The Anatomy of Trading Algorithms

Abstract

We study the anatomy of four widely used standardized institutional trading algorithms representing $675 billion in demand from 961 institutions between 2012 and 2016. The central tradeoff in these algorithms is between the desire to trade and transaction costs. Large parent orders generate hundreds of child orders which strategically employ the price, time, and display priority rules embodied in market structure to navigate this tradeoff. The distribution of child orders is non-random, generating strategic runs which oscillate between providing and taking liquidity. Price impact occurs both at the time an order is submitted to the book (regardless of whether it is filled), and at the time of execution. Passive child orders have much lower likelihood of execution but still incur substantial price impact. Conversely, marketable orders, even though immediately executable, do not necessarily guarantee execution and generate even larger price impact.

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 making firms still attempt to set regret-free quotes. But the nature of information and adverse selection is intimately tied to horizon and order splitting. Recognizing these issues, and especially the advent of high frequency trading, O'Hara (2015) issues a clarion call for research to update learning models and traditional empirical methods in this new trading environment.

Much of the recent research that follows this direction has focused on the activities and implications of 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 information and protection from adverse selection are even more important in an environment in which HFTs are often viewed as adversarial.[2] 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

[1] The list of papers that seek to understand the behavior and impact of HFTs is large (see surveys by Biais and Wolley (2012), O'Hara (2015), and Menkveld (2016)). Much of this research focusses on the effects of the speed advantage of HFT's and its implications for other agents' behavior and market equilibrium.

[2] Studies that use data on institutional trading provided by Ancerno do not reveal the microstructure of institutional trading strategies because they are based on end-of-day trading tickets, not individual orders or executions, and lack timestamps in the majority of the data (see Hu et al. (2018)).

total trading volume (see, for example, Pucket and Yan (2011) and various surveys conducted by the Tabb Group).

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 features back to the changes in concepts and constructs underlying market microstructure research. Our hope is that our findings are helpful to researchers who seek to build intuition in developing models of price formation, as well as guide research design in empirical tests that examine aspects of market structure. Second, we seek to understand the tradeoff between transaction costs and the likelihood of execution, a fundamental issue in all trading, and a building block of order choice models.[3] Because wayward trading algorithms are sometimes associated with market structure induced volatility (e.g. flash crashes), 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.

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 trading algorithms used by 961 unique institutions representing over $675 billion in aggregate demand. These are widely-used non-bespoke agency algorithms, produced in some form by almost all brokerage firms and deployed worldwide. The four differ in the importance that they attach to trading versus the transaction costs of achieving their volume objective in a timely fashion. Their anatomical structures are similar across brokerage firms which means that insights from the data are generalizable. Most importantly, the data are rich in the detail required to draw sharp inferences. 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. Our 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

[3] This tradeoff is also described as a price-time tradeoff by Demsetz (1968), and Cohen et al. (1981), who argue that impatience has a positive price.

that allow us to examine fundamental tradeoffs in electronic trading. All data are time-stamped to the millisecond, allowing precise computations of short horizon price movements.

We start with gross anatomy. 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. Institutional traders impose price caps/floors and/or volume constraints at the parent level. These restrictions are independent of the pre-programmed behavior of an algorithm. The frequency of price, volume, and price-plus-volume constraints in our sample average 21, 17, 10 percent respectively, with enormous variation across algorithms: in some algorithms, price or volume constraints average over 50 percent, sometimes reaching levels of 90 percent. The average cancellation rate is 52 percent, although again, there is considerable variation across algorithms. The variation across algorithms is consistent with differing tradeoffs with respect to volume versus transaction costs. These features also imply that while institutions' turn over their order flow to a black box to the handle the arcane details of market structure, this does not prohibit them from maintaining risk controls.

Of the 300 million child orders, less than 0.40 percent 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 in our sample is limit orders (81.5 percent), followed by PEG orders (18.1 percent), which are dark, exchange-disseminated limit orders that are dynamically "pegged" to the NBBO (i.e., move lock-step with it). This is in stark contrast to most sequential or strategic trade models in which trades takes place when market orders interact with limit orders (see, for example, Holden and Chakravarty (1995), Harris and Hasbrouck (1996), Parlour (1998), Large (2009), and many others). Such assumptions are also common in early limit order models, along with an information structure 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, our data suggest that the majority of information is impounded into prices via limit, not market, orders.[4]

[4] 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

The implication is that instead of modeling the dichotomous market versus limit order choice, it may be more fruitful to examine the price-time tradeoff within limit and PEG orders. On the theory front, recent models by Goettler, Parlour, and Rajan (2005, 2009) relax the symmetric information constraint, but they are forced to give up analytic solutions and use numerical techniques. Riccò, Rindi, and Seppi (2018) come even closer to what we observe in the data, modelling a dynamic limit order market with asymmetric information where the price impact and information content of newly arriving orders is conditional on the order book history. On the empirical front, we examine the price-time tradeoff at the child order level.[5]

Of the 169 million limit orders in our sample, only 24 percent are marketable, 65 percent are passive, and the remainder are priced inside the spread. For PEG orders, only 9 percent of the 130 million orders are marketable, 67 percent are priced at the midpoint, and the remainder (24 percent) passive. Fill rates are not 100 percent, even for marketable orders. Only 76 (10) percent of marketable limit (PEG) orders are filled and the remainder are cancelled. In limit order models, it is common practice to regard marketable limit orders as observationally equivalent to market orders because they are immediately executable (an exception is Peterson and Sirri (2002)). Our data show that this is not the case. Marketable limit orders do not guarantee execution, potentially for a variety of reasons. For instance, a dark marketable limit order is not guaranteed execution. Even a displayed marketable order may not get executed if it is sent to a venue that does not have a contra order at the limit price and the order is non-routable; it is either rejected (at NYSE and ARCA), or repriced (Nasdaq) to prevent locked markets. Another possibility is that a marketable dark order specifies a minimum fill instruction and is larger than available quantity. Regardless of the specific reason, matching engine protocols matter for our understanding of the endogenous choices made by traders, and therefore for price formation. In this context, the abundance of order

[5] By focusing on child orders, the maintained assumption is that market participants cannot infer the underlying parent order, or the sequence of future child orders, from the past sequence of publicly visible child orders. The counterfactual would imply that order anticipation strategies are successful. Beyond anecdotes, we know of no systematic evidence that order anticipation strategies are successful. Theoretical models of such strategies are sensitive to assumptions about market structure and information availability (see, for example, Brunnermeier and Pedersen (2005), and Bessembinder et al. (2016)).

