

The Anatomy of Trading Algorithms*

Tyler Beason
WP Carey School of Business
Arizona State University
Tempe, AZ 85287

Sunil Wahal
WP Carey School of Business
Arizona State University
Tempe, AZ 85287
Sunil.Wahal@asu.edu

September 2020

* We are grateful to Lili Ge for research assistance and an anonymous firm for providing algorithmic trading data. We thank the Center for Investment Engineering at ASU for financial support. Wahal is a consultant to Avantis Investors. Avantis did not provide any data or funding for this study. We thank Robert Almgren, Hank Bessembinder, Terry Hendershott, Marc Lipson, Ananth Madhavan, and seminar participants at the European Finance Association meeting, the Microstructure Exchange, Microstructure Online Asia-Pacific, the NBER Big Data workshop, Purdue University, Southern Methodist University, and the University of Virginia for helpful comments.

The Anatomy of Trading Algorithms

Abstract

We study the anatomy of four widely used institutional trading algorithms representing \$675 billion in demand from 961 institutions between 2012 and 2016. Parent orders generate hundreds of child orders which strategically employ price, time-in-force, and display priority rules to navigate the tradeoff between the desire to trade and minimizing transaction costs. Child orders incur price impact at the time they are submitted to the book regardless of whether or not they are (ex post) filled, and even when passively priced relative to the prevailing quote. The intra-parent distribution of child orders is non-random, generating strategic runs which oscillate between the aggressive or passive side of the spread. Despite algorithmic attempts to reduce their influence, programmatic child-level price, time-in-force, and display choices aggregate up to parent-level trading costs borne by investors.

1. Introduction

Classical models of market microstructure were developed prior to the advent of electronic trading. In these models, market makers were distinct parties with no inherent trade motive who set quotes passively under a zero profit condition. Trade sizes were equal to order sizes, and price discovery occurred when active traders hit quotes. In modern markets, trades are no longer the key unit of analysis because institutional trading involves order splitting and the submission of large numbers of passive orders, many of which go unexecuted. High frequency market makers still attempt to set regret-free quotes, but the nature of information and adverse selection is tied to trading horizons. Recognizing these issues, O’Hara (2015) issues a clarion call for research to update learning models and evidence in this trading environment.

Recent research that follows this direction has focused on issues related to speed differentials and high frequency traders (HFTs).¹ But electronic trading has also changed the trading processes of large institutional investors, a counterparty to HFTs. Again, O’Hara (2015, pg. 258) points out that while, “much has been made of the activities of high frequency traders, the behavior of non high frequency traders is also now radically different...”. These changes are, of course, endogenous. Instead of manually working orders to find counterparties, brokerage firms now provide suites of algorithmic execution services that institutions can access directly from their Execution Management Systems (EMS). Our understanding of precisely how these large institutions trade, their impact on prices, and implications for how information is incorporated into prices, is severely inhibited by lack of data. Large institutions do not wish to release data on their trading practices because safeguarding trading information is even more important in an environment in which HFTs are viewed as adversarial. Brokerage firms also do not release data for fear of revealing proprietary designs. The upshot is opacity in the trading behavior of a group of investors that, by some estimates, generate over 80 percent of total trading volume.

In this paper, we employ proprietary data to study buy-side trading algorithms, a modern-day analogue to Keim and Madhavan’s (1995) anatomical study of (manual) institutional trading. Our purpose is two-fold. First, we open up the black box of trading algorithms, linking design

¹ The list of papers that seek to understand the behavior and impact of HFTs is large (see surveys by Biais and Wolley (2011), O’Hara (2015), and Menkveld (2016)). Much of this research focuses on the effects of the speed advantage of HFT’s and its implications for other agents’ behavior and market equilibrium.

features and choices back to constructs underlying market microstructure research. To the extent that wayward trading algorithms are sometimes associated with market structure induced volatility, our analysis is also of interest to practitioners and regulators concerned with the stability of the trading environment – anatomy informs diagnosis, and if necessary, cure. Second, we study three important choice variables for child orders generated by an algorithm: submission price, time-in-force, and order display, all of which are at the heart of an algorithm’s attempt to trade while minimizing trading costs. We quantitatively tie these primitive choice variables to child and parent-level trading costs borne by investors.

The data we employ are powerful, from both a statistical and economic perspective. The time series is long (2012-2016), and cross-sectional coverage is large (over 5,000 securities). We study four widely used non-bespoke trading algorithms used by 961 unique institutions representing over \$675 billion in aggregate demand. The anatomical structures of a given algorithm (e.g., VWAP) are similar across brokerage firms which means that insights from the data are generalizable. We observe parent orders and *all* downstream child orders, both unexecuted and those that result in fills. Being able to observe trading intentions, instead of just realizations, is particularly important for understanding the tradeoff between trading costs and execution risk. The data also capture important details associated with each child order, including limit prices, time-in-force qualifications, venue decisions, trading fees, and other such attributes that allow us to examine fundamental tradeoffs in electronic trading. All data are time-stamped to the millisecond, allowing computations of short horizon price movements uncontaminated by latency issues.

The average parent order attempts to trade \$287,000 over 84 minutes, equivalent to 4.80 percent of volume over the duration of the order. The 2.3 million parent orders in the sample generate over 300 million child orders, indicative of the enormous velocity generated in attempts to trade. Less than 0.40 percent of child orders are market orders. By comparison, data from Rule 605 and Rule 606 reports show that retail investors usage of market orders is over 50 percent (Boehmer, Jones, and Zhang (2017), and Kelley and Tetlock (2013)). The dominant order type is limit orders (81.5 percent), followed by PEG orders (18.1 percent), which are dark, exchange-disseminated limit orders dynamically “pegged” to the NBBO. This usage is in stark contrast to

most sequential or strategic trade models in which trade takes place when market orders interact with limit orders (Harris and Hasbrouck (1996), Parlour (1998), Large (2009), and others), or first generation limit order models in which limit orders are uniformed (Glosten (1994), Seppi (1997), Parlour (1998), Foucault (1999), Biais, Martimort and Rochet (2000), Foucault, Kadan and Kandel (2005), Roşu (2009)). To the extent that the institutional users of algorithmic trading are informed, usage data suggest that the majority of information is impounded into prices via limit (not market) orders, which begs the question of limit order price impact.²

We estimate price impact both at the time a child order is submitted to the book, and subsequent to its execution. We distinguish between orders submitted at various price points of the prevailing spread, defining aggressive orders as those submitted at the far side of the NBBO (e.g. buy orders at the National Best Ask), intra-spread orders as those in between the NBBO, and passive orders as those submitted at or behind the near side of the NBBO.³ Of the 169 million limit (PEG) orders in our sample, 65 percent (24 percent) are passive, 11 percent (67 percent) are inside the spread, and 24 percent (9 percent) are aggressive. At the 10 second horizon, aggressive orders move quotes by 2.03 basis points post-submission, but even passive orders move prices by 0.84 basis points post-submission. This has important consequences for theory and practice. It is inconsistent with traditional models in which price discovery takes place exclusively via trades. From a practical perspective, it indicates that the mere submission of a child order, even a passive one that is not guaranteed execution, contributes to the overall cost of trading. This non-zero cost of potentially failed attempts to trade is ignored in standard models of optimal execution (e.g. Bertsimas and Lo (1988), Almgren and Chriss (2000), and others). The post-execution price impact of orders is also large, as liquidity providers move quotes in response to trading. Cross-sectional variation in these price movements is positively related to the size of the order, and whether the order is displayed.⁴ It is also larger in high VIX periods, an outcome is that is outside

² Bloomfield, O'Hara, and Saar (2003) provide experimental evidence that suggests that limit orders can be informed. Kaniel and Liu (2006), and Collin-Dufresne and Fos (2015) provide equivalent empirical evidence. Brolley and Malinova (2020) construct a model in which informed investors use both market and limit orders, and in which limit orders have positive price impact.

³ Aggressively priced orders may or may not be marketable, depending on other order-specific and matching engine details. We discuss these issues in Section 5.1.

⁴ The lower price impact and higher time-to-fill of non-displayed orders is consistent with price discovery modelled as by Zhu (2014). Interestingly, our data show no evidence of dark pool usage that provides a size discovery function (Duffie and Zhu (2017)).

of an algorithm's locus of control but related to aggregate risk bearing capacity. Broadly, our child-level empirical results are remarkably consistent with Riccò, Rindi, and Seppi (2020) in which informed trading takes place via limit orders with varying degrees of price aggressiveness, and in which price impact is related to the order book and market volatility.

Child orders are strung together in strategic “runs” in which they are consecutively passive or aggressive.⁵ The average parent order contains about 63 such runs lasting about 566 seconds, with each run containing almost nine consecutive child orders. The price aggressiveness of a new run depends on executions in the prior run; as with child orders, these runs tradeoff the desire to trade with cost mitigation. The outcome is alternating phases of providing and taking liquidity. The child and run-level data indicate that while agency algorithms trade directionally, they provide liquidity in the process, although not in the classical sense. Unlike the market makers in Glosten and Milgrom (1985), they are motivated to trade. They do not attempt to profit from round-trip trades over very short horizons like electronic market makers. They also do not appear to lean against the wind in the sense of Grossman and Miller (1988) and Weill (2007), and have no affirmative obligation to supply liquidity. Instead, liquidity is supplied almost incidentally, as a byproduct of not wanting to pay the spread. These distinctions are economically important. When posting resting quotes, agency algorithms compete with electronic market makers. This competition can have consequences for markets, as illustrated by Li, Wang and Ye (2019) who show more complex equilibria than Budish, Cramton, and Shim (2015) and Menkveld and Zoican (2017) in which only HFTs supply liquidity. Given this behavior, it is unsurprising that agency algorithms and high frequency traders trade with and against each other (Van Kervel and Menkveld (2019), Korajczyk and Murphy (2019)).

The *raison d'être* of a trading algorithm is to trade while adding sufficient noise so that parent intentions are not perfectly revealed by child characteristics and submission protocols. To investigate this, we estimate regressions of parent-level trading costs on child-level choices, focusing on variables tied to child-level price, time-in-force, and display choices. To our knowledge, we are the first to study this linkage because it requires us to observe intent (parent

⁵ Hasbrouck and Saar (2013) impute strategic runs using linked messages from Nasdaq TotalView-ITCH data and suggest that their runs largely capture the activity of high frequency traders. They find that their runs are associated with improvements in market quality measures.

orders), process (all submitted child orders), and realizations (executed child orders). Controlling for market conditions and security-specific characteristics, a one standard deviation increase in the percentage of aggressive child orders per parent increases expected trading costs by between 10 and 23 percent. Similarly, allowing for more patient child orders reduces parent-level trading costs; a one standard deviation increase in the percentage of day orders (rather than impatient immediate-or-cancel (IOC) orders) reduces parent trading costs between 15 and 72 percent. The influence of order display is algorithm dependent. Although all four algorithms share the common goal of wishing to trade desired quantities at the lowest possible cost, they differ in the relative importance attached to volume versus cost. These differences show up in their anatomical structures, particularly in the usage of passive versus aggressive orders, and in display frequencies. In the two algorithms that are especially cost-sensitive and more likely to use passive child orders, increasing the percentage of displayed child orders lowers parent trading costs by about 30 percent. In the other two algorithms that are more volume sensitive, increasing the percentage of displayed orders raises parent trading costs by about the same amount. Notwithstanding these algorithm-specific fixed effects, there is no escaping price impact: child level choices are reflected in parent-level costs ultimately borne by investors.

The remainder of the paper is organized as follows. Section 2 describes the electronic trading process for institutional investors, moving downstream from order creation to execution. Section 3 discuss the proprietary data and sample. Section 4 outlines the anatomy of trading algorithms, from parent to child orders. Section 5 focuses on child orders, assessing execution likelihood and transaction costs. Section 6 examines the link between parent-level trading costs and child-level choices. Section 7 concludes.

2. The Electronic Institutional Trading Process

We provide a brief description of the trading process at buy-side institutional trading desks to facilitate the analysis of trading algorithms. Individual investment management firms customize their processes to account for variations in investment styles, management structure, portfolio turnover, and other such firm-specific features. The description of the work flow below is deliberately generic so that it highlights key decisions in the process. We describe the process in

a downstream manner, sequentially from order generation to submission, execution, and trade reporting. We focus on key nodes of transmission of information between counterparties (brokerage firms, execution venues, etc.) that are relevant to the economic analysis later in the paper.

At most institutions, orders are generated at the portfolio manager level. These portfolios can represent individual funds, fund classes, separate accounts, commingled accounts, or some combination of the above. Orders are typically entered into an Order Management System (OMS), whose functionality includes position management, cash management, communication between portfolio managers and the trading desk, and ex-post allocation of trades back to portfolios. The OMS may or may not combine orders for the same security from multiple portfolio managers into one blocked order. Orders are “staged” and also subject to compliance requirements to satisfy internal rules as well as regulatory obligations. In some firms, two or more portfolios seeking to buy and sell the same security on the same day, may be internally crossed. This type of internal cross, referred to as a Rule 17A-7 transaction, reduces brokerage costs and the price impact of trading but is subject to strict regulatory requirements. In cases where the buy order is larger than the sell (or vice versa), the post-cross residual is sent onward to execution systems.

From the OMS, block orders are routed to the execution management system (EMS), which interfaces with market data and allows for electronic routing of orders to brokers, as well as monitoring of child orders. In some firms, the OMS and EMS can be integrated (referred to as automated staging) which minimizes errors and speeds up the trading process.⁶ In many cases, the EMS may provide Direct Market Access (DMA) to a trader, which allows him/her to push trades directly out to the marketplace instead of going through brokerage firms. In virtually all cases, the communication between institutions, the marketplace, and algorithmic brokerage firms takes place via the Financial Information Exchange (FIX) protocol. FIX requires communication of trade details in FIX tags so that all details of an order are captured in a standardized manner regardless of the counterparty. From our perspective, this is crucial because it ensures that data with particular FIX tags always represent the same fields and values.

