30)
Agency×ColoHFT ^{Pass.} _{PM}						−0.350 (−1.63)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	✓	✓	✓	✓	✓	✓
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	✓	✓	✓	✓	✓	✓

$N = 11,724$

Table C1: Baseline Relative Fees Analysis

This table shows parent order level regressions for the sum of exchange fees as dependent variables, taking any rebates into account. The regressions are otherwise similar to Table 5, but while excluding execution characteristics as independent control variables. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)
<i>log</i> Order Size	−0.015*** (−4.44)	−0.015*** (−4.40)
<i>log</i> Market Trading		−0.003 (−0.44)
Buy	0.002 (0.38)	0.002 (0.39)
Volatility $t-10\text{min}$	−0.008 (−0.29)	−0.008 (−0.30)
Return $t-10\text{min}$	0.000 (−1.48)	0.000 (−1.48)
Information	0.000 (−0.07)	0.000 (−0.09)
Duration	0.008*** (4.80)	0.008*** (3.85)
Stock-Day FE	✓	✓
Intraday FE	✓	✓
Inv. Prob. Weighting	✓	✓
Exchange Member FE	−	−
$N = 11,724$		

Table D1: Baseline Execution Cost Analysis excluding Execution Characteristics

This table shows parent order level regressions for the execution cost in basis points similar to Table 5, but while excluding execution characteristics as independent control variables. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)
<i>log</i> Order Size	5.450*** (7.03)	5.527*** (7.12)
<i>log</i> Market Trading		-3.237 (-1.53)
Buy	2.339* (1.67)	2.394* (1.72)
Volatility _{t-10min}	14.878 (1.09)	14.712 (1.08)
Return _{t-10min}	0.002 (0.08)	0.003 (0.09)
Information	0.145*** (28.40)	0.145*** (28.33)
Duration	-1.342*** (-3.84)	-0.717 (-1.32)
Stock-Day FE	✓	✓
Intraday FE	✓	✓
Inv. Prob. Weighting	✓	✓
Exchange Member FE	-	-
<i>N</i> = 11,724		

Table D2: Exchange Member Colocation excluding Execution Characteristics

This table shows parent order level regressions for the execution cost in basis points similar to Table 6 and Table 7, but while excluding execution characteristics as independent control variables. Models (4)-(6) include exchange member fixed effects. As being collocated on the MTF does not vary within exchange members, only collocation at the stock's primary market is considered in these models. Standard errors are clustered by exchange member and stock. t-statistics are given in parentheses. ***, **, * denotes significance at the 1%, 5%, 10%-level, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Agency	3.851** (2.54)	3.784** (2.49)	-2.145 (-1.07)	7.601*** (3.88)	7.573*** (3.88)	3.508 (1.41)
Colo		-1.656 (-1.06)	-7.359*** (-3.35)		-3.569 (-1.09)	-6.852* (-1.94)
Colo×Agency			11.168*** (3.94)			7.936** (2.24)
Order Controls	✓	✓	✓	✓	✓	✓
Execution Controls	–	–	–	–	–	–
Stock-Day FE	✓	✓	✓	✓	✓	✓
Intraday FE	✓	✓	✓	✓	✓	✓
Inv. Prob. Weighting	✓	✓	✓	✓	✓	✓
Exchange Member FE	–	–	–	✓	✓	✓
$N = 11,724$						