types offered by exchanges can be interpreted as a way for traders to fine tune execution likelihood and trading costs, rather than introducing unnecessary complexity that benefits HFTs.[6]

Since fills and cancellations are related, we model them jointly using the time-to-fill accelerated failure time model of Lo, MacKinlay and Zhang (2012). The approach accounts for the censoring that occurs when limit orders are cancelled and allows us to focus on three choice variables determined programmatically by the trading algorithm: submission price, submission size, and whether an order is displayed. Our estimates suggest that controlling for market and security effects, the most significant effect on time to fill comes from the price decision. A movement in the submission price from the aggressive side of the spread (i.e. a marketable order) to the passive side increases the time to fill by about 500 seconds. Interestingly, order size, or asymmetry in the depth of limit order book between the bid and ask side does not appear to influence time-to-fill. However, non-displayed orders, either designated as non-displayed on exchanges, or submitted to dark pools, double the time-to-fill.

Price impact can occur in two stages, at the time an order is submitted to the limit order book, and subsequent to execution (which obviously influences subsequent orders and executions). We measure price impact using quote midpoints after both events, separately for unfilled and filled orders, at horizons ranging from 100 milliseconds to 10 seconds post-submission and post-trade. At the 10 second horizon, marketable orders move quotes by 2.03 basis points post-submission. But passive orders also move prices by 0.84 basis points post-submission. This has 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 unexecuted orders contribute to the overall cost of trading a particular parent order. This cost, along with execution risk, is ignored in models of optimal execution (e.g. Bertsimas and Lo (1988), Almgren and Chriss (2000), and others). The post-execution price impact of marketable orders is large, as liquidity providers move quotes in response to aggressive limit orders. The post-execution price impact of passive orders reflects the adverse selection faced by these resting orders. This implies that the tradeoff for agency algorithms is between the cost of paying the spread and the sum of price

[6] Many market participants consider the sheer number of order types made available by exchanges and their complexity as generating trading advantages for HFTs. See, for example, SEC (2015, page 22-23), and references therein.

movements associated with both submissions and execution. Rather intuitively, cross-sectional variation in these price movements is also positively related to the size of the order relative to the prevailing depth, and whether the order is displayed.[7] It is also larger in high VIX periods, an outcome is that is outside of an algorithm's locus of control, but related to the aggregate risk bearing capacity.

Finally, we examine the intra-parent time series distribution of child orders in large parent orders. We find that child orders are strung together in strategic "runs" in which they are systematically on the passive side of the spread, inside the spread, or are marketable.[8] The average parent order contains about 63 such runs lasting about 566 seconds with each run containing almost nine consecutive child orders. The 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 between providing and taking liquidity.

Our results 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 like electronic market makers, 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. Liquidity provision can change to liquidity extraction rapidly, with implications for liquidity risk, an issue that we leave for subsequent work. Moreover, by posting resting quotes, agency algorithms compete with electronic market makers. This competition can have intriguing 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.

---

[7] 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)).

[8] 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.

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 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 and 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.[9] 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.

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"). Assuming appropriate communication systems, 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 time-stamp. 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, ranging from

[9] 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.

buy-side long-only investment managers to multi-asset hedge funds. 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) or 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.[10] Large providers of such algorithms also white-label their offerings to smaller brokerage firms. Moreover, we restrict our attention to single-stock algorithms so that we can focus our attention on stock-specific execution issues without having to concern 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. Although we are not at liberty to identify or describe each algorithm, we can report that they reflect differences in the importance that they attach to trading volume versus transaction costs. As such, they are an economically meaningful source of variation in the data and revelatory in our tests. Other parent order information include a stock identifier

[10] See the 2019 Algorithmic Trading Survey (https://www.thetradenews.com/surveys/algorithmic-trading-survey-long-results-2019/).

(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.[11] 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) which shows whether the execution added liquidity, removed liquidity, or was routed out, and the trading fee paid or rebate earned by the order.[12]

Two other 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.[13] 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

[11] These algorithms employ exchange-supplied PEGs rather than synthetic PEGs so that if the NBBO changes, the order's PEG price is immediately updated.
[12] 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.
[13] Two recent studies observe parent orders and fills, but not submissions. Saglam et al. (2019), and Battalio, Hatch, and Saglam (2018) observe parent orders and child order fills (but not submissions) for a small sample of 22,000 parent orders in S&P 500 stocks over 14 months.

We match the algorithmic trading data with market data from daily Trade and Quote (TAQ) files with millisecond timestamps. For each relevant timestamp, 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.

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 Depository 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 57 percent of parent demand, and 18 percent of notional child order volume.

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 by dollar volume also shows no particular spikes or patterns.

4. Algorithm Anatomy

4.1 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. Algorithm D is most widely used, deploying over 1.7 million parent orders with an aggregate value of $407 billion. Buy orders are generally more frequent than sells. The number and dollar value of parent orders that are short sales are smaller than for sells, but not by a large margin.

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 parent order size as a percentage of average daily dollar volume over the prior 20 days. By this metric, average parent size is also quite large, 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, often 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. Across all algorithms 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 duration is 84.19 minutes. Algorithm A is has shorter duration (averaging between 17 and 20 minutes), followed by algorithm C (between 231 and 36 minutes). Algorithms B and D attempt to trade for substantially longer intervals, between 81 and 104 minutes. There is also considerable variation in parent order duration with each algorithm, likely linked to the variation in 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 2 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. Overall, 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.

4.2 Child Orders

Panel A of Table 2 shows the numerical and dollar distribution of child orders. We first discuss these data across all parent orders, before examining variation across algorithms and buy/sell/short sale categories. 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. Child orders posted to the limit order book generate quote updates in exchange feeds and the SIP. The child to fill ratios implied by these statistics are lower than the quote-to-trade ratios commonly observed in market data. For instance, Conrad et al. (2015) report quote update to trade ratios ranging from 10 to 20 across the cross-section of securities (see also the cancel to trade and trade to order volume ratios reported by SEC's Midas system at https://www.sec.gov/marketstructure). The difference is likely driven by the much higher quote updates originating from high frequency liquidity providers instead of buy-side trading algorithms.