⁶ A single OMS can typically handle multiple asset classes such as equities and fixed income, and can therefore be connected to multiple EMS platforms that allow for market data and execution of different asset classes.

The order received by an algorithmic brokerage firm from an institution's EMS is referred to as a parent order. A single parent order is assigned to one trading algorithm and not split among different algorithms. Once the parent order is received, the algorithm goes to work in a mostly automated fashion, submitted child orders to trading venues either sequentially or concurrently ("spraying"). With appropriate systems and technology, institutional traders can monitor parent orders, child orders, and fills on their EMS. Once fills are received by the brokerage firm's servers, they are sent back to the institution's OMS via FIX, with each step of the transmission receiving a separate timestamp. Finally, the institution's OMS allocates shares to individual funds and accounts based on a pre-defined set of rules that may vary across institutions. Our data record all activity downstream from the parent order but not the ex post allocations to individual funds.

3. Data and Sample

3.1 Data sourcing

Our data are provided by a large algorithmic trading firm that provides execution services to institutional clients. The firm is well-established and widely regarded as providing superior algorithmic execution services to institutional investors. Its client base is diverse, including buy-side long-only investment managers, long-short investment managers, and hedge funds that run a diverse set of investment strategies. It is among the top ten brokerage firms by volume.

The data consist of all parent orders received by the firm from its clients to be executed using four trading algorithms for US stocks between 2012 and 2016. As with many proprietary datasets, implicit selection bias and generalizability of conclusions are important concerns. Since customized algorithms are built to reflect the order flow and preferences of specific institutions, we request data from non-bespoke single stock trading algorithms to mitigate these concerns. These are standardized algorithms widely used by many algorithmic trading providers in markets throughout the world. For instance, volume-weighted average price (VWAP), time-weighted average price (TWAP), implementation shortfall (IS), target close (TC), volume target (VT), and percent of volume (POV) are emblematic algorithms widely used by buy-side firms; variants of

these are provided off-the-shelf by many brokerage firms.⁷ Large providers of such algorithms also white-label their offerings to smaller brokerage firms. We restrict our attention to single-stock algorithms to focus attention on stock-specific execution issues without concerning ourselves with the covariance structure of short horizon price movements across firms. This would be the case, for example, if we considered pairs trading or basket trading algorithms.

3.2 Data Elements

The data consist of daily files that correspond to parent and child orders. Each parent order is uniquely identified with a client ID, appropriately masked so that we cannot associate parent orders with particular institutions. The client ID is unique, however, so that institutions can be tracked over time. The parent order also identifies the type of algorithm used, which we label A, B, C, and D to ensure confidentiality. Other parent order information include a stock identifier (symbol), side (buy, sell, or short-sale indicators), the number of shares desired, start and end times, and parameters that pertain to price/volume constraints that buy-side traders can customize prior to algorithm initiation. Price constraints are represented by limit price beyond which the institution does not want to buy or sell. Volume constraints indicate the maximum percentage of volume that the parent order can participate in over a particular duration.

Each parent order is uniquely linked to the sequence of child orders and fills that it generates. Each child order is associated with fields that specify submission, cancellation and fill times alongside a host of other features. These include the order type (market, limit or PEG order), limit prices for limit orders, the PEG price (primary, midpoint, or far side) for PEG orders, display or non-display instructions, execution instructions which have to do with whether the order is to be held, traded over the day etc. (FIX Tag 18), time-in-force (FIX Tag 59) which specifies whether the order is immediate or cancel (IOC), day, etc., and the venue to which the child order is sent including specific dark pools.⁸ If the child order results in one or multiple fills, the data indicate the price and number of shares traded, a last liquidity indicator (corresponding to FIX Tag 851),

⁷ See the 2019 Algorithmic Trading Survey (<https://www.thetradenews.com/surveys/algorithmic-trading-survey-long-results-2019/>).

⁸ These algorithms employ exchange-supplied PEGs rather than synthetic PEGs so that if the NBBO changes, the order's PEG price is immediately updated.

which shows whether the execution added liquidity, removed liquidity, or was routed out, and the trading fee paid or rebate earned by the order.⁹

Two particular aspects of the child order data are important. First, we observe *all* child orders generated by a parent, regardless of whether they result in fills. This is critical because it allows us to assess price movements generated by the revelation of trading intentions, as opposed to only realized trades. Second, the algorithms use direct exchange feeds, not the consolidated SIP, and all of the timestamps that we observe are in milliseconds. This minimizes latency induced errors both in execution and in matching with market data.

3.3 Market Data

We match the algorithmic trading data with market data from daily Trade and Quote (TAQ) files with millisecond timestamps. We compute the NBBO following the procedures in Holden and Jacobsen (2014) with appropriate modifications for changes in data structures over the sample period. We also require total depth at the NBBO. Since only one trading venue can be the official NBBO at any point in time, depth at that venue does not necessarily represent total depth available at that price point. To compute the true total depth, we sum all depth available in all trading venues that are at the best bid or offer, regardless of whether they represent the official NBBO.

The algorithms use direct exchange feeds which are faster than the consolidated feed through the Securities Information Processor (SIP). Therefore, it is possible that latency in the NBBO (or the BBO for each venue) identified from the TAQ data is enough to affect inferences. To check if this is the case, we examine the Participant Timestamps field reported by TAQ. This timestamp, only available after August 2015, is the time at which an exchange's matching engine processes a quote update. In our sample, the median lag between the SIP timestamp and the Participant timestamp is about 0.5 milliseconds. This lag is small relative to quote update frequencies and unlikely to change the NBBO or BBO significantly. Nonetheless, we also compute the NBBO assuming a conservative 5 millisecond lag and assess whether it influences

⁹ A child order can be partially filled if it trades with the residual of a larger counterparty order or result in multiple fills if the residual is held in place.

the assigned price aggressiveness of child orders. We find that it only influences price aggressiveness classifications in 0.1 percent of all child orders in our sample.

3.4 Sample Statistics

The data consist of 2.3 million parent orders sent by 961 unique buy-side firms over the 2012-2016 period. Cross-sectional coverage is quite comprehensive, including over 5,000 US-traded securities, including American Depositary Receipts (ADRs) and Exchange Traded Funds (ETFs). Parent orders represent over \$675 billion in aggregate demand over the period. These parent orders generate over 300 million child submissions, which represent \$2.1 trillion in notional volume. The fact that child order notional volume is much larger than parent volume is not surprising since many child orders go unexecuted. The aggregate amount of trading generated by these parent orders is \$388 billion, about 18 percent of notional child order volume.¹⁰

Figure 1 shows the distribution of parent order dollar volume across all 961 institutions. No single institution dominates the data. Even the two largest users of these trading algorithms constitute only 7.91 and 7.05 percent of total parent volume. As such, single-institution selection biases that can confound inferences do not plague these data. The daily time series of algorithm use also shows no particular spikes or patterns.

4. Algorithm Anatomy

4.1 Algorithm Types

All algorithms seek to trade desired quantities as inexpensively as possible. Despite this common objective, each algorithm embeds within it differing degrees of sensitivity to volume versus trading costs. As described earlier, we are not permitted to reveal algorithm names or specific objectives. To provide economic context, Figure 2 places each of the four algorithms in a two-by-two volume and cost sensitivity grid. Algorithms B and D (which are similar and enclosed in the same circle) are more sensitive to trading costs than algorithms A and C. Algorithm

¹⁰ Incomplete parent-level execution could occur for a variety of reasons. For instance, it could be because the trading desk is not proficient at setting parameters of the algorithm, or perhaps because portfolio managers submit orders without the benefit of a sufficiently accurate trading cost model so that price movements are larger than the value of the expected benefit. Regardless, incomplete execution represents opportunity costs, which we defer to future work.

C is more sensitive to volume than algorithms A, B, and D. We stress that these are relative, not absolute, differences. As becomes apparent in subsequent tests, these differences represent a meaningful source of variation in the usage of various types of child orders, as well as realized trading costs.

4.2 Parent Orders

Panel A of Table 1 shows the number and aggregate value (in \$ billions) of all parent orders, as well as separately for each algorithm for buys, sells, and short sales. Both by number and dollar value, all four algorithms receive considerable usage. Even the lowest usage algorithm (A) generates 184,000 parent orders with an aggregate value of \$49 billion. The largest, algorithm D, generates over 1.7 million parent orders with an aggregate value of \$407 billion. Buy orders are generally more frequent than sells, but short sales are as frequent as sells. This is likely due to the composition of the client base which includes long-short investment managers.

The first few rows of Panel B shows statistics on the dollar value of the parent orders. Across all algorithms, average parent order size is \$287,000. There is considerable skewness as median parent size is substantially smaller (\$17,000). The standard deviations of parent order sizes are also quite large, often five times the mean. Another common way to measure order size is by scaling parent size with average daily volume over the prior 20 days. By this metric, average parent size is 38 basis points of average daily volume. Once again, skewness and large variation is apparent, with the median being substantially smaller than the mean, and the standard deviation four to five times the mean. We also compute a measure of order size by scaling parent dollar volume over the duration of the order (from the start time to the end time of the order) by actual (realized) volume, referred to as interval volume or participation volume. This measure is often used to obtain a sense of footprint the algorithm generates during the time it is active. Average parent volume as a percentage of interval volume is 4.80 percent. Within each algorithm there is very little variation across buys, sells, and short sales.

Panel C provides statistics on the duration of parent orders. The average parent order lasts 84.19 minutes. Algorithm A has shorter duration (averaging between 17 and 20 minutes), followed by algorithm C (between 31 and 36 minutes). Algorithms B and D attempt to trade for substantially

longer intervals, between 81 and 104 minutes, consistent with higher cost sensitivity shown in Figure 2. There is also considerable variation in parent order duration with each algorithm, likely linked to order size. The length of time an algorithm is active in the marketplace is important because it affords the algorithm more time to manage the tradeoff between execution likelihood and transaction costs. Presaging our results on strategic runs, the longer a parent order is “live”, the more opportunity it has to flip between providing and taking liquidity.

Institutional traders can impose price and/or volume floors and caps on trading algorithms at the parent level. These are implemented via price limits beyond which the algorithm cannot trade, and/or the percentage of rolling volume in which they participate. Figure 3 shows the percent of parent orders constrained by price limits, volume limits, and price-volume limits for each algorithm. The figure also shows the percentage of parent orders that are cancelled before trading the desired number of shares. The data show significant usage of price and volume floors and caps, as well as high cancellation rates. Variation in cancellation rates suggests that it is endogenous to algorithm design and execution expectations. These design mechanisms indicate that while institutions turn over the mechanics of trading to an algorithm over which they have no influence, they can (and do) exercise risk controls directly.

Panel D of Table 1 contains parent-level implicit trading costs, computed by scaling the weighted average transactions prices of all executed child orders within a parent by the prevailing midpoint before the start of the parent order. The first row shows average parent-level trading costs in basis points, weighted by the size of the parent order for the entire sample. Across all algorithms, the weighted average cost of parent orders is 7.33 basis points. By comparison, Frazzini, Israel, and Moskowitz (2018) report an average equivalent cost of about 10 basis points. Since our subsequent analysis focuses on a set of parent orders with at least 50 child orders and trading at least one basis point of daily volume, we also report costs for this set of “large orders”. The average cost across all orders rises to 9.04 basis points.

4.3 Child Orders

Panel A of Table 2 shows the numerical and dollar distribution of child orders. On average, a parent order spawns 126 child orders, of which 38 result in fills. In dollar terms, the average

total dollar value of child orders is \$918,000 resulting in \$165,000 in executed trades. Both metrics indicate patience in the submission process, as many child orders do not execute. Panel B shows selected characteristics of child orders. For each parent order, we compute the percentage of child orders with a particular characteristic based on dollar values. We then report the average percentage across parents in a group. In the section of the panel titled “order type”, un-indented rows show the average percentage of market, limit and PEG orders across all parent orders. Market orders are extremely rare, constituting less than 0.38 percent of all parent orders. By comparison, Boehmer, Jones, and Zhang (2017), and Kelley and Tetlock (2013) report that market orders are 50 percent of order flow from retail investors. The vast majority of algorithmic child orders are either limit orders (81.48 percent), followed by PEG orders (18.12 percent).¹¹ The latter are immediately repriced when the NBBO moves and are therefore subject to lower execution risk than equivalently priced limit orders. Indented rows show the percentage of limit and PEG orders with various time-in-force qualifications. Time-in-force (FIX code 59) can take on seven different values, but in our data over 99 percent fall into two categories: day or IOC orders.¹² Across all algorithms, limit orders are more likely to have full day discretion (61.9 percent) than be IOC (38.1). For PEG orders, the opposite is true, so that the majority of PEG orders are IOC (65.5 percent). Child orders that result in executions earn rebates from exchanges with non-inverted make-take fee schedules if they added liquidity but pay trading fees if they remove liquidity. We report the percentage of limit and PEG order executions from exchanges that add versus remove liquidity. Across all algorithms, 45.7 percent of limit order executions and 17.0 percent of PEG executions add liquidity.

The last three rows of Panel B show characteristics of child orders, pertaining to price, display, and venue choice. Aggressive orders are those priced at the far side of prevailing NBBO (e.g. buy orders priced at the best ask and sell orders priced at the best bid or lower). Over 42 percent of child orders are aggressively priced. About 75 percent of child orders from a parent are

¹¹ A limit order arriving to the book which contains a resting order at the same price but in the opposite direction generates an immediate execution (subject to minimum fill and other such qualifications).

¹² If not executed, day orders expire at the end of the trading day and are inherently more patient than IOC orders which must be executed in entirety as soon as they represent trading interest. FIX code 59 also allows for good till cancel orders, at the open orders, fill or kill orders, good till crossing orders, and good till date orders. These other categories represent less than 1 percent of the child orders in our data.

visible on exchange feeds and the remainder (25 percent) are either hidden orders on exchanges or posted to dark pools. Although we observe the complete set of venues utilized by each parent, we group them into lit versus dark categories and report the percent of child orders that use lit venues. Exchange use is widespread with over 77 percent of child orders going to Lit venues.

There are no meaningful differences in order attributes across buys, sells, and short sales. The remainder of the paper therefore aggregates all trade sides. There are, however, differences in order attributes across algorithms. These differences correlate with the cost and volume tradeoffs in Figure 2. For example, algorithms B and D, which are similar in the sense that they have heightened sensitivity to trading costs and low sensitivity to volume, have similar child-to-fill ratios. They are also less likely to use aggressively priced orders but more likely to use displayed orders. In contrast, algorithms A and C, which are more sensitive to volume, make greater use of aggressively priced orders, and use PEG orders to maintain queue priority. The data suggest that tactical order type decisions are related to the algorithms' objectives.

5. Intra-Parent Analysis

In this section, we use child orders as the unit of analysis. We take the perspective that market participants do not observe parent level information, and cannot infer future order flow from current orders or their characteristics. Attempts to do so, termed order anticipation strategies, are often regarded as predatory. While the incentives to engage in order anticipation are clear, the price consequences depend on modelling assumptions (see, for example, Brunnermeier and Pedersen (2005), Bessembinder et al. (2016), and Yang and Zhu (2019)), and the empirical evidence is mixed. While interesting, empirically assessing order anticipation is outside the scope of our paper.

5.1 Child Order Choices

We start by examining child order choices. To focus the analysis, we study three economic primitives of child orders: submission price, order size, and whether the order is displayed. We place each child order in a five point grid similar to Biais, Hillion, and Spatt (1995), based on the aggressiveness of its submission price relative to the prevailing NBBO: (a) aggressive orders

placed at the far-side quote (ask price for buys, bid price for sells), (b) orders between the far-side quote and the midpoint, (c) midpoint orders, (d) orders between the near-side quote and midpoint, and (e) passive orders placed at or further away from the near-side quote (bid price or below for buys, ask price or above for sells). We use this nomenclature for expositional convenience to distinguish between various price points; the labels aggressive and passive are meant to be indicative rather than definitional. For instance, limit orders inside the spread are ‘passive’ in the sense that they are resting orders but they can also be viewed as ‘aggressive’ in the sense that improving the prevailing quote shows a keener willingness to trade.¹³ Orders that offer price improvement can only be placed when the NBBO is greater than the minimum tick size. As a result, the sampling distribution of securities for such orders is different from the remaining orders. To compare order sizes across securities, we scale the size of the order in shares by total depth available at the NBBO at the time the order is submitted to the market. Total depth at the NBBO is computed using depth available in all trading venues that are the best bid or offer, regardless of whether the venue is the official NBBO. Non-displayed limit orders are either orders routed to dark pools or exchange designated non-displayed orders. PEG orders are, by definition, non-displayed.

Table 3 shows the 25th percentile, the median, and the 75th percentile of scaled child order sizes by price aggressiveness categories, for all child orders, and separately for orders that are unfilled and filled. Panel A presents results for all limit orders, as well as separately for non-displayed and displayed limit orders. Panel B contains results for PEG orders.

Differences in order size across the submission price grid are systematic. The median scaled order size for aggressive limit orders is 13.3 percent of visible depth but only 3.0 percent for passive limit orders. Similarly, the median order size for PEG orders falls from 20.0 percent for aggressive orders to 9.0 percent for passive orders. In both panels, passive orders are smaller because lower order sizes reduce the exposure to adverse selection, the major cost faced by such

¹³ Aggressively priced limit orders may or may not be marketable. If order instructions are such that they match an exchange’s matching engine rule system, the order receives an immediate execution. If not, the order may not get executed right away, despite the fact that it is aggressively priced. For instance, aggressively priced dark limit orders are not guaranteed execution because they are dark. Similarly, a displayed but aggressively priced limit order may not receive immediate execution because of a minimum fill instruction. Unfortunately, we do not observe the full set of order instructions or exchange matching engine protocols to isolate immediately executable marketable orders.

orders.¹⁴ Another way to limit adverse selection is by not displaying passive orders. Consistent with this, Panel A shows that non-displayed passive orders are systematically larger than displayed passive limit orders (28.0 versus 10.0 percent of depth). We also examine the time series distribution of price-size-display decisions. We do not show the results to conserve space but can report that time variation in these choices is not large. Overall, price, size, and display choices are systematic, endogenous, and jointly determined with an algorithm.

5.2 Execution Likelihood and Time-to-Fill

Table 4 shows execution frequencies, time-to-cancel (for unfilled child orders), and time-to-fill (for filled child orders). The sample consists of day orders (excludes IOC orders) that are permitted to rest on the book. The first column shows the number of child orders in each submission price category. Panel A contains counts for limit orders, separately, for all orders, non-displayed orders, and displayed orders. Panel B provides equivalent information for PEG orders. Columns labelled “N” show counts for unfilled and filled child orders.

Of the 121 million day child limit orders in our sample, 88 percent (106.8 million) are passive. Only 1 percent (1.4 million) are aggressive, and the remaining are priced inside the spread. In contrast, the majority of PEG orders are tied to the midpoint (36.9 million or 53 percent), although the number of passive orders (28.5 million) is more than seven times the number of aggressive orders (3.8 million). The number of limit orders with submission prices between the far-side quote and the midpoint, or between the near-side quote and the midpoint, is relatively small, 0.2 and 8.3 million respectively. As described earlier, orders posted between the near (passive) side of the quote and the midpoint definitionally narrow quotes. This is only possible if the NBBO is greater than the minimum tick size. To verify this, we calculate the NBBO (in pennies) for all orders in price aggressiveness categories and examine the distribution. The median spreads, moving from aggressive to passive orders are 0.01, 0.05, 0.01, 0.07, and 0.1 respectively. More importantly, in all 8.9 million orders that are priced between the near side quote and the midpoint, the quoted spread is greater than one penny.

¹⁴ Orders submitted just ahead of the prevailing quote, labelled (Agg., Midpoint), are substantially larger. The sample size for these orders is extremely small so we caution against drawing conclusions from this group.

Of the 106.8 million passive limit orders, 31 percent (33.3 million) are filled while the remaining are cancelled. For passive PEG orders, the fill rate is lower, at 19 percent (5.5 million out of 28.5 million). The fill rates for aggressive limit and PEG orders are 79 and 81 percent respectively. There are many reasons why aggressively priced orders may not get executed. Matching engine protocols, timing differentials, and other such factors can result in non-execution. For instance, if a displayed order is sent to a venue that does not have a contra-side order at the limit price and the order is non-routable (by instruction), it may not execute. Another possibility is that the order has a minimum fill instruction and is of a size that is larger than the contra-side order. Further, timing differentials can result in non-execution, or “lost races” as described in Aquilina, Budish, and O’Neill (2020).

Table 4 also reports the 25th percentile, median, and 75th percentile of time-to-cancel (for unfilled orders) and time-to-fill (for filled orders).¹⁵ As before, Panels A and B report statistics for limit and PEG orders respectively. For aggressive limit orders, the median time-to-cancel and time-to-fill are quite different, 14.77 and 0.12 seconds respectively. Similarly, for aggressive PEG orders, the median time-to-cancel 8.12 seconds but the time-to-fill is 0.79 seconds. These differences reinforce the notion that aggressively priced orders that are not immediately executable by an exchange’s rules have to sit and await an execution.

For passive limit orders, the median time-to-cancel is over three times as long as for aggressive orders (45.00 seconds compared to 14.77 seconds), and the median time-to-fill is over 225 times as long (27.78 seconds compared to 0.12 seconds). In PEG orders, the ratio of the median time-to-fill for passive to aggressive orders is 2.5 (20.21 seconds compared to 8.12 seconds), and the ratio of time-to-cancel is 14.3 (11.63 compared to 0.79 seconds). These simple statistics are sufficient to indicate a clear price-time tradeoff in the data. We also assess the effect of submission price, size, and display on the time to execution in a multivariate setting using the accelerated failure time limit order model of Lo, MacKinlay and Zhang (2002). The key conclusion from these models is precisely as above: submission price and display decisions are the

¹⁵ A single child order can result in multiple fills depending on the order type, matching protocols, and reporting conventions of the trading venue. For example, a child order of 100 shares, could result in two fills of 40 shares and 60 shares. For simplicity, we focus on the time-to-fill for the first fill.

most meaningful when considering time-to-fill, even after controlling for market conditions. For readers interested in parameter estimates, the full results are in Appendix Table A1.

5.3 Child Order Price Impact

5.3.1 Measurement

We measure the price impact of a child order j as

$$cpi_{jt\tau} = q_{jt}(m_{j,t+\tau} - m_{jt}) / m_{jt} \quad (1)$$

where q_{jt} is equal to +1 for buys and -1 for sells and short sales, m_{jt} is the prevailing quote midpoint, and $m_{j,t+\tau}$ is the quote midpoint at some subsequent time τ . Following Conrad and Wahal (2020), we use a variety of short horizons for τ , corresponding to 100 milliseconds, 500 milliseconds, 1 second, 5 seconds, and 10 seconds.

Ideally, we would measure price movements for an order after its submission as well as its execution (for the same order). This is complicated by the fact that the time between submission and execution is sometimes so short that post-trade price movements conflate post-submission price movements. For example, the median time-to-fill for aggressively priced day orders is 12 milliseconds (Table 4), which implies that even at the shortest horizon ($\tau=100$ milliseconds), post-submission and post-trade price impact overlap by 90 milliseconds. To disentangle the two, we calculate price impact separately for child orders that are unfilled and filled. Although this is ex post, it allows us to cleanly estimate post-submission and post-execution price impact. The time subscript t in equation 4 corresponds to the submission time for unfilled orders and execution time for filled orders.

To aggregate, we calculate dollar-weighted average price impact for each day, and then average across days. In addition to averages, we also calculate the standard deviation, the percentage of child orders with zero price impact, and the percentage of orders with positive price impact for all orders in a day. We use this daily approach rather than averaging across the entire sample because we explore the time series variation in price impact later in the paper.

5.3.2 The Distribution of Price Impact

Panel A of Table 5 shows time series averages of price impact for various price aggressiveness categories, separately for unfilled and filled orders across all horizons. Standard errors based on the daily time series are about $1/20^{\text{th}}$ of the mean so we do not report them. For unfilled orders, average price impact is positive and declines steadily with price aggressiveness. At the one second horizon, for example, the price impact of aggressive limit orders is 1.28 basis points, declining monotonically to 0.18 basis points for passive orders. These results imply price discovery occurring through the submission of limit orders, irrespective of execution.¹⁶ A comparison of point estimates is interesting. Brogaard, Hendershott, and Riordan (2019) estimate the price impact of limit orders from 15 securities traded on the TSX that are not cross-listed in the US over a 9-month period between October 2012 and June 2013. Their estimate of price impact over a 10-second horizon is 0.69 basis points (Table 6 in their paper). At the same horizon, our data show the price impact of aggressive and passive orders are 2.03 and 0.84 basis points respectively. Interestingly, the average price impact of PEG orders show in Panel B are lower, although still positive. At the same 10 second horizon, aggressive PEG orders incur a price impact of 1.06 basis points, roughly half that of aggressive limit orders. Similarly, passive limit orders have a price impact of 0.31 basis points, compared to 0.84 basis points for limit orders.

The second, third, and fourth blocks in Panels A and B show the time series average of the daily standard deviation of dollar weighted price impact, and the time series average of the percentage of orders with zero and positive price impact. Focusing again on limit orders at the 10 second horizon, the average standard deviation for aggressive and passive orders is 5.02 and 3.37 basis points respectively, more than twice the mean. Similarly, the percentage of orders with zero price impact is 54.67 percent for aggressive orders and 61.64 percent for passive orders. Thus variation in price impact is large and we examine it in more detail in section 5.3.3.

Table 5 also contains equivalent statistics for filled child orders, where price impact is measured from the time of the fill (not submission). As expected, aggressive orders incur positive price impact. Prices move by 0.80 basis points at the 10 second post-trade horizon for limit orders, and 0.14 basis points for PEG orders. Passive orders, on the other hand, incur negative price

¹⁶ Spoofing strategies, regardless of whether they are successful or not, rely on the belief that order arrival moves prices.

impact, or positive adverse selection costs since prices continue to decline after buys and rise after sells. At the 10 second horizon, the average price movement is -2.01 basis points for limit orders and -1.45 basis points for PEG orders.

For a subset of orders in which the time to execution is greater than 10 seconds, we calculate post-submission and post-execution price impact. Although there is an obvious selection bias in this subsample, it allows us to measure post-submission and post-execution price movements for the same order. We do not report the results in a table, but they are easily summarized. For aggressive orders, the sample sizes are too small to be meaningful. However, for approximately 1.2 million midpoint limit orders and 2.4 million midpoint PEG orders, post-submission price impact at the 10 second horizon is 1.21 and 0.31 basis points respectively. Post-execution price impact is -2.00 and -1.61 basis points respectively, indicating adverse selection. For 2.5 million passive limit orders and 3.0 million passive PEG orders, the post-submission price impacts are considerably smaller (-0.07 and 0.22 basis points respectively), but the post-execution price movements are larger (-3.36 and -2.78 basis points respectively).