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 retail investors regularly use market orders over half the time. The vast majority of algorithmic child orders are either limit orders (81.48 percent), followed by PEG orders (18.12 percent). The latter are immediately repriced when the NBBO moves and are therefore subject to lower execution risk than 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. 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 three other key characteristics of child orders, pertaining to price, display, and venue choice. Marketable orders are those priced at or above the far side of prevailing NBBO (e.g. buy orders priced at the best ask or higher are immediately marketable). Over 42 percent of child orders are marketable, implying that the majority (58 percent) provide liquidity in the sense that they rest on the book for a certain duration and provide trading options for others. About 75 percent of child orders from a parent are visible on exchange feeds and the remainder (25 percent) are either hidden orders on exchanges or posted to dark pools. While 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.

Table 2 also shows the above statistics for each algorithm and for buy, sell, and short sale categories. 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 is, however, much variation across algorithms. Variation in limit versus PEG order roughly mirrors child to fill ratios across algorithms. Algorithms A and D which have higher child to fill ratios are more likely to use PEG orders as a way of maintaining queue priority. In contrast, algorithms B and C largely employ limit orders (over 75 percent and almost 90 percent respectively). Algorithms A and C are much more likely to use marketable orders (over 60 percent and 70 percent respectively), compared with algorithms B and D (slightly less than 40 percent). Algorithms A and C are less likely to use displayed orders (between 28 and 35 percent), and less likely to use Lit exchanges (between 36 and 43 percent). This suggests that these two algorithms are more sensitive to price impact of their child orders. Algorithms B and D, on the other hand, use displayed orders the vast majority of time (between 69 and 87 percent), and generally favor Lit venues (between 73 and 88 percent).

5. Intra-Parent Analysis

Thus far we have focused on the parent level. In this section, we study child orders, treating them as independent. We take the perspective that market participants do not observe parent level information, and cannot easily aggregate child order arrival, trade, and cancellation into expected future order flow. 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, modelling dependence and order anticipation strategies is outside the scope of our paper. Our simpler econometric approach permits straightforward inference with respect to execution likelihood and price impact.

5.1 Child Order Choices

We start by examining 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 in spirit to Biais, Hillion, and Spatt (1995) based on the aggressiveness of its submission price relative to the prevailing NBBO: (a) marketable 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 the near-side quote (bid price for buys, ask price for sells). Note that orders between the near-side quote and the midpoint (i.e. (d) above), offer price improvement and 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 reaches the market. Total depth at the NBBO is computed using all 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 in the sample, 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 marketable limit orders is 13.3 percent of available depth but only 3.0 percent for passive limit orders. Similarly, the median order size for PEG orders falls from 20.0 percent for marketable orders to 9.0 percent for passive orders. In both panels, passive orders are smaller because lower order sizes reduce the economic exposure to adverse selection, the major cost faced by such orders. 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.

5.2 Execution Likelihood

The purpose of buy-side trading algorithms is to trade while minimizing execution costs. Execution likelihood is therefore a key outcome of interest. We begin by examining frequency distributions of order usage and their fill rates. The first column of Table 4 shows the number of child orders in each submission price category. Panel A contains counts for limit orders (for all orders, non-displayed orders, and displayed orders). Panel B provides equivalent information for PEG orders. Subsequent columns labelled "N" show counts for unfilled and filled child orders.

Of the 169 million child limit orders in our sample, 65 percent (111 million) are passive. About 24 percent (41.6 million) are marketable, and only 5.9 million are at the midpoint. The majority of PEG orders, in contrast, are tied to the midpoint (88.2 million). Here too, the number of passive orders is more than twice the number of marketable orders. 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, 1.9 and 9.1 million respectively. As described earlier, orders posted between the near (passive) side of the quote and the midpoint definitionally narrow quotes. Orders can only do this 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 marketable to passive orders are 0.01, 0.05, 0.01, 0.07, and 0.1 respectively. More importantly, in all 9.1 million orders that are priced between the near side quote and the midpoint, the quoted spread is greater than one penny.

Of the 111 million passive limit orders, 31 percent (34.9 million) are filled while the remaining are cancelled. For passive PEG orders, the fill rate of 19 percent (5.8 million out of 30.4 million) is even lower. The fill rates for marketable limit and PEG orders are 76.9 percent (32 out of 41.6 million), and 89 percent (10.9 out of 12.2 million) respectively.

These statistics reveal several interesting features of the data. First, pricing limit or PEG orders aggressively so that they are marketable does not eliminate execution risk. There are potentially many reasons why a marketable order may not guarantee even partial execution. Certain order types are not guaranteed execution, and matching engine protocols, timing differentials and other such factors can also result in non-execution. For instance, if a displayed marketable 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. Yet another possibility is that the marketable order has a minimum fill instruction and is of a size that is larger than the contra-side order. Regardless of the specific trader choices and reasons, market design rules generate execution risk that cannot be overcome with aggressive limit prices. Second, the fact that the majority of child orders are passive suggests, by revealed preference, that these algorithms are concerned with the joint distribution of transaction costs and execution likelihood. Although passive child orders attempt to trade directionally, they do not "pay" the spread and instead provide liquidity in the sense that they post resting orders. Aside from non-execution, an important cost associated with resting orders is adverse selection, an issue that we address in section 5.3.2.

Table 4 also reports the 25th percentile, median, and 75th percentile of the life of child orders, measured in seconds from submission time. This is the time-to-cancel for unfilled child orders, and time-to-fill for filled orders.[14] As before, Panels A and B report statistics for limit and PEG orders respectively.

For marketable limit orders, the median time-to-cancellation and time-to-fill are identical (10 milliseconds). The 75th percentile of the time-to-fill is 27 milliseconds but for time-to-cancel it is 483 milliseconds. For marketable PEG orders, the distributions for time-to-fill and time-to-cancel are quite similar, suggesting that if the order is priced aggressively and does not execute, it is cancelled at roughly the same time at which one would expect it to execute. For passive orders, the median time-to-cancel is almost twice as long (45 seconds) as time to fill (27.28 seconds). Similarly, the median time-to-cancel for passive PEG orders is 20.26 seconds, compared to 11.28 seconds for time-to-fill.