Some readers may worry that the price movements we detect are part of the normal market dynamic, as opposed to being causally related to child order submission. We do not believe this to be the case because of the precision of time stamps. Nonetheless, we perform a simple placebo test to rule out this alternative. For all security-date pairs in our sample, we generate 1 million draws with replacement of a random time t . For these random draws, we then calculate price impact using the same values of τ . The average price impact in this sample is zero (to the 4th decimal place), implying that the positive price impact that we observe in our data is conditional on an order arrival event.

5.3.3 Variation in Price Impact

We examine variation in price impact with respect to order choice attributes using a triple sort procedure as well as cross-sectional regressions. On each day, we independently sort limit orders into three price aggressiveness categories, scaled order size quartiles, and display groups.¹⁷

¹⁷ For limit orders, we restrict our attention to aggressive, midpoint and passive orders, ignoring orders between quote boundaries and the midpoint. Since the sorts are done each day, the sample size of the latter group can be quite small.

This 3x4x2 sorting procedure generates 24 groups. The equivalent sorts for PEG orders only use price aggressiveness categories and size quartiles because all PEG orders are hidden. For each group, we calculate average price impact within a day. Table 6 shows time series averages of the daily group averages, separately for unfilled and filled orders.

Holding price aggressiveness and display constant, the price impact of unfilled limit orders (Panel A) rises monotonically with scaled order size. For displayed aggressive orders, average price impact at the 10-second horizon for order size quartiles Q1 through Q4 are 0.88, 1.14, 1.48, and 2.80 basis points respectively. Even in displayed passive orders, average price impact rises from 0.48 basis points for Q1 to 1.20 basis points for Q4. The level of price impact is lower for non-displayed limit orders, but the pattern across order size quartiles is the same: for aggressive orders, average price impact rises from 0.74 basis points in Q1 to 2.30 basis points in Q4. For passive orders, the equivalent increase is from 0.25 to 0.98 basis points. Interestingly, in aggressive PEG orders (Panel B), the levels of average price impact are lower than for non-displayed limit orders, but the increase is still monotonic in size quartiles, from 0.56 (Q1) to 1.32 basis points (Q4). In contrast, order size is unrelated to price impact for passive PEG orders. Holding price aggressiveness and order size constant, unfilled displayed orders have systematically higher price impact. In aggressive orders in the highest size quartile, the difference between non-displayed and displayed orders is 0.50 basis points (2.30 versus 2.80 basis points). For passive orders in the largest size quartile, price impact rises from 0.98 to 1.20 basis points. These results indicate that variation in submission price, scaled order size, and display is systematically related to post-submission price impact.

A similar pattern emerges with post-fill variation in price impact. Holding price aggressiveness and order size constant, displayed orders have larger absolute price impact than non-displayed orders. Within price aggressiveness categories, however, the effect of order size and the sign of the price impact changes. For aggressive limit orders, holding display constant, average post-fill price impact is positive and increases with order size. Within the other two price aggressiveness categories average price impact is negative, but conditional on display, order size plays less of a role. For passive displayed orders, for example, the average price impact for order

size quartiles 1 through 4 are -2.03, -1.93, -2.02, -2.18 basis points respectively. There is similar lack of variation in price impact across order size quartiles in non-displayed passive orders.

We also estimate daily price impact regressions separately for unfilled and filled orders. The regressions have two advantages relative to triple sorts. They allow us to control for other covariates of interest (such as asymmetry in the depth of the limit order book) and examine the time series of coefficients. Table 7 presents average parameter estimates from these regressions. Standard errors appear in parentheses, based on the time series of coefficients. The average sample size for regressions with unfilled limit and PEG orders is greater than 60,000 orders. For filled orders limit and PEG orders, the sample sizes are 45,125 and 12,228 respectively.

The parameter estimates from the regressions confirm the evidence from triple sorts: order price aggressiveness, order display, and order size significantly impact post-submission and post-execution price movements. Price aggressiveness has the largest effect, with estimates decreasing in an almost monotonic fashion across the five price aggressiveness bins for both filled and unfilled orders. Holding all else constant, displaying an order is associated with 0.2 basis points additional price impact. A one standard deviation increase in scaled order size is associated with a 0.15 basis points increase in price impact.

We also examine time series variation in price impact. Our interest is driven by the potential influence of market-wide variation in risk bearing capacity (Adrian and Shin (2010) and Brunnermeier and Pedersen (2009)). We define low, medium and high market risk periods based on 33rd and 67th percentiles of the daily distribution of the VIX over our sample period. We then average the time series of coefficients from the regressions in Table 7, and also average the univariate price impact statistics in Table 6 over these periods. We do not present the results in a table but can report that there is a monotonic relation between average price impact for both unfilled and filled orders and VIX levels. For example, the average price impact for unfilled aggressive limit orders measured at $\tau=10$ is 1.84, 2.13, and 2.67 basis points in low, medium and high VIX periods, respectively. For passive limit orders, the equivalent averages are 0.67, 0.95 and 1.16 basis points respectively. There is a similar pattern in PEG orders. In filled orders, there is again a monotonic relation between post-trade price impact and the level of the VIX. These

results imply that price impact at the child level, and therefore parent-level trading costs, are influenced by market-risk, well outside the scope of a trading algorithm.

5.4 Strategic Runs

Each child order is generated in pursuit of a common goal (i.e. to fill the desired demand) from a shared codebase, cognizant of the tradeoff between time-to-fill and price impact. In this section, we provide an exploratory analysis of the dependence structure that arises from strategic behavior as an algorithm navigates the tradeoff between the need to trade and the cost of trading.

We define a “run” as a sequence of consecutive child orders emanating from a parent within a particular price aggressiveness category. We simplify the analysis by collapsing the five price aggressiveness categories to three groups: passive, aggressive, and inside the spread. The last group includes child orders posted at the midpoint, between the far side of the spread and the midpoint, and between the near side of the spread and the midpoint. We restrict the analysis to parent orders that seek to trade at least 1 basis point of average daily volume and with at least 50 child orders (referred to as “Large Orders” in Table 1). This ensures that the analysis is not driven by small parent orders in which the notion of a run is less economically meaningful.

Panel A of Table 8 shows summary statistics of the restricted sample, which consists of 812,132 parent orders. On average, parent orders in this sample contain 63 runs. Each run includes nine consecutive child orders in the same aggressiveness class, lasts 567 seconds, and submits over 1,500 shares. The percentage of runs that are passive, aggressive, and inside the spread are 45, 16, and 38 percent respectively. There is variation in the distribution of these features across algorithms, reflecting the cost-volume tradeoff in Figure 2. For instance, Algorithm A, with the highest cost and volume sensitivity, has the highest number of child orders per run, the shortest run duration, and the largest run volume (number of shares submitted in the run). In contrast, Algorithms B and D, which are more sensitive trading costs, have the lowest percentage of aggressively priced orders per run.

Panel B shows transition matrices between successive passive, aggressive, and inside runs. Across all algorithms, the probability that an algorithm in a passive run at $t-1$ moves to an aggressive state in t is 31.2 percent. The equivalent transition probability from an aggressive run

to a passive run is 29.3. Despite the fact that each algorithm is quite different in terms of the percentage of runs that are passive or aggressive, the transitions probabilities between extreme phases (passive to aggressive, and vice versa) are quite similar and symmetric. And conditional on a run being in a passive or aggressive state in $t-1$, the probability that the subsequent run is at the other extreme is 2-5 times more than being intra-spread.

The transition matrices show the unconditional probabilities of moving from one type of run to another. Since these probabilities likely depend on price movements and whether the prior run received an execution, we also estimate conditional logistic regressions. The regressions model the probability that the run is aggressive or passive, conditional on the prior run. We use two dependent variables: (a) the dependent variable ($\text{Pr}(A_t)$) is equal to one for aggressive runs, and zero otherwise if the prior run is either passive or inside the spread, (b) the dependent variable ($\text{Pr}(P_t)$) is equal to one if the current run is passive and zero otherwise, if the prior run is aggressive. Effectively, we model the switch from taking to providing liquidity, and vice versa. We include two independent variables which capture the success of the prior run in trading and the associated price movement. Fill_{t-1} is an indicator variable equal to one if the prior run received a fill. SRet_{t-1} is a signed return, the midpoint to midpoint return from the start of the prior run to its end, multiplied by +1 for buys and -1 for sells. The signed return reflects the implicit cost incurred during the prior run.

Panel C presents the results of these regressions. The top row shows the dependent variables, $\text{Pr}(A_t)$ or $\text{Pr}(P_t)$. The row immediately below, labeled Run_{t-1} , shows whether the prior run was passive (P), intra-spread (I), or aggressive (A). Standard errors appear in parentheses. Given the large sample size, the coefficients are estimated with considerable precision so we focus on marginal effects, reported in square brackets. For Fill_{t-1} the marginal effect is the change in probability based on whether the prior run received a fill or not. For the signed return SRet_{t-1} , the change in probability is based on a 1 basis point change in the signed return.

The regressions suggest that the probability that a run switches from being aggressive to passive, or vice-versa, is unrelated to price movements during the prior run. It is possible that this lack of sensitivity is because average price movements are quite small and therefore non-binding for the switching mechanism. But it also implies that these runs do not lean against the wind, at

least over these short horizons. In contrast, the probability of switches is significantly related to whether the prior run results in a fill. In Algorithm A, for example, if the prior run was passive and received a fill, the probability that the subsequent run is aggressive rises by 7.59 percent. In the Li, Wang, and Ye (2019) model this sensitivity is due to the inelastic need to trade. Our results emphasize that this inelasticity is important, but we note that asymmetry and heterogeneity are also important – there are differences in the changes in probabilities between taking and providing liquidity both within and across different algorithms. For Algorithm D, a passive run receiving a fill has a minuscule effect on the probability that the subsequent run is aggressive (0.8 percent). However, for the same algorithm, if the fill is received in an aggressive run, the probability that the subsequent run is passive rises by 11.21 percent. This type of asymmetry is different across algorithms, which we interpret as differential sensitivity to the tradeoff between the need to trade and the cost of trading. In other words, the probability of “switches” is very much algorithm related and hearkens back to volume and cost sensitivity in Figure 2.

6. Linking Algorithmic Choices to Parent Order Costs

We turn our attention to the influence of child and run-level choices on parent-level trading costs. The linkage between the two is not innocuous. The purpose of trading algorithms is to slice and dice parent orders to hide parent size and directional trading intentions; order breakup adds noise so that observable child-level characteristics do not perfectly reveal parents. Studying this attempted obfuscation requires us to be able to observe all child orders, not just those that are executed. It is perhaps because of the observability of intent (parent orders), process (submitted child orders), and outcomes (executions) that, to our knowledge, we are the first to study this link.

The early literature on institutional trading costs employs data from SEI, the Plexus Group, or more recently Ancerno, all of which are data aggregators from multiple institutions (for examples, see Chan and Lakonishok (1995), Keim and Madhavan (1995), and Hu et al. (2018) respectively). Notably, parent and child orders are not observable in these data sources since they are collated using end-of-day trade tickets. This literature is largely devoted to understanding the influence of market structure and security-specific characteristics on trading costs. More recently, Frazzini, Israel, and Moskowitz (2018) examine parent order costs from one large asset manager

that employs trading algorithms, observing parent orders and child executions. Their interest is in calibrating trading costs models and on market structure. The methodological approach in these papers is to estimate regressions of parent-level trading costs on security characteristics and market conditions. We adapt this empirical strategy to suit our purpose. As control variables, we use parent order size (scaled by average daily volume), the market capitalization of the firm, the VIX index, lagged idiosyncratic volatility measured using market model residuals over the prior year, and indicator variables that reflect the presence of price or volume constraints. The novelty in our regressions is, of course, the ability to incorporate key aspects of child orders. Our primary interest is in the price, time, and display choices of all child orders. We use three variables to capture each of these tradeoffs: (a) the percentage of child orders with aggressive submission prices, (b) the percentage of child orders that are day orders, and (c) the percentage of child orders that are displayed. Because the variables are specified as percentages, one minus the variable reflects order passivity, IOC orders and non-displayed orders. The sample is the same as that for strategic runs: parent orders attempting to trade at least one basis point of average daily volume with at least 50 child orders. This ensures that intra-parent information is meaningfully reflected in the regressions.

Table 9 reports parameter estimates with standard errors (clustered by day) in parentheses. The control variables in the regressions take on signs consistent with the existing literature: parent-level costs are positively correlated with parent order size, idiosyncratic volatility and market volatility, and negatively correlated with firm size. In terms of magnitude, parent order size is of particular interest. The point estimates in Table 9 indicate that a one standard deviation increase in parent order size increases expected trading costs by about 33, 31, 35 and 38 percent for algorithms A, B, C and D respectively. These estimates are in the same range as those reported by Frazzini, Israel, and Moskowitz (2018).

In Panel A, we add the percentage of aggressively priced child orders per parent as an explanatory variable. For all algorithms, child-level price aggressiveness is positively related to parent-level trading costs. For algorithms A, B, and D, a one standard deviation increase in price aggressiveness is associated with about a 10 percent increase in trading costs (reported in the row labelled “Impact (%)” and roughly corresponding to one basis point). For Algorithm C, which has

the highest sensitivity to volume but the lowest sensitivity to trading costs, the increase is much larger, 23 percent (1.83 basis points).

In Panel B, we add the percentage of day orders to the baseline regression, recalling that the other major category for time-in-force is IOC orders. For three of the four algorithms (A, C, and D), the effect is negative, indicating that the use of day orders reduces parent-level trading costs. In algorithm B, the coefficient is indistinguishable from zero. The magnitudes are particularly large for algorithms A and C (50 and 72 percent reductions, corresponding to 2.75 and 5.81 basis points respectively), implying that the use of more patient orders can substantially lower parent-level trading costs. In Panel C, we include the percentage of displayed orders as an explanatory variable. The effect of displayed orders is algorithm dependent. In algorithms B and D, the use of displayed orders reduces trading costs by 32 and 17 percent respectively. In contrast, for algorithms A and C, displayed orders increase trading costs by 30 and 32 percent. These differences correlate with the cost and volume sensitivity of the respective algorithms; B and D are more cost sensitive than A and C (Figure 2), and as is evident from Table 2, are more likely to use passive, displayed orders.

Since child orders are linked together, it is also interesting to examine the role of strategic runs on parent-level trading costs. In Table 10, we estimate equivalent regressions but add run-level information. If orders that are more difficult to trade require more runs that oscillate between taking and providing liquidity, then the number of runs should be positively related to parent-level trading costs. Consistent with this, the coefficient on the (logarithm of) the number of runs is positive across all four algorithms. We use two variables to capture the effect of price aggressiveness within runs. In Panel A, we add the proportion of traded volume generated by aggressive runs (aggressive run volume), and in Panel B we include the fraction of time the parent order employs aggressively priced runs (aggressive run duration). Both measures are positively related to parent trading costs. In algorithms A and C (which are more sensitive to volume), a one standard deviation increase in aggressive run volume increases parent trading cost by 23 percent. Similarly, a one standard deviation increase in aggressive run duration increases parent trading costs by 13 and 21 percent respectively. In contrast, in algorithms B and D, the effect of aggressive

run volume is substantially attenuated, only 9 and 8 percent respectively. In algorithm B, the duration of aggressive runs has no discernible effect on trading costs.

We explore alternative constructs of the above variables to assess their sensitivity. For instance, the price aggressiveness of child orders can be computed in dollar terms instead of the number of child orders, or instead of aggressiveness, one could use the frequency of passive submissions. One could also use the percentage of child orders that add or remove liquidity (as specified in FIX Tag 851). None of these alternative formulations makes a difference to inferences and we do not report them in the interest of brevity. Overall, the regressions suggest two conclusions. First, algorithmic behavior at the child- or run-level tradeoffs influences parent-level trading costs. In particular, there is no escaping the price impact of aggressive pricing. Second, the influence of child-level choices varies across algorithms, suggesting that design features endogenously correspond closely to algorithmic objectives and latent economic tradeoffs.

7. Conclusion

We study the anatomy of four trading algorithms widely used in global equity markets. Parent orders from these algorithms spawn hundreds of child orders that attempt to balance transaction costs with execution likelihood. They do so via complex order features that endeavor to optimize price-time-display priority rules in fragmented markets. Our evidence suggests that while the basic nature of trading – the tradeoff between the need to trade and the cost of trading – remains unchanged, the manner and characteristics of trade in modern equity markets are quite different. Price aggressiveness, time-in-force, and display instructions are key levers that trading algorithms use in pursuit of best execution. The differences in the usage of these levers across algorithms presents opportunities for traders to appropriately express their desire to trade.

The algorithms make extensive use of passive limit or PEG orders to minimize transaction costs by not paying the quoted bid-ask spread, but increase exposure to execution risk and adverse selection in the process. Child orders move prices, not only conditional on execution but also at the time of submission. This implies that price discovery also takes place via the revelation of trading interest in addition to actual trading. The extent to which child order submissions move prices largely depends on the price aggressiveness of the order, whether it is displayed on the order

book and the VIX. Child orders are not independent, generating strategic runs in which algorithms maintain their presence on the same side of the spread. These runs oscillate between providing and taking liquidity, effectively competing with high frequency market-making firms. This movement between liquidity provision and extraction is reflective of the underlying tradeoff between execution risk and transaction costs.

The breakup of parent orders into hundreds of child orders ultimately represents algorithmic attempts at best execution, an effort to deliver desired quantities at the lowest possible cost. Notwithstanding attempts to hide trading intentions, the price, time, and display choices in child orders aggregate to parent-level trading costs in quantitatively meaningful terms.

References

- Adrian, T. and Shin, H.S., 2010. Liquidity and leverage. *Journal of Financial Intermediation* 19, 418-437.
- Almgren, R.F. and N. Chriss, 2000. Optimal execution of portfolio transactions. *Journal of Risk* 18, 57-62.
- Aquilina, M., E. Budish, P. O'Neill, 2020. Quantifying the high frequency trading “arms race”: A simple new methodology and estimates. Working paper, University of Chicago.
- Bertsimas, D. and Lo, A.W., 1998. Optimal control of execution costs. *Journal of Financial Markets* 1, 1-50.
- Bessembinder, H., A. Carrion, L. Tuttle, K. Venkataraman, 2016. Liquidity, resiliency and market quality around predictable trades: Theory and evidence. *Journal of Financial Economics* 121, 142-166.
- Biais, B., Hillion, P., and Spatt, C., 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655-1689.
- Biais, B., Martimort, D. and Rochet, J.C., 2000. Competing mechanisms in a common value environment. *Econometrica* 68, 799-837.
- Biais, B. and Woolley, P., 2011. High frequency trading. *Manuscript, Toulouse University, IDEI*.
- Bloomfield, R., O'Hara, M. and Saar, G., 2005. The “make or take” decision in an electronic market: Evidence on the evolution of liquidity. *Journal of Financial Economics* 75, 165-199.
- Boehmer, E., Jones, C.M., and Zhang, X., 2017, Tracking retail investor activity. Working paper, Columbia Business School.
- Brogaard, Jonathan, Terrance Hendershott and Ryan Riordan, 2019, Price discovery without trading: Evidence from limit orders, *Journal of Finance* 74, 1621-1658.
- Brolley, M. and Malinova, K., 2020, Informed liquidity provision in a limit order market, *Journal of Financial Markets*, forthcoming.
- Brunnermeier, M.K., and L.H. Pedersen, 2005, Predatory trading, *Journal of Finance* 60, 1825-1863.
- Budish, E., Cramton, P. and Shim, J., 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics* 130, 1547-1621.

- Chan, L., and J. Lakonishok, 1995, The behavior of stock prices around institutional trades, *Journal of Finance* 50, 1147-1174.
- Collin- Dufresne, P. and Fos, V., 2015. Do prices reveal the presence of informed trading? *Journal of Finance* 70, 1555-1582.
- Conrad, J. and Wahal, S., 2020. The term structure of liquidity provision. *Journal of Financial Economics* 135, 239-259.
- Duffie, D. and Zhu, H., 2017. Size discovery. *Review of Financial Studies* 30, 1095-1150.
- Foucault, T., Kadan, O. and Kandel, E., 2005. Limit order book as a market for liquidity. *Review of Financial Studies* 18, 1171-1217.
- Frazzini, A., Israel, R., and Moskowitz, T., 2018. Trading costs. *SSRN*
- Glosten, L.R. and Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 71-100.
- Glosten, L.R., 1994. Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127-1161.
- Grossman, S.J. and Miller, M.H., 1988. Liquidity and market structure. *Journal of Finance* 43, 617-633.
- Harris, L. and Hasbrouck, J., 1996. Market vs. limit orders: the SuperDOT evidence on order submission strategy. *Journal of Financial and Quantitative Analysis* 31, 213-231.
- Holden, C.W. and Jacobsen, S., 2014. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance* 69, 1747-1785.
- Hu, G., K.M. Jo, Y. A. Wang, J. Xie, 2018, Institutional trading and Abel Noser data. *Journal of Corporate Finance* 52, 143-167.
- Kaniel, R. and Liu, H., 2006. So what orders do informed traders use? *Journal of Business* 79, 1867-1913.
- Keim, D.B. and Madhavan, A., 1995. Anatomy of the trading process empirical evidence on the behavior of institutional traders. *Journal of Financial Economics* 37, 371-398.
- Kelley, E., and Tetlock, P., 2013. How wise are crowds? Insights from retail orders and stock returns. *Journal of Finance* 68, 1229-1265.
- Korajczyk, R. and Murphy, D., 2019, High frequency market-making to large institutional trades. *Review of Financial Studies* 32, 1034-1067.

- Large, J., 2009. A market-clearing role for inefficiency on a limit order book. *Journal of Financial Economics* 91, 102-117.
- Li, S., Wang, X., and Ye, M., 2019. Who Provides Liquidity and When. *SSRN*
- Lo, A.W., MacKinlay, A.C. and Zhang, J., 2002. Econometric models of limit-order executions. *Journal of Financial Economics* 65, 31-71.
- Menkveld, A.J., 2016. The economics of high-frequency trading: taking stock. *Annual Review of Economics* 8, 1-24.
- Menkveld, A.J. and Zoican, M.A., 2017. Need for speed? Exchange latency and liquidity. *Review of Financial Studies* 30, 1188-1228.
- O'Hara, Maureen, 2015, High frequency market microstructure, *Journal of Financial Economics* 116, 257-270.
- Parlour, C.A., 1998. Price dynamics in limit order markets. *The Review of Financial Studies* 11, 789-816.
- Peterson, M. and Sirri, E., 2003. Evaluation of the biases in execution cost estimation using trade and quote data. *Journal of Financial Markets* 6, 259-280.
- Riccò, Roberto, Rindi, Barbara and Seppi, Duane J., 2020. Information, Liquidity, and Dynamic Limit Order Markets. *SSRN*
- Roşu, I., 2009. A dynamic model of the limit order book. *Review of Financial Studies* 22, 4601-4641.
- Securities and Exchange Commission, 2015, Maker-Taker fees on equities exchanges.
- Seppi, D.J., 1997. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies* 10, 103-150.
- Van Kervel, V., Menkveld, A., 2019, High frequency trading around large institutional orders. *Journal of Finance* 74, 1091-1137.
- Weill, P.O., 2007. Leaning against the wind. *Review of Economic Studies* 74, 1329-1354.
- Yang, L. and Zhu, H., 2019. Back-running: Seeking and hiding fundamental information in order flows. *Review of Financial Studies*, forthcoming.
- Zhu, H., 2014. Do dark pools harm price discovery? *Review of Financial Studies* 27, 747-789.

Table 1

Algorithm usage and parent orders

Panel A reports the number of parent orders and their aggregate value in billions of dollars. Panel B reports means, medians and standard deviations of three measures of parent order size. The first is the dollar value of the parent in thousands of dollars. The second measure is the dollar value of the parent order, scaled by average daily trading volume in the security over the prior 20 days. The third measure is the dollar value of the parent order, scaled by interval volume. Interval volume is the dollar volume over the duration of the order (i.e. from the start time to the end time of the order). Panel C shows statistics for the duration of the parent order in minutes. Panel D average parent-level implicit costs (in basis point), weighted by the dollar value of the parent order. Implicit costs are computed as weighted average transaction prices scaled by the prevailing midpoint before the start of the parent order. The first row in the panel shows average implicit costs for the full sample of parent orders. The second row shows implicit costs for large orders, define as parent orders with at least 50 child orders and whose order size is at least 1 basis point of average daily volume.

	All	Algo A			Algo B			Algo C			Algo D		
		Buy	Sell	Short Sale	Buy	Sell	Short Sale	Buy	Sell	Short Sale	Buy	Sell	Short Sale
Panel A: Parent orders													
Number	2,348,390	95,345	59,427	32,159	93,576	53,637	42,520	92,576	71,817	25,637	897,978	453,661	430,116
Value (\$B)	675	23.3	17.1	8.7	33.7	20.9	15.5	65.2	65.9	16.2	196.3	136.6	75.6
Panel B: Measures of parent order size													
Value of parent (\$000)													
Mean	287	245	287	272	360	391	366	701	917	635	218	301	175
Median	17	24	28	23	61	70	64	192	210	183	10	12	8
Std. Dev.	1,337	987	1,068	984	1,516	1,768	1,174	2,022	2,795	1,586	1,088	1,570	862
Parent volume / Average prior daily volume (%)													
Mean	0.38	0.24	0.27	0.23	0.72	0.83	0.59	0.72	0.77	0.60	0.31	0.43	0.23
Median	0.01	0.04	0.04	0.04	0.06	0.07	0.05	0.17	0.17	0.17	0.01	0.01	0.00
Std. Dev.	2.48	1.00	1.25	1.12	4.69	5.81	2.14	2.27	2.50	1.37	2.02	3.25	1.21
Parent volume / Interval volume (%)													
Mean	4.80	14.04	14.68	12.04	3.76	5.04	3.28	12.05	12.05	10.43	3.13	3.48	2.96
Median	0.24	6.03	6.25	4.87	0.60	0.81	0.59	9.83	10.00	8.08	0.11	0.13	0.10
Std. Dev.	14.13	22.79	23.15	21.17	9.56	12.16	8.26	11.66	11.06	11.38	12.86	13.07	13.03
Panel C: Parent duration (minutes)													
Mean	84.19	18.52	17.88	20.36	93.55	96.46	81.81	31.18	32.53	36.69	95.34	104.06	88.04
Median	18.89	3.41	3.25	3.18	50.45	45.45	44.76	8.73	8.19	12.45	23.41	28.81	16.87
Std. Dev.	121.96	45.52	44.64	51.98	106.07	112.12	93.90	59.68	64.53	61.00	129.37	134.44	126.18
Panel D: Weighted average implicit costs (basis points)													
All Orders	7.33	6.92	4.95	9.02	8.13	5.35	7.97	5.22	2.34	7.24	9.51	6.99	9.00
Large Order	9.04	8.34	5.82	11.14	9.83	6.51	9.56	6.94	3.28	8.77	11.49	8.40	11.09

Table 2

Child order characteristics and fills

Each parent order generates a sequence of child orders. Child orders are unfilled (and therefore cancelled), or filled. Panel A shows the average number of child orders and fills per parent, and the average dollar value of child orders (in \$000s) and fills per parent. Panel B shows order characteristics. For each parent order, we compute the percentage of child orders with a particular characteristic based on dollar values. We then report the average percentage across parent orders in a group. For example, we calculate the percentage of child orders (by dollar volume) that are market, limit or PEG orders. Indented order percentages corresponding to day orders, immediate or cancel (IOC) orders, orders that add liquidity, and orders that remove liquidity are based on limit or PEG order subgroups. Day and IOC orders are identified based on FIX tag 59. Add versus remove liquidity indications are based on FIX tag 851. For some characteristics, we only report one category of a mutually exclusive group so that the omitted category can be inferred. The omitted category for displayed orders are non-displayed orders, the omitted category for orders directed to Lit venues is dark venues. Aggressive orders are those priced at or above the far side of the prevailing NBBO (buy orders priced at the best ask or higher, sell orders priced at the best bid or lower).