The results in Table 4 are indicative of the fact that limit orders cancellations affect the conditional distribution of time-to-fill. Limit order execution is not only a function of the endogenous choice variables of interest but also a random function of market conditions. To assess the effect of submission price, size, and display on the time to execution, we adapt the accelerated failure time limit order model of Lo, MacKinlay and Zhang (2012). The model is part of a class of parametric survival models particularly well suited to this setting. Following their approach, T is a nonnegative random variable that represents the life of a limit order, so that $f(t)$ and $F(t)$ are the PDF and CDF of T. The sequence of realization of T are $(t_1, \ldots, t_n)$, which may be censored when the limit order is cancelled or expires. Define δ_i as an indicator variable equal to 1 if the observation is censored, and zero otherwise. Given the pairs (t_i, δ_i), the likelihood function is given by

[14] 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.

$$\prod_{i=1}^{n} f(t_i)^{\delta_i} S(t_i)^{1-\delta_i} = \prod_U f(t_i) \prod_C S(t_i) \qquad (1)$$

where U and C correspond to uncensored and censored observations, and $S(t_i)$ is the survivor function (*1-F(t)*). The accelerated failure time model is

$$T = T_0 e^{X'\beta+\varepsilon} \qquad (2)$$

where T_0 is the baseline failure time rate, T is the time to execution, β is a vector of parameters, and X is a vector of explanatory variables. The conditional quantile q_τ is given by

$$q_\tau = \exp(X'\beta)(\nu^2)^{(\sigma/\nu)} G^{-1}(\tau, \nu^{-2}))^{(\sigma/\nu)} \qquad (3)$$

where G is the gamma function, and σ and ν correspond to scale and shape parameters from the gamma distribution.

Lo, MacKinlay and Zhang (2012) present a persuasive case for using the generalized gamma distribution for estimating the accelerated failure time model, rather than the widely used proportional hazard model. We follow their suggestion, estimating separate models for limit and PEG orders. The primary explanatory variables of interest are price aggressiveness, submission size, and whether the order is displayed. We employ four dummy variables for price aggressiveness categories, leaving marketable orders in the intercept. As in prior tables, submission size is scaled by total depth available at the NBBO. Differences in depth between the bid and ask side of the market can be significant, and placing orders on one side can either exacerbate or ameliorate asymmetry, influencing time to execution.[15] Therefore, we also estimate a model with book asymmetry (measured in percentage terms as depth at the bid, minus depth at the ask, divided by average depth) and a buy/sell indicator variable. To investigate the effects of order visibility, we use to an indicator variable equal to one when an order is displayed. We also estimate a specification in which we use an indicator variable if the order was sent to a lit exchange (and zero if was routed to a dark pool). This allows us to determine the incremental effect of routing from display decisions. Finally, we include three variables that correspond to market and stock specific covariates. The absolute value of returns over the prior 5 minutes ($|R_{-5,0}|$) measures stock volatility, relevant for picking off risk (Foucault (1999)). The logarithm of the market

[15] In Parlour (1998) and other models, changes in book depth on one side of the market influences limit order submission probabilities endogenously.

capitalization, and the logarithm of trading volume over the prior 20 days, account for stock-specific effects.

Table 5 presents parameter estimates with standard errors in parentheses. Consistent with the results in Lo, MacKinlay and Zhang (2012), the scale (σ) and shape parameter (ν) have small standard errors, implying that using simpler distributions (such as the Weibull or lognormal) would be inappropriate. The last two rows show the number of censored (cancelled limit orders) and uncensored (filled limit orders) observations employed in the estimation. The large size of the sample imposes computational burdens but allows the parameters to be precisely estimated. As a result, we do not discuss the statistical significance of the parameter estimates, focusing instead on their economic content.

Price aggressiveness stands out the main determining factor in time-to-fill. Moving from a marketable limit order to a passive order is associated with over a 450,000 percent increase in uncensored time-to-fill, which is the equivalent of a move from 1 millisecond to 450 seconds. The coefficient on the display indicator is -0.77, suggesting the effect of displaying an order approximately halves its time-to-fill. The estimates for scaled order size and book asymmetry, while statistically significant, are economically tiny. Trading in larger stocks, in stocks with higher volume, and during periods of increased return volatility all shorten the time-to-fill. Our estimates suggest that price and display decisions are among the most meaningful when considering time to execution.

5.3 Price Impact

5.3.1 Measurement

We measure the price impact of a child order *j* as

$$cpi_{jt\tau} = q_{jt}(m_{j,t+\tau} - m_{jt}) / m_{jt} \tag{4}$$

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 (2019), we use a variety of horizons for τ, corresponding to 100 milliseconds, 500 milliseconds, 1 second, 5 seconds, and 10 seconds.

Ideally, we would measure price movements subsequent to the submission of an order, and after its execution. That is complicated by the fact that the time between submission and execution is so often short that post-trade price movements conflate post-submission price movements. For example, the median time-to-fill for marketable orders is 10 milliseconds (Table 4), which implies that even at the shortest horizon (τ=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 nonetheless allows us to cleanly estimate price impact. To the extent that non-executed orders are likely to be less informative, this biases us against find that order arrival per se has price impact. For unfilled orders, 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 in a 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

The upper part of Panel A of Table 6 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 1/20th 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 marketable 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 limit orders, irrespective of execution.[16] A comparison of point estimates is also 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

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

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 marketable 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, marketable PEG orders incur a price impact of 1.06 basis points, roughly half that of marketable limit orders. Similarly, passive limit orders have a price impact of 0.31 basis points, compared to 0.84 basis points for limit orders.

Since the price impact of each individual order is small, readers may worry that the price movements we detect are just part of the normal market dynamic, as opposed to being causally related to child order submission. We do not believe this is the case because of the precision of time stamps. Nonetheless, we perform 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 from this placebo test 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.

It is tempting to interpret the average price impact of child orders in the context of parent level information. The average parent order in Table 1 generates 88 unfilled child orders (126 child orders minus 38 fills). In this prototypical parent, approximately 42 percent of child orders are marketable and the remaining are passive. If the average marketable order has a price impact of 2.03 basis points, and the average passive order incurs a price impact of 0.84 basis points, this implies that all child orders submissions move prices by 117 basis points (0.42*88*2.03 + 0.58*88*0.84). This seems implausibly high. One explanation resides in the distribution of price impact. 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 of for marketable and passive orders is 5.02 and 3.37 basis points respectively, over the twice that of the mean. Similarly the percentage of orders with zero price impact is 54.67 percent for marketable orders and 61.64 percent for passive orders. These statistics imply that while some limit orders clearly move prices, a large fraction do not.

A second explanation of why it may be inappropriate to aggregate these results to the parent order level is that unfilled child order submissions are only one source of price impact at the parent level. The right set of columns in Table 6 provide equivalent statistics for filled child orders, where price impact is measured from the time stamp of the fill (not submission). As expected, marketable 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 impact, or equivalently 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. These are comparable to the aggregate price impact statistics for trades reported in Conrad and Wahal (2019).