	All	Algo A			Algo B			Algo C			Algo D		
		Buy	Sell	Short Sale	Buy	Sell	Short Sale	Buy	Sell	Short Sale	Buy	Sell	Short Sale
Panel A: Distribution of child orders and fills per parent													
Number of child orders	126	100	107	102	145	150	142	139	143	145	122	149	106
Number of fills	38	18	18	20	51	51	50	37	39	37	39	45	34
\$ Value of child orders	918	1,042	1,498	964	1,524	1,405	1,626	2,150	2,056	2,951	638	985	478
\$ Value of fills	165	131	152	141	241	246	245	330	368	368	135	178	108
Panel B: Selected characteristics of child orders													
					Order type								
Market	0.38	0.5	1.6	0.0	0.1	0.2	0.0	1.4	1.3	3.1	0.2	0.3	0.2
Limit	81.48	38.5	36.5	37.5	78.3	76.9	77.7	45.8	47.5	43.9	90.1	88.6	91.9
Limit: Day	61.9	40.8	40.1	41.8	69.5	69.2	69.8	52.1	50.8	52.1	64.5	62.9	64.8
Limit: IOC	38.1	59.2	59.9	58.2	30.5	30.8	30.2	47.9	49.2	47.9	35.5	37.1	35.2
Limit: Add Liq.	45.7	26.1	26.4	27.4	56.3	56.3	56.7	40.4	39.9	39.3	47.0	45.8	47.6
Limit: Remove Liq.	54.3	73.9	73.6	72.6	43.7	43.7	43.3	59.6	60.1	60.7	53.0	54.2	52.4
PEG	18.12	61.0	61.9	62.4	21.6	22.9	22.2	52.6	51.2	53.0	9.5	11.1	7.9
PEG: Day	34.5	43.1	43.7	50.1	39.7	40.1	40.6	47.4	46.4	48.8	24.5	25.5	26
PEG: IOC	65.5	56.9	56.3	49.9	60.3	59.9	59.4	52.6	53.6	51.2	75.5	74.5	74
PEG: Add Liq.	17.0	10.6	13.4	10.6	22.0	22.6	23.3	24.6	24.2	28.1	14.0	14.5	15.4
PEG: Remove Liq.	83.0	89.4	86.6	89.4	78.0	77.4	76.7	75.4	75.8	71.9	86.0	85.5	84.6
					Order characteristics								
Aggressive	42.61	67.9	64.4	61.1	38.8	39.1	39.3	70.9	70.9	73.9	37.1	38.8	36.9
Displayed	75.63	28.0	27.2	25.6	71.1	69.2	70.0	34.5	35.4	34.7	85.5	83.6	87.9
Lit venues	77.95	37.3	36.5	36.7	75.0	73.2	74.2	42.2	43.0	42.9	86.4	84.8	88.0

Table 3**Scaled order size in price aggressiveness and display categories**

In Panel A, child limit orders are categorized into five categories based on their price aggressiveness relative to the prevailing NBBO. A limit order is aggressive if its limit price is at the far side of the quote, ask or higher for limit buys, bid or lower for limit sells. A limit order is passive if its limit price is at the near side of the quote, bid or lower for limit buys, ask or higher for limit sells. Orders posted between these are categorized using open intervals in parentheses. PEG orders (Panel B), can only be priced at the aggressive side, midpoint, and passive side and are categorized appropriately. The table shows the 25th percentile, the median, and the 75th percentile of the size of the order scaled by the average depth at the NBBO. Depth at the NBBO is computed using all depth available in all trading venues that are at the best bid or offer, regardless of whether they are the official NBBO. Average depth is computed the simple average of total depth at the bid and ask.

	All orders			Unfilled orders			Filled orders		
	25 th Perc.	Med.	75 th Perc.	25 th Perc.	Med.	75 th Perc.	25 th Perc.	Med.	75 th Perc.
Panel A: Child limit orders									
	All orders								
Aggressive	2.75	13.3	40.0	4.1	16.6	43.2	1.3	8.6	33.3
(Agg., Midpoint)	20.0	40.0	66.6	25.0	50.0	80.0	5.6	28.0	64.0
Midpoint	1.5	11.8	33.3	8.7	25.0	50.0	0.8	6.0	23.5
(Midpoint, Pass.)	0.6	2.4	18.0	0.7	3.6	25.0	0.6	2.0	14.6
Passive	0.6	3.0	11.7	0.6	2.9	11.0	0.6	3.4	1.4
	Non-displayed orders								
Aggressive	3.5	14.5	40.0	5.4	20.0	50.0	2.5	10.0	33.3
(Agg., Midpoint)	25.0	50.0	85.7	28.5	50.0	100	22.2	40.0	66.6
Midpoint	13.3	28.5	57.1	16.6	33.3	66.6	10.7	22.2	47.0
(Midpoint, Pass.)	13.7	27.5	56.0	16.4	33.3	66.6	12.1	23.0	48.6
Passive	5.1	12.5	28.0	5.3	13.0	28.5	4.7	11.6	26.0
	Displayed orders								
Aggressive	2.2	12.5	40.0	2.8	12.5	38.5	0.9	7.8	33.3
(Agg., Midpoint)	16.6	40.0	66.6	22.2	40.0	70.0	4.0	25.0	58.0
Midpoint	1.0	7.2	28.5	4.2	22.2	50.0	0.6	3.7	19.4
(Midpoint, Pass.)	0.6	1.3	12.0	0.6	1.9	15.0	0.5	1.3	10.5
Passive	0.5	2.5	10.0	0.5	2.4	9.0	0.5	3.0	13.0
Panel B: Child PEG orders									
Aggressive	5.8	20.0	50.0	6.2	22.2	50.0	3.4	13.3	40.0
Midpoint	4.8	18.1	50.0	5.5	20.0	50.0	2.8	11.4	33.3
Passive	2.5	9.0	25.0	2.7	10.0	28.5	2.1	6.4	16.6

Table 4**Frequency, time-to-cancel, and time-to-fill of day child orders**

The sample includes all non-IOC (immediate or cancel) child orders, as identified by FIX tag 59, over 99 percent of which are ‘day’ orders. In Panel A, child limit orders are categorized into five categories based on their price aggressiveness relative to the prevailing NBBO. A limit order is aggressive if its limit price is at the far side of the quote, ask for limit buys, bid for limit sells. A limit order is passive if its limit price is at the near side of the quote, bid for limit buys, ask for limit sells. Orders posted between these are categorized using open intervals in parentheses. Orders are non-displayed if they are posted on dark pools, or designated non-displayed at exchanges. PEG orders (Panel B), can only be priced at the far side, midpoint, and near side and are categorized appropriately. The table shows the 25th percentile, the median, and the 75th percentile of the time-to-cancellation of unfilled child orders (in seconds), of the time-to-fill for filled child orders (in seconds). “N” is the number of child orders in each category, in millions.

	All orders	Unfilled child orders: time-to-cancel				Filled child orders: time-to-fill			
	N	25 th Perc.	Med.	75 th Perc.	N	25 th Perc.	Med.	75 th Perc.	N
Panel A: Child limit orders									
		All orders							
Aggressive	1.4M	4.19	14.77	47.14	1.1M	0.01	0.12	10.53	0.2M
(Agg., Midpoint)	0.2M	0.04	5.06	31.76	0.1M	0.01	0.36	10.16	0.1M
Midpoint	3.9M	15.66	44.09	75.15	0.9M	0.73	5.49	20.88	3.0M
(Midpoint, Pass.)	8.3M	29.88	59.39	103.07	3.1M	3.88	15.38	43.33	5.1M
Passive	106.8M	29.50	45.00	84.79	73.4M	8.90	27.78	74.00	33.3M
		Non-displayed orders							
Aggressive	0.9M	4.17	14.31	45.67	0.9M	0.01	0.48	7.88	0.02M
(Agg., Midpoint)	0.05M	6.88	18.26	42.64	0.03M	0.44	4.53	18.86	0.01M
Midpoint	0.6M	13.51	43.69	77.50	0.2M	2.76	9.70	27.40	0.4M
(Midpoint, Pass.)	0.9M	26.80	59.27	89.75	0.4M	6.47	19.08	44.84	0.4M
Passive	8.9M	17.03	51.30	77.76	6.5M	6.59	19.76	46.91	2.3M
		Displayed orders							
Aggressive	0.4M	4.26	17.46	51.69	0.2M	0.01	0.10	11.20	0.2M
(Agg., Midpoint)	0.2M	0.02	0.25	18.54	0.06M	0.01	0.30	9.62	0.1M
Midpoint	3.2M	17.81	44.13	75.00	0.6M	0.58	4.88	19.73	2.6M
(Midpoint, Pass.)	7.3M	29.92	59.40	104.79	2.7M	3.65	14.99	43.13	4.6M
Passive	97.8M	29.56	44.95	85.84	66.8M	9.15	28.53	77.09	31.0M
Panel B: Child PEG orders									
Aggressive	3.8M	0.34	8.12	26.66	3.1M	0.02	0.79	9.23	0.7M
Midpoint	36.9M	0.10	2.81	23.22	30.6M	0.71	5.32	19.40	6.2M
Passive	28.5M	5.12	20.21	44.63	23.0M	3.84	11.63	28.89	5.5M

Table 5

Price impact of unfilled and filled child orders

We calculate the price impact of each child order as $cpi_{j\tau} = q_{j\tau}(m_{j,t+\tau} - m_{j\tau})/m_{j\tau}$ where m_t is the prevailing quote midpoint, $q_{j\tau}$ is the an indicator equal to +1 (-1) for buys (sells), and τ takes on values from 100 milliseconds to 10 seconds. For unfilled child orders, t is the start time of the child order. For filled child orders, t is the fill time; price impact measures the post-trade movement in quotes. We calculate dollar-weighted average price impact, standard deviation, the percentage of child orders with zero price impact and the percentage of orders with positive price impact for all orders in a day. The table shows time series averages of each of these variables. Panels A and B show these statistics for limit and PEG orders respectively. All numbers are in basis points.

	Unfilled child orders					Filled child orders				
	0.1	0.5	1	5	10	0.1	0.5	1	5	10
Panel A: Limit orders										
	Mean									
Aggressive	1.19	1.25	1.28	1.66	2.03	0.41	0.41	0.43	0.64	0.80
(Agg., Mid.)	0.57	0.67	0.71	1.25	1.72	-0.12	-0.15	-0.16	0.09	0.26
Midpoint	0.30	0.33	0.36	0.73	1.06	-0.43	-0.49	-0.55	-0.75	-0.83
(Mid., Pass.)	0.67	0.77	0.85	1.43	1.97	-0.91	-1.08	-1.27	-1.75	-2.02
Passive	0.08	0.14	0.18	0.51	0.84	-1.04	-1.23	-1.44	-1.82	-2.01
	Std. Dev.									
Aggressive	2.05	2.28	2.47	3.80	5.02	1.40	1.67	1.91	3.17	4.24
(Agg., Mid.)	1.67	2.08	2.35	3.95	5.29	1.96	2.41	2.73	4.34	5.62
Midpoint	0.97	1.21	1.42	2.69	3.72	1.64	1.92	2.15	3.15	4.00
(Mid., Pass.)	1.47	1.82	2.07	3.51	4.69	2.39	2.81	3.13	4.32	5.29
Passive	0.62	0.96	1.20	2.40	3.37	2.46	2.85	3.15	4.16	4.95
	Average percentage of orders with zero price impact									
Aggressive	77.33	74.88	72.90	62.23	54.67	80.00	76.89	74.29	62.44	54.67
(Agg., Mid.)	81.96	76.75	72.50	51.59	40.74	65.58	58.71	53.98	37.81	29.67
Midpoint	81.22	77.57	73.53	53.31	41.79	62.33	55.08	49.79	37.46	30.29
(Mid., Pass.)	67.88	62.89	59.39	43.74	34.67	51.19	42.74	36.80	25.80	20.24
Passive	96.35	92.53	89.08	72.74	61.64	56.11	48.37	42.71	34.57	30.11
	Average percentage of order with positive price impact									
Aggressive	21.93	23.43	24.56	30.76	34.77	10.39	11.77	13.22	20.08	24.29
(Agg., Mid.)	14.24	17.02	19.23	32.38	38.75	9.74	13.16	15.37	24.63	28.83
Midpoint	17.03	18.88	21.33	35.24	42.96	7.38	11.16	12.91	19.12	23.24
(Mid., Pass.)	28.69	32.04	34.16	45.11	51.94	8.47	12.57	14.04	18.39	21.36
Passive	2.68	5.13	7.30	18.04	25.62	4.67	7.49	8.52	12.08	14.98
Panel B: PEG orders										
	Mean									
Aggressive	0.25	0.51	0.54	0.82	1.06	-0.06	-0.04	-0.05	0.05	0.14
Midpoint	0.28	0.40	0.42	0.59	0.76	-0.21	-0.25	-0.31	-0.49	-0.60
Passive	0.04	0.06	0.07	0.17	0.31	-0.79	0.89	-1.00	-1.29	-1.45
	Std. Dev.									
Aggressive	1.04	1.64	1.88	3.27	4.45	1.49	1.92	2.23	3.63	4.77
Midpoint	1.07	1.45	1.65	2.89	3.97	1.36	1.72	2.00	3.26	4.29
Passive	0.42	0.75	0.99	2.19	3.19	2.00	2.30	2.54	3.47	4.25
	Average percentage of orders with zero price impact									
Aggressive	96.14	90.98	88.32	75.28	66.39	83.69	79.31	76.35	64.37	56.69
Midpoint	92.11	86.78	83.77	69.16	59.62	84.08	80.06	77.13	65.44	57.72
Passive	97.36	93.66	90.45	75.32	65.05	68.03	62.60	58.84	48.95	42.93
	Average percentage of orders with positive price impact									
Aggressive	3.00	6.82	8.37	16.18	21.47	5.03	7.52	9.05	15.88	20.27
Midpoint	6.88	10.61	12.22	20.41	25.72	3.08	4.84	5.92	10.78	14.33
Passive	1.82	3.95	5.67	14.09	20.18	2.99	4.99	6.06	10.18	13.45

Table 6

Child order price impact across price aggressiveness, order size, and display categories

We calculate the price impact of each child order as $cpi_{jt\tau} = q_{jt}(m_{j,t+\tau} - m_{jt}) / m_{jt}$ where m_t is the prevailing quote midpoint, q_{jt} is an indicator equal to +1 (-1) for buys (sells), and τ takes on values from 100 milliseconds to 10 seconds. For unfilled child orders, t is the start time of the child order. For filled child orders, t is the fill time; price impact measures the post-trade movement in quotes. We scale the size of each order scaled by the average depth at the NBBO. Depth at the NBBO is computed using all depth available in all trading venues that are at the best bid or offer, regardless of whether they are the official NBBO. Average depth the simple average of total depth at the bid and ask. Each day, we form quartiles based on scaled order size. The table shows time series means of daily average price impact for each quartile for displayed and non-displayed orders in three price aggressiveness categories. To conserve space, we only report statistics for aggressive, midpoint and passive categories. All numbers are in basis points.

Size Quartile		Display	Unfilled orders					Filled orders				
			0.1	0.5	1	5	10	0.1	0.5	1	5	10
Panel A: Limit orders												
Aggressive	Q1	Non-Displayed	0.07	0.13	0.18	0.49	0.74	0.01	0.02	0.03	0.13	0.22
	Q1	Displayed	0.12	0.20	0.26	0.59	0.88	0.04	0.06	0.08	0.21	0.32
	Q2	Non-Displayed	0.20	0.27	0.32	0.61	0.82	0.02	0.03	0.04	0.12	0.19
	Q2	Displayed	0.21	0.29	0.36	0.79	1.14	0.12	0.15	0.18	0.36	0.49
	Q3	Non-Displayed	0.46	0.54	0.59	0.87	1.09	0.04	0.05	0.06	0.15	0.25
	Q3	Displayed	0.43	0.50	0.59	1.08	1.48	0.24	0.27	0.31	0.51	0.65
	Q4	Non-Displayed	1.40	1.52	1.55	1.97	2.30	0.26	0.26	0.27	0.44	0.58
	Q4	Displayed	1.39	1.39	1.45	2.15	2.80	0.72	0.69	0.72	1.02	1.24
Midpoint	Q1	Non-Displayed	0.23	0.26	0.28	0.62	0.88	-0.47	-0.54	-0.62	-0.79	-0.91
	Q1	Displayed	0.31	0.37	0.45	0.92	1.36	-0.48	-0.52	-0.57	-0.76	-0.80
	Q2	Non-Displayed	0.10	0.11	0.12	0.33	0.51	-0.37	-0.42	-0.52	-0.65	-0.76
	Q2	Displayed	0.31	0.36	0.42	0.84	1.21	-0.46	-0.50	-0.57	-0.73	-0.83
	Q3	Non-Displayed	0.11	0.11	0.13	0.37	0.57	-0.30	-0.37	-0.46	-0.61	-0.74
	Q3	Displayed	0.37	0.41	0.46	0.80	1.11	-0.47	-0.53	-0.63	-0.79	-0.90
	Q4	Non-Displayed	0.20	0.23	0.26	0.64	0.99	-0.43	-0.50	-0.57	-0.78	-0.88
	Q4	Displayed	0.55	0.60	0.65	1.14	1.53	-1.09	-1.24	-1.43	-2.02	-2.27
Passive	Q1	Non-Displayed	0.02	0.03	0.05	0.13	0.25	-1.11	-1.29	-1.45	-1.78	-1.96
	Q1	Displayed	0.02	0.05	0.08	0.28	0.48	-1.12	-1.31	-1.49	-1.83	-2.03
	Q2	Non-Displayed	0.03	0.03	0.04	0.15	0.29	-0.77	-0.93	-1.08	-1.36	-1.39
	Q2	Displayed	0.04	0.08	0.11	0.36	0.62	-1.07	-1.25	-1.46	-1.77	-1.93
	Q3	Non-Displayed	0.06	0.06	0.05	0.21	0.37	-0.75	-0.92	-1.08	-1.31	-1.48
	Q3	Displayed	0.07	0.12	0.17	0.48	0.81	-1.06	-1.26	-1.49	-1.84	-2.02
	Q4	Non-Displayed	0.17	0.19	0.22	0.62	0.98	-0.64	-0.78	-0.92	-1.18	-1.27
	Q4	Displayed	0.14	0.22	0.28	0.73	1.20	-1.05	-1.26	-1.50	-1.95	-2.18
Panel B: PEG orders												
Aggressive	Q1		0.03	0.10	0.14	0.38	0.56	-0.17	-0.19	-0.19	-0.12	-0.04
	Q2		0.05	0.11	0.14	0.39	0.59	-0.09	-0.10	-0.11	-0.02	0.06
	Q3		0.07	0.15	0.15	0.44	0.67	-0.08	-0.08	-0.08	0.02	0.11
	Q4		0.36	0.70	0.70	1.05	1.32	-0.02	0.02	0.01	0.10	0.21
Midpoint	Q1		0.02	0.05	0.08	0.25	0.39	-0.22	-0.26	-0.30	-0.48	-0.59
	Q2		0.06	0.11	0.14	0.29	0.43	-0.23	-0.27	-0.32	-0.47	-0.57
	Q3		0.13	0.20	0.22	0.37	0.50	-0.25	-0.29	-0.34	-0.51	-0.61
	Q4		0.38	0.53	0.55	0.73	0.93	-0.17	-0.22	-0.27	-0.48	-1.61
Passive	Q1		0.01	0.02	0.04	0.14	0.25	-1.04	-1.19	-1.33	-1.72	-1.94
	Q2		0.02	0.04	0.06	0.19	0.33	-0.83	-0.92	-1.02	-1.26	-1.41
	Q3		0.03	0.05	0.07	0.17	0.31	-0.74	-0.84	-0.93	-1.20	-1.35
	Q4		0.04	0.07	0.08	0.16	0.28	-0.77	-0.89	-1.02	-1.42	-1.65

Table 7**Average slopes from daily child order price impact regressions**

We estimate daily cross-sectional regressions of child order price impact. The dependent variable (child price impact) is measured at $\tau=5$ seconds. We use four indicator variables to assess the price aggressiveness of the order as described in Table 4. The omitted category is aggressive child orders and is represented in the intercept. The display variable is a dummy variable equal to one if an order is displayed. All orders in dark pools are undisplayed. The scaled size variable is as described in tables 3 and 6. Buy is an indicator variable equal to +1 for buys, and -1 for sales or short sales. Book asymmetry (in percent) is calculated as the depth of book at bid, minus the depth of the book at the ask, scaled by the average depth of the book. Depth is computed using all depth available in all trading venues that are at the best bid and ask, regardless of whether they are the official NBBO. $|\text{Ret}_{5,0}|$ is absolute return over the prior 5 second interval. Average “N” is the average number of orders in each daily regression. The table shows average slopes from the daily regressions, with standard errors based on the time series of coefficients.

	Limit				PEG			
	Unfilled		Filled		Unfilled		Filled	
Intercept (Aggressive)	1.215 (0.017)	1.360 (0.028)	0.388 (0.008)	0.490 (0.009)	0.581 (0.022)	0.688 (0.026)	0.093 (0.015)	0.110 (0.016)
(Agg., Mid.)	-0.634 (0.028)	-0.555 (0.031)	-0.646 (0.028)	-0.573 (0.031)	-	-	-	-
Midpoint	-0.894 (0.019)	-0.912 (0.020)	-1.404 (0.050)	-1.404 (0.032)	-0.160 (0.025)	-0.199 (0.026)	-0.557 (0.015)	-0.549 (0.016)
(Mid., Pass.)	-0.355 (0.025)	-0.403 (0.025)	-2.372 (0.025)	-2.407 (0.026)	-	-	-	-
Passive	-1.110 (0.016)	-1.220 (0.020)	-2.458 (0.008)	-2.502 (0.026)	-0.508 (0.022)	-0.585 (0.025)	-1.357 (0.019)	-1.660 (0.019)
Display	0.170 (0.010)	0.224 (0.018)	0.217 (0.008)	0.226 (0.007)	-	-	-	-
Scaled size	0.099 (0.004)	-	0.151 (0.006)	-	0.060 (0.003)	-	0.038 (0.005)	-
Book Asymmetry*Buy	-	0.002 (0.000)	-	0.001 (0.000)	-	0.003 (0.000)	-	0.001 (0.000)
Display * Scaled size	0.005 (0.001)	-	-0.037 (0.007)	-	-	-	-	-
$ \text{Ret}_{5,0} $	0.104 (0.001)	0.117 (0.014)	-0.024 (0.001)	-0.017 (0.001)	0.058 (0.001)	0.063 (0.001)	-0.042 (0.002)	-0.041 (0.002)
Buy	0.026 (0.012)	0.004 (0.014)	0.006 (0.007)	0.001 (0.008)	0.011 (0.009)	-0.002 (0.010)	0.006 (0.008)	0.015 (0.010)
Average N	63,815	63,815	45,125	45,125	66,478	66,478	12,228	12,228
Average adj-R ²	0.103	0.076	0.110	0.104	0.038	0.029	0.037	0.027

Table 8

Strategic runs of child orders

We define a strategic run as a sequence of child orders emanating from a parent order in three price aggressiveness categories: the passive side of the spread (P), inside the spread (midpoint or otherwise, I), or aggressive (far side of the spread, A)). The sample consists of parent orders with order size scaled by average daily volume of at least 1 basis point, and with at least 50 child orders. Panel A shows the number of parent orders, and the average number of runs per parent, the average number of child orders per run, the average duration of the run in seconds, the average run volume measured as the number of shares submitted by all child orders in the run, and the average percentage of runs at the three aggressiveness categories. Panel B shows transition probabilities for sequential runs between price aggressiveness categories. Panel C contains estimates from logistic regressions estimated separately for each algorithm, conditional on the price aggressiveness of the prior run. The regressions use two dependent variables. If the prior run is either passive or inside the spread, the dependent variable is equal to 1 for aggressive runs and zero otherwise. If the prior run contains aggressive child orders, the dependent variable is equal to 1 if the current run is passive, and zero otherwise. $Fill_{t-1}$ is an indicator variable equal to one if the prior run generated a fill, zero otherwise. $SRet_{t-1}$ is the midpoint to midpoint price movement from the beginning to the end of the prior run. Ret_{t-1} is multiplied by +1 for buys and -1 for sells/short sales so that it represents a cost to the algorithm. Standard errors appear in parentheses below parameter estimates. Probability changes implied by the model appear in square brackets.

Panel A: Summary statistics for strategic runs															
	All Algos			Algo A			Algo B			Algo C			Algo D		
Number of parents	812,132			50,582			82,903			87,751			590,896		
Runs per parent	63.09			53.59			51.65			41.23			68.76		
Child per run	8.84			11.92			7.45			8.16			8.87		
Run duration (sec)	566.99			102.04			452.81			158.58			683.46		
Run volume (shs)	1515.26			5503.08			3024.40			3063.60			732.23		
Percent runs															
Passive	45.58			40.58			42.78			42.00			46.94		
Inside spread	16.04			16.34			19.98			10.94			16.21		
Aggressive	38.37			43.06			37.23			47.04			36.84		

Panel B: Transition matrices between runs, probabilities in percent															
	All Algos			Algo A			Algo B			Algo C			Algo D		
	P _t	I _t	A _t	P _t	I _t	A _t	P _t	I _t	A _t	P _t	I _t	A _t	P _t	I _t	A _t
P _{t-1}	-	13.9	31.2	-	10.3	33.0	-	17.0	27.2	-	3.0	41.1	-	14.6	30.7
I _{t-1}	12.9	-	5.2	9.2	-	7.1	13.7	-	6.6	2.8	-	6.9	13.9	-	4.8
A _{t-1}	29.3	7.4	-	31.3	9.1	-	25.7	9.8	-	38.4	7.8	-	28.9	7.0	-

Panel C: Conditional logistic regressions of run price aggressiveness												
	Algo A			Algo B			Algo C			Algo D		
	Pr(A _t)	Pr(A _t)	Pr(P _t)	Pr(A _t)	Pr(A _t)	Pr(P _t)	Pr(A _t)	Pr(A _t)	Pr(P _t)	Pr(A _t)	Pr(A _t)	Pr(P _t)
Run _{t-1}	P	I	A	P	I	A	P	I	A	P	I	A
Fill _{t-1}	0.52	-0.07	0.06	0.02	-0.37	0.09	0.20	-0.02	0.36	-0.04	-1.35	0.65
	(0.00)	(0.01)	(0.05)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)
	[7.6]	[1.7]	[0.9]	[0.5]	[-8.3]	[1.7]	[1.1]	[-0.3]	[5.1]	[0.8]	[26.9]	[11.2]
SRet _t	0.01	-0.02	-0.01	0.01	-0.01	-0.00	0.01	-0.01	-0.01	0.00	-0.01	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	[0.1]	[0.4]	[-0.1]	[0.0]	[0.1]	[0.0]	[.04]	[0.0]	[-0.1]	[0.0]	[0.2]	[0.0]

Table 9**Regressions of parent order trading costs on child order price, time and display decisions**

Parent-level implicit trading costs are calculated as weighted average prices of executed child orders scaled by the prevailing midpoint before the start of the parent order. The table shows regressions of parent-level costs on the following set of regressors common to all specifications: the size of the parent order scaled by average daily volume (“parent size”), the logarithm of market value of equity (ME), an indicator variable if the parent order was subject to a volume constraint (“vol. constraint”), an indicator variable if the parent order was subject to a price constraint (“price constraint”), the level of the VIX, and idiosyncratic volatility measured using the standard deviation of residuals from a daily regression of firm-level returns on the market over the prior year. Panels A, B and C show specifications that add the percentage of aggressively priced child orders (“child agg.”), the percentage of child orders classified as good for the day (“child day”), and the percentage of child orders that are displayed (“child display”). The row labelled “Impact (%)” shows the percentage impact of a one standard deviation change in the price aggressiveness, the percentage of day orders, and the percentage of displayed orders on expected trading costs. The sample consists of all parent orders from algorithms A, B, C, and D with at least 50 child orders trading at least 1 basis point of average daily volume. Standard errors clustered by day appear in parentheses.