Finally, 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, 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 marketable 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).

5.3.3 Variation in Price Impact

We first examine variation in price impact with respect to order choice attributes using a simple triple sort procedure. On each day, we independently sort limit orders into three price aggressiveness categories, scaled order size quartiles, and two categories that correspond to whether the order was displayed or not.[17] The equivalent sorts for PEG orders only use price

[17] For limit orders, we restrict our attention to marketable, 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.

aggressiveness and size quartiles because all PEG orders are hidden. For each group, we calculate average price impact within a day. Table 7 shows time series averages of the daily group averages, separately for unfilled and filled orders. As before, Panel A and B correspond to limit and PEG orders respectively.

Holding price aggressiveness and display constant, the price impact of unfilled orders rises monotonically with scaled order size. For displayed marketable 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. In non-displayed limit orders, the level of price impact is lower, but the pattern across order size quartiles is the same: for marketable 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 marketable 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 marketable 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. The bottom line is that variation in submission price, scaled order size, and display is systematically related to post-submission price impact.

The post-fill variation in price impact is a bit more complex. In marketable limit orders, holding display constant, price impact increases with order size. And holding order size constant, price impact is higher for displayed orders. But for passive orders, the effects are not quite so clear cut. 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. There is no such inconsistency in the spread between displayed versus non-displayed orders (holding order size constant): in all four size quartiles, displayed passive orders have more negative price impact than displayed orders.

Given the above results, we estimate daily price impact regressions separately for unfilled and filled orders. The regressions estimate the incremental effects of price aggressiveness, scaled order size, and display, and have two other advantages. 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 8 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 adjusted R^2 is noticeably higher for limit order regressions.

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 8, 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 marketable limit orders measured at τ=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

The analysis thus far treats each child order as independent. However, child orders belonging to the same parent order are clear not independent; each child order is being generated in pursuit of a common goal (i.e. to fill the desired demand) from a shared codebase. We are interested in quantifying the dependence structure that arises from strategic behavior as the 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 in Table 5 to three groups: passive, marketable, 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. 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 9 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 surprisingly long (566 seconds), and involves submitting over 1,500 shares. The percentage of runs that are passive, marketable, and inside the spread are 45, 16, and 38 percent respectively. There is large variation in all these features across algorithms. For instance, strategic runs generated by Algorithm A have lower duration (102 seconds) but submit a larger number of shares, and are equally likely to be passive or marketable (40 versus 43 percent). In contrast, strategic runs from Algorithm D are six times as long (683 seconds), but submit far fewer shares (732), and are far more likely to be passive (47 percent) than marketable (37 percent). This variation reflects differences in the need to trade and the sensitivity of each algorithm to trading costs.

Panel B shows transition matrices between successive passive, marketable, and inside runs. Across all algorithms, the probability that an algorithm in a passive run moves to a marketable state is 31.2 percent. The equivalent transition probability from a marketable 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 marketable, the transitions probabilities between phases are quite similar and symmetric.

The transition matrices descriptive of the unconditional probability of moving from one type of run to another. We investigate whether these transitions depend on the prior run using logistic regressions. Specifically, we estimate separate regressions conditional on the price aggressiveness of the prior run. We use two dependent variables: (a) if the prior run is either passive or inside the spread, the dependent variable is equal to one for marketable orders, and zero otherwise, (b) if the prior run is marketable, the dependent variable is equal to one if the current run is passive, and zero otherwise. 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, multiplied by +1 for buys and -1 for sells. The signed return reflects the implicit cost incurred with the prior run.

Panel C presents the results of these regressions. Standard errors appear in parentheses under the coefficients. 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, the change in probability is based on a 1 basis point change in the signed return. Two results stand out. First, the probability that a run is marketable or passive appears to be largely unrelated to price movements during the prior run, at least by our measure. This is the case regardless of whether the price movement occurs during a passive, inside, or marketable run. Second, the price aggressiveness of the run is influenced by whether the prior run results in a fill. For instance, in Algorithm A, if the prior run was passive and received a fill, the probability that the subsequent run is marketable rises by 7.59 percent. There are, however, differences in the changes in probabilities across algorithms. For Algorithm D, a passive run receiving a fill has a minuscule effect on the probability that the subsequent run is marketable (0.83 percent). But for the same algorithm, if the fill is received in a marketable run, the probability that the subsequent run is passive rises by 11.21 percent. This type of asymmetry is evident across all algorithms, though it differs in magnitude, which we interpret as differential sensitivity to the tradeoff between the need to trade and the cost of trading.

6. Conclusion

We study the anatomy of four non-bespoke trading algorithm widely used throughout global equity markets. Parent orders in these algorithms spawn hundreds of child orders that attempt to balance transaction costs with execution likelihood. They do so via complex order features, most of which go unexecuted, 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 characteristics of trade in modern equity markets are dramatically different.

Market orders are almost never deployed. Marketable limit orders are widely used as substitutes, and although they are immediately executable, do not guarantee execution. Passive limit or PEG orders attempt to minimize transaction costs by not paying the quoted bid-ask spread, but increase exposure to execution risk and adverse selection in the process. Therefore, microstructure theory models where order choice takes a backseat to price aggressiveness would seem to be a step in the right direction.

Unsurprisingly, 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.

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.

Battalio, R., Hatch, B. and Saglam, M., 2018. The Cost of Exposing Large Institutional Orders to Electronic Liquidity Providers. *SSRN*

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.

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*, 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.

Chakravarty, S. and Holden, C.W., 1995. An integrated model of market and limit orders. *Journal of Financial Intermediation* 4, 213-241.

Collin- Dufresne, P. and Fos, V., 2015. Do prices reveal the presence of informed trading? *Journal of Finance* 70, 1555-1582.

Cohen, K.J., Maier, S.F., Schwartz, R.A. and Whitcomb, D.K., 1981. Transaction costs, order placement strategy, and existence of the bid-ask spread. *Journal of Political Economy* 89, 287-305.

Conrad, J., Wahal, S. and Xiang, J., 2015. High-frequency quoting, trading, and the efficiency of prices. *Journal of Financial Economics* 116, 271-291.

Conrad, J. and Wahal, S., 2019. The term structure of liquidity provision. *Journal of Financial Economics*, forthcoming.

Demsetz, H., 1968. The cost of transacting. *Quarterly Journal of Economics* 82, 33-53.

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.

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.

Goettler, R.L., Parlour, C.A. and Rajan, U., 2005. Equilibrium in a dynamic limit order market. *Journal of Finance* 60, 2149-2192.

Goettler, R.L., Parlour, C.A. and Rajan, U., 2009. Informed traders and limit order markets. *Journal of Financial Economics* 93, 67-87.

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.

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 retial orders and stock returns. *Journal of Finance* 68, 1229-1265.

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. 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., 2018. 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.

Saglam, M., Moallemi, C., Sotiroupolous, M., 2019, Short-term trading skill: An analysis of investor heterogeneity and execution quality. Forthcoming, *Journal of Financial Markets*.

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.

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 shows the percentage of parent orders that volume-constrained, price-constrained, and price-and-volume constrained. The last row shows the percentage of parent orders that are cancelled before completion.

	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: Frequency of parent-level constraints and cancellations (%)													
Volume	17.29	26.39	21.11	32.09	35.93	33.10	38.72	93.34	90.22	96.39	6.43	7.16	5.53
Price	21.87	61.10	61.38	76.80	45.07	49.91	49.64	55.14	55.20	57.10	11.33	12.21	9.67
Price & Vol.	10.25	20.92	17.88	28.23	22.87	20.24	27.34	51.95	51.31	54.75	3.27	3.42	3.05
Cancel	52.10	16.25	17.40	17.87	30.33	36.39	30.89	45.68	48.16	36.86	58.69	55.37	61.96

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 marketable orders is non-marketable orders, the omitted category for displayed orders are non-displayed orders, the omitted category for orders directed to Lit venues is dark venues.

	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	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	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							
Marketable	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 marketable 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 as between the far side and the midpoint, at the midpoint, and between the near side and the midpoint. 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 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		
	25th Perc.	Med.	75th Perc.	25th Perc.	Med.	75th Perc.	25th Perc.	Med.	75th Perc.
Panel A: Child limit orders									
				All orders					
Marketable {far}	2.75	13.3	40.0	4.1	16.6	43.2	1.3	8.6	33.3
{far, 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
{near, midpoint}	0.6	2.4	18.0	0.7	3.6	25.0	0.6	2.0	14.6
Passive {near}	0.6	3.0	11.7	0.6	2.9	11.0	0.6	3.4	1.4
				Non-displayed orders					
Marketable {far}	3.5	14.5	40.0	5.4	20.0	50.0	2.5	10.0	33.3
{far, 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
{near, midpoint}	13.7	27.5	56.0	16.4	33.3	66.6	12.1	23.0	48.6
Passive {near}	5.1	12.5	28.0	5.3	13.0	28.5	4.7	11.6	26.0
				Displayed orders					
Marketable {far}	2.2	12.5	40.0	2.8	12.5	38.5	0.9	7.8	33.3
{far, 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
{near, midpoint}	0.6	1.3	12.0	0.6	1.9	15.0	0.5	1.3	10.5
Passive {near}	0.5	2.5	10.0	0.5	2.4	9.0	0.5	3.0	13.0
Panel B: Child PEG orders									
Marketable {far}	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 {near}	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 child orders
Child orders are unfilled (and therefore cancelled), or filled. 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 marketable 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 as between the far side and the midpoint, at the midpoint, and between the near side and the midpoint. 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	25th Perc.	Med.	75th Perc.	N	25th Perc.	Med.	75th Perc.	N
Panel A: Child limit orders									
All orders									
Marketable {far}	41.6M	0.006	0.010	0.483	19.6M	0.000	0.010	0.027	32.0M
{far, midpoint}	1.9M	0.006	0.010	0.160	1.1M	0.000	0.004	0.043	0.8M
Midpoint	5.9M	0.007	0.390	54.100	2.1M	0.016	2.530	15.284	3.8M
{near, midpoint}	9.1M	29.690	58.900	103.600	3.5M	2.670	13.706	41.900	5.6M
Passive {near}	111M	29.480	45.000	85.710	76.5M	8.303	27.280	74.140	34.9M
Non-displayed orders									
Marketable {far}	21.3M	0.003	0.010	0.099	11.8M	0.000	0.010	0.023	9.5M
{far, midpoint}	0.5M	0.007	0.013	0.070	0.4M	0.000	0.005	0.024	0.1M
Midpoint	1.2M	0.007	0.047	20.29	0.7M	0.010	5.11	20.04	0.5M
{near, midpoint}	0.9M	13.333	50.00	85.264	0.5M	5.985	18.571	44.524	0.4M
Passive {near}	9.5M	14.530	48.820	76.870	7.1M	6.133	19.263	46.590	2.4M
Displayed orders									
Marketable {far}	30.2M	0.006	0.010	41.100	7.7M	0.000	0.010	0.027	22.5M
{far, midpoint}	1.3M	0.006	0.010	0.034	0.6M	0.000	0.004	0.050	0.7M
Midpoint	4.6M	0.007	3.67	62.000	1.4M	0.001	2.186	14.390	3.2M
{near, midpoint}	8.0M	29.84	59.39	107.14	2.9M	2.370	13.210	41.220	5.1M
Passive {near}	101.9M	29.56	45.00	87.08	69.5M	8.520	28.029	77.274	32.4M
Panel B: Child PEG orders									
Marketable {far}	12.2M	0.006	0.030	2.000	10.9M	0.003	0.023	2.216	1.3M
Midpoint	88.2M	0.007	0.010	15.400	74.4M	0.006	0.023	3.866	14.2M
Passive {near}	30.4M	4.983	20.260	45.537	24.6M	3.560	11.283	28.634	5.8M

Table 5

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

The table contains parameter estimates from accelerated failure time models for limit and PEG orders under the generalized gamma distribution. The intercept captures marketable 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 to 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. Scale and shape parameters correspond to σ and ν in equation3 in section 5.2.