	Panel A				Panel B				Panel C			
	A	B	C	D	A	B	C	D	A	B	C	D
Intercept	77.02 (11.62)	137.35 (13.10)	89.04 (11.82)	167.10 (8.87)	76.26 (11.55)	144.28 (13.13)	72.23 (12.75)	169.86 (8.70)	68.12 (11.35)	148.65 (12.41)	101.35 (12.19)	182.34 (8.69)
Parent Size	94.11 (30.30)	57.93 (17.75)	90.74 (32.39)	97.44 (10.30)	105.83 (30.61)	57.97 (17.80)	112.50 (33.72)	96.05 (10.16)	102.14 (30.67)	51.27 (16.36)	94.06 (32.54)	91.06 (9.87)
Log (ME)	-24.06 (3.89)	-50.76 (4.36)	-32.30 (4.14)	-61.51 (3.04)	-21.24 (3.91)	-53.23 (4.54)	-19.71 (4.43)	-60.77 (3.07)	-20.86 (3.82)	-50.52 (4.19)	-35.08 (4.20)	-65.10 (2.99)
Vol. Constr.	4.90 (1.07)	16.94 (0.93)	3.26 (0.71)	6.58 (0.84)	4.40 (1.08)	16.44 (0.94)	6.72 (0.75)	6.60 (0.84)	5.17 (1.12)	15.48 (0.89)	5.07 (0.74)	6.08 (0.83)
Price Constr.	-10.89 (0.75)	-21.65 (0.89)	-20.82 (0.72)	-20.09 (0.60)	-9.25 (0.71)	-21.89 (0.90)	-20.66 (0.69)	-20.21 (0.60)	-10.95 (0.71)	-23.17 (0.85)	-22.06 (0.68)	-20.48 (0.61)
VIX	-0.25 (0.12)	0.40 (0.13)	0.11 (0.08)	0.30 (0.09)	-0.31 (0.12)	0.40 (0.13)	0.19 (0.09)	0.30 (0.09)	-0.31 (0.12)	0.31 (0.13)	0.03 (0.08)	0.28 (0.09)
Idio. Vol.	-0.02 (0.03)	0.20 (0.03)	0.11 (0.03)	0.19 (0.03)	0.00 (0.03)	0.19 (0.03)	0.15 (0.03)	0.19 (0.03)	0.00 (0.03)	0.18 (0.03)	0.10 (0.03)	0.18 (0.03)
Child Agg.	2.20 (1.08)	3.70 (1.32)	10.35 (1.52)	3.60 (0.96)	-	-	-	-	-	-	-	-
Child Day	-	-	-	-	-10.49 (1.18)	2.90 (1.50)	-25.79 (1.31)	-5.27 (1.00)	-	-	-	-
Child Display	-	-	-	-	-	-	-	-	11.92 (2.11)	-12.05 (1.34)	18.51 (1.62)	-5.01 (0.77)
Impact (%)	11	10	23	10	-50	7	-72	-15	30	-32	32	-17
N	52,486	91,365	93,308	627,506	52,486	91,365	93,308	627,506	52,486	91,365	93,308	627,506
Adj-R ²	0.02	0.04	0.04	0.02	0.02	0.04	0.05	0.02	0.02	0.04	0.04	0.02

Table 10**Regressions of parent order trading costs on run-level characteristics**

Parent-level implicit trading costs are calculated as weighted average prices of executed child orders scaled by the prevailing midpoint before the start of the parent order. The table shows regressions of parent-level costs on the following set of regressors common to all specifications: the size of the parent order scaled by average daily volume (“parent size”), the logarithm of market value of equity (ME), an indicator variable if the parent order was subject to a volume constraint (“vol. constraint”), an indicator variable if the parent order was subject to a price constraint (“price constraint”), the level of the VIX, and idiosyncratic volatility measured using the standard deviation of residuals from a daily regression of firm-level returns on the market over the prior year. The first set of specifications include the proportion of traded volume generated by aggressive runs. The second set of specifications include the fraction of time the parent order deploys aggressively priced runs. The sample consists of all parent orders from algorithms A, B, C and D with at least 50 child orders trading at least 1 basis point of average daily volume. Standard errors clustered by day appear in parentheses.

	Panel A				Panel B			
	A	B	C	D	A	B	C	D
Intercept	79.38 (11.59)	133.51 (12.09)	102.52 (12.25)	165.61 (8.86)	80.38 (11.66)	135.81 (12.04)	114.36 (12.41)	170.02 (8.64)
Parent size	43.90 (31.99)	42.69 (14.49)	56.95 (32.58)	81.00 (9.28)	45.36 (32.10)	42.71 (14.50)	60.97 (32.81)	81.84 (9.32)
Log (ME)	-26.65 (3.97)	-54.96 (4.21)	-36.24 (4.43)	-63.11 (3.22)	-27.27 (3.98)	-56.72 (4.21)	-39.75 (4.48)	-67.35 (3.12)
Vol. constraint	3.52 (1.07)	14.11 (0.87)	1.63 (0.72)	5.86 (0.82)	4.13 (1.10)	14.13 (0.88)	1.83 (0.72)	5.83 (0.82)
Price constraint	-9.64 (0.69)	-20.64 (0.86)	-19.45 (0.69)	-18.74 (0.62)	-9.58 (0.72)	-20.81 (0.87)	-18.96 (0.73)	-18.18 (0.61)
VIX	-0.14 (0.11)	0.31 (0.13)	0.09 (0.07)	0.29 (0.09)	-0.17 (0.11)	0.32 (0.13)	0.05 (0.07)	0.27 (0.09)
Idio. Vol.	-0.03 (0.03)	0.16 (0.03)	0.10 (0.03)	0.18 (0.03)	-0.03 (0.03)	0.16 (0.03)	0.10 (0.03)	0.17 (0.03)
Log (Number of runs)	2.15 (0.26)	6.57 (0.40)	4.33 (0.37)	2.49 (0.23)	2.21 (0.26)	6.77 (0.40)	3.40 (0.36)	3.42 (0.22)
Aggressive run volume (%)	12.91 (1.40)	5.81 (1.41)	25.94 (1.67)	3.90 (1.23)	-	-	-	-
Aggressive run duration (%)	-	-	-	-	4.69 (1.31)	1.02 (1.38)	17.64 (1.37)	4.14 (1.07)
Impact (%)	24	9	23	8	13	2	21	12
N	52,486	91,365	93,308	627,506	52,486	91,365	93,308	627,506
Adj-R ²	0.02	0.04	0.05	0.02	0.02	0.04	0.05	0.02




Figure 1: The pie chart shows the distribution of the total dollar value of parent orders from each institution across all trading algorithms. Each institution is represented by a unique color.

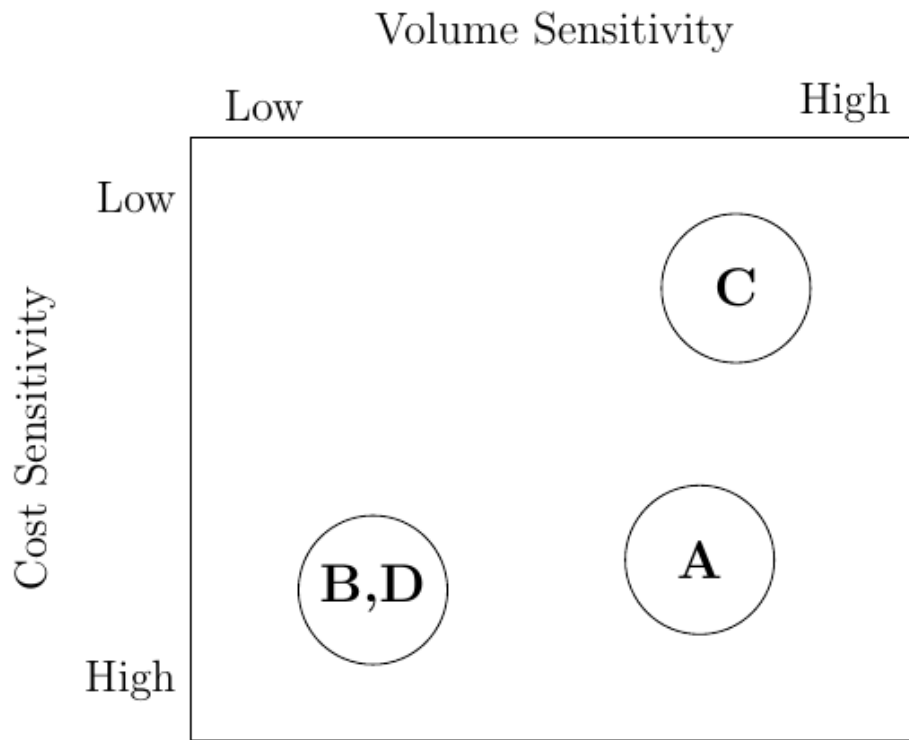


Figure 2: The figure shows the relative sensitivity of trading algorithms (labelled A through D) to trading costs and trading volume. The circles are meant to show indicative relative sensitivities and are not precise or absolute descriptions.




Figure 3: The chart shows the percent of parent orders constrained by price limits, volume limits, price and volume limits, and cancelled before completion. The statistics are shown separately for buys, sells, and short sales for each algorithm.

Table A1**Parameter estimates from accelerated failure time models of limit and PEG order execution**

The sample includes all day child orders. The table contains parameter estimates from accelerated failure time models for limit and PEG orders under the generalized gamma distribution (see Lo, MacKinlay and Zhang (2002) for details). The intercept captures aggressively priced orders, followed by indicator variables for each price aggressiveness category. Scaled order size is the size of the submission, scaled by total depth at the NBBO. Book asymmetry is measured (in percent) as total depth at the bid minus total depth at the ask, scaled by average depth at the bid and ask. The parameter estimates for scaled order size and book asymmetry are multiplied by 100 for presentation clarity. The buy indicator is equal to one for buy orders, zero for sells and short sales. The display order indicator is equal to one for displayed orders on lit exchanges, zero for non-displayed orders on exchanges or dark pool orders. The lit indicator is equal to one for orders sent to lit exchanges. The logarithm of market capitalization is measured as of the day prior to the order. The logarithm of volume is calculated from average dollar volume over the prior 20 days. $|R_{-5,0}|$ is the absolute value of returns 5 minutes prior to the submission of the orders.

	Limit orders			PEG orders	
Intercept (Aggressive)	1.011 (0.008)	0.741 (0.002)	1.015 (0.002)	1.308 (0.007)	1.361 (0.007)
(Agg., Mid.)	1.798 (0.003)	1.733 (0.003)	1.798 (0.003)	-	-
Midpoint	5.296 (0.001)	5.254 (0.001)	5.297 (0.003)	-0.461 (0.003)	-0.459 (0.003)
(Mid., Pass.)	7.377 (0.001)	7.326 (0.001)	7.377 (0.001)	-	-
Passive	9.380 (0.001)	9.320 (0.001)	9.380 (0.001)	3.752 (0.001)	3.754 (0.003)
Scaled order size	0.000 (0.000)	0.000 (0.000)	-0.040 (0.000)	0.000 (0.000)	0.000 (0.000)
Book Asymmetry	-	-	-0.010 (0.000)	-	0.000 (0.000)
Buy indicator	-	-	-0.008 (0.001)	-	-0.101 (0.001)
Display order indicator	-0.542 (0.000)	-	-0.542 (0.001)	-	-
Lit indicator	-	-0.195 (0.001)	-	-	-
Log (market cap)	0.549 (0.000)	0.532 (0.000)	0.549 (0.000)	2.245 (0.001)	2.247 (0.001)
Log (volume)	-0.323 (0.000)	-0.320 (0.000)	-0.323 (0.000)	-0.379 (0.000)	-0.380 (0.000)
$ R_{-5,0} $	-0.021 (0.000)	-0.022 (0.000)	-0.021 (0.000)	0.004 (0.000)	0.004 (0.000)
Scale parameter (σ)	2.458 (0.000)	2.473 (0.000)	2.458 (0.000)	7.041 (0.001)	7.041 (0.001)
Shape parameter (ν)	0.025 (0.000)	0.086 (0.000)	0.025 (0.000)	-3.137 (0.001)	-3.136 (0.001)