	Limit orders			PEG orders	
Intercept (Marketable)	1.599	1.603	1.601	4.330	4.388
	(0.002)	(0.002)	(0.002)	(0.007)	(0.007)
{far, midpoint}	1.628	1.610	1.628		
	(0.003)	(0.003)	(0.003)	-	-
Midpoint	3.489	3.566	3.490	-0.459	-0.457
	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)
{near, midpoint}	5.954	6.050	5.955		
	(0.001)	(0.001)	(0.001)	-	-
Passive	8.420	8.511	8.420	3.705	3.708
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)
Scaled order size	-0.036	-0.030	-0.040	0.010	0.010
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Book Asymmetry			-0.010		-0.010
	-	-	(0.000)	-	(0.000)
Buy indicator			-0.003		-0.101
	-	-	(0.001)	-	(0.001)
Display order indicator	-0.774		-0.775		
	(0.000)	-	(0.001)	-	-
Lit indicator		-0.825			
	-	(0.001)		-	-
Log (market cap)	-0.078	-0.076	-0.079	-0.033	-0.033
	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Log (volume)	-0.186	-0.186	-0.187	-0.185	-0.186
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\|R_{-5,0}\|$	-0.035	-0.035	-0.036	0.002	0.002
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Scale parameter (σ)	3.416	3.422	3.416	7.136	7.136
	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Shape parameter (ν)	-0.786	-0.788	-0.786	-3.141	-3.140
	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
N (censored)	101 M	101 M	101 M	101M	101M
N (uncensored)	68 M	68 M	68 M	19M	19M

Table 6

Price impact of unfilled and filled child orders

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 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 average of each of these variable for the 2012-2016 sample period. 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					
Marketable{far}	1.19	1.25	1.28	1.66	2.03	0.41	0.41	0.43	0.64	0.80
{far, midpoint}	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
{near, midpoint}	0.67	0.77	0.85	1.43	1.97	-0.91	-1.08	-1.27	-1.75	-2.02
Passive {near}	0.08	0.14	0.18	0.51	0.84	-1.04	-1.23	-1.44	-1.82	-2.01
					Std. Dev.					
Marketable{far}	2.05	2.28	2.47	3.80	5.02	1.40	1.67	1.91	3.17	4.24
{far, midpoint}	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
{near, midpoint}	1.47	1.82	2.07	3.51	4.69	2.39	2.81	3.13	4.32	5.29
Passive {near}	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					
Marketable{far}	77.33	74.88	72.90	62.23	54.67	80.00	76.89	74.29	62.44	54.67
{far, midpoint}	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
{near, midpoint}	67.88	62.89	59.39	43.74	34.67	51.19	42.74	36.80	25.80	20.24
Passive {near}	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					
Marketable{far}	21.93	23.43	24.56	30.76	34.77	10.39	11.77	13.22	20.08	24.29
{far, midpoint}	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
{near, midpoint}	28.69	32.04	34.16	45.11	51.94	8.47	12.57	14.04	18.39	21.36
Passive {near}	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					
Marketable{far}	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 {near}	0.04	0.06	0.07	0.17	0.31	-0.79	0.89	-1.00	-1.29	-1.45
					Std. Dev.					
Marketable{far}	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 {near}	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					
Marketable{far}	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 {near}	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					
Marketable{far}	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 {near}	1.82	3.95	5.67	14.09	20.18	2.99	4.99	6.06	10.18	13.45

Table 7

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 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 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 separate in three price aggressiveness categories. To conserve space, we only report statistics for marketable, midpoint and passive categories. All numbers are in basis points.

	Size	Display	Unfilled orders					Filled orders				
	Quartile		0.1	0.5	1	5	10	0.1	0.5	1	5	10
Panel A: Limit orders												
Marketable	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												
Marketable	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 8

Average slopes from daily price impact regression

We estimate daily cross-sectional regressions of child order price impact. The dependent variable (child price impact) is measured at τ=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 venues are undisplayed by definition. 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. $|Ret_{-5,0}|$ 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	1.215	1.360	0.388	0.490	0.581	0.688	0.093	0.110
(Marketable)	(0.017)	(0.028)	(0.008)	(0.009)	(0.022)	0.026)	(0.015)	(0.016)
{far, midpoint}	-0.634	-0.555	-0.646	-0.573	-	-	-	-
	0.028)	(0.031)	(0.028)	(0.031)				
Midpoint	-0.894	-0.912	-1.404	-1.404	-0.160	-0.199	-0.557	-0.549
	(0.019)	(0.020)	(0.050)	(0.032)	(0.025)	(0.026)	(0.015)	(0.016)
{near, midpoint}	-0.355	-0.403	-2.372	-2.407	-	-	-	-
	(0.025)	(0.025)	(0.025)	(0.026)				
Passive {near}	-1.110	-1.220	-2.458	-2.502	-0.508	-0.585	-1.357	-1.660
	(0.016)	(0.020)	(0.008)	(0.026)	(0.022)	(0.025)	(0.019)	(0.019)
Display	0.170	0.224	0.217	0.226	-	-	-	-
	(0.010)	(0.018)	(0.008)	(0.007)				
Scaled size	0.099	-	0.151	-	0.060	-	0.038	-
	(0.004)		(0.006)		(0.003)		(0.005)	
Book Asymmetry*Buy	-	0.002	-	0.001	-	0.003	-	0.001
		(0.000)		(0.000)		(0.000)		(0.000)
Display * Scaled size	0.005	-	-0.037	-	-	-	-	-
	(0.001)		(0.007)					
$\lvert Ret_{-5,0}\rvert$	0.104	0.117	-0.024	-0.017	0.058	0.063	-0.042	-0.041
	(0.001)	(0.014)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Buy	0.026	0.004	0.006	0.001	0.011	-0.002	0.006	0.015
	(0.012)	(0.014)	0.007)	(0.008)	(0.009)	(0.010)	(0.008)	(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 9

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, inside the spread (midpoint or otherwise), or marketable (far side of the spread). For this table, we restrict the analysis to 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 simple averages of the number of runs per parent, the number of child orders per run, the duration of the run in seconds, the run volume measured as the number of shares submitted by all child orders in the run, and the percentage of runs at the three aggressiveness categories. Panel B shows unconditional 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 marketable runs and zero otherwise. If the prior run contains marketable 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 at					
Passive	45.58	40.58	42.78	42.00	46.94
Inside spread	16.04	16.34	19.98	10.94	16.21
Marketable	38.37	43.06	37.23	47.04	36.84

Panel B: Transition matrices between runs

	All Algos			Algo A			Algo B			Algo C			Algo D		
	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_t
Passive (P_{t-1})	-	13.9	31.2	-	10.3	33.0	-	17.0	27.2	-	3.0	41.1	-	14.6	30.7
Inside (I_{t-1})	12.9	-	5.2	9.2	-	7.1	13.7	-	6.6	2.8	-	6.9	13.9	-	4.8
Marketable (M_{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	
	$Fill_{t-1}$	$SRet_{t-1}$	$Fill_{t-1}$	$SRet_{t-1}$	$Fill_{t-1}$	$SRet_{t-1}$	$Fill_{t-1}$	$SRet_{t-1}$
Aggressiveness of prior run								
Passive	0.524	0.005	0.022	0.001	0.205	0.006	-0.041	0.001
	(0.005)	(0.000)	(0.003)	(0.000)	(0.006)	(0.000)	(0.001)	(0.000)
	[7.59]	[0.09]	[0.50]	[0.03]	[1.07]	[0.04]	[0.83]	[0.01]
Inside	-0.067	-0.017	-0.374	-0.006	-0.017	-0.018	-1.359	-0.010
	(0.006)	(0.000)	(0.004)	(0.000)	(0.007)	(0.000)	(0.001)	(0.000)
	[1.65]	[0.44]	[-8.32]	[0.14]	[-0.35]	[0.01]	[26.95]	[0.19]
Marketable	0.063	-0.007	0.098	-0.001	0.366	-0.014	0.659	-0.000
	(0.05)	(0.004)	(0.004)	(0.000)	(0.004)	(0.000)	(0.001)	(0.001)
	[0.92]	[-0.11]	[1.71]	[-0.02]	[5.06]	[-0.18]	[11.21]	[0.00]

Table 9

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 marketable (far side of the spread, M)). For this table, we restrict the analysis to 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 simple averages of the number of runs per parent, the number of child orders per run, the duration of the run in seconds, the run volume measured as the number of shares submitted by all child orders in the run, and the percentage of runs at the three aggressiveness categories. Panel B shows unconditional 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 marketable runs and zero otherwise. If the prior run contains marketable 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 at					
Passive	45.58	40.58	42.78	42.00	46.94
Inside spread	16.04	16.34	19.98	10.94	16.21
Marketable	38.37	43.06	37.23	47.04	36.84

Panel B: Transition matrices between runs

	All Algos			Algo A			Algo B			Algo C			Algo D		
	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_t	P_t	I_t	M_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
M_{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		
	$Pr(M_t)$	$Pr(M_t)$	$Pr(P_t)$	$Pr(M_t)$	$Pr(M_t)$	$Pr(P_t)$	$Pr(M_t)$	$Pr(M_t)$	$Pr(P_t)$	$Pr(M_t)$	$Pr(M_t)$	$Pr(P_t)$
Run_{t-1}	P	I	M	P	I	M	P	I	M	P	I	M
$Fill_{t-1}$	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]
$SRet_{t-1}$	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]

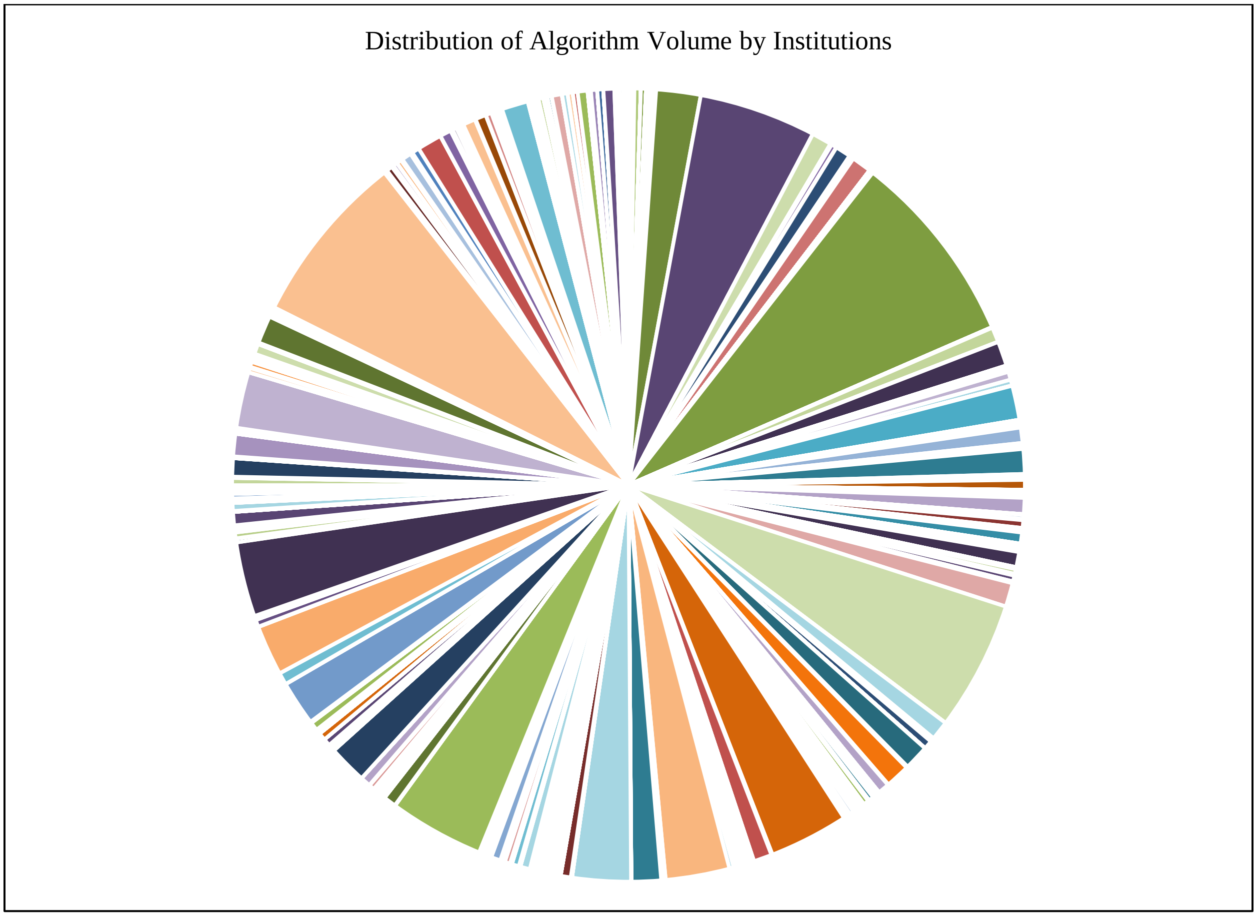

Figure 1: The pie chart shows the distribution of dollar value parent orders of each institution across all trading algorithms. Each institution is represented by a unique color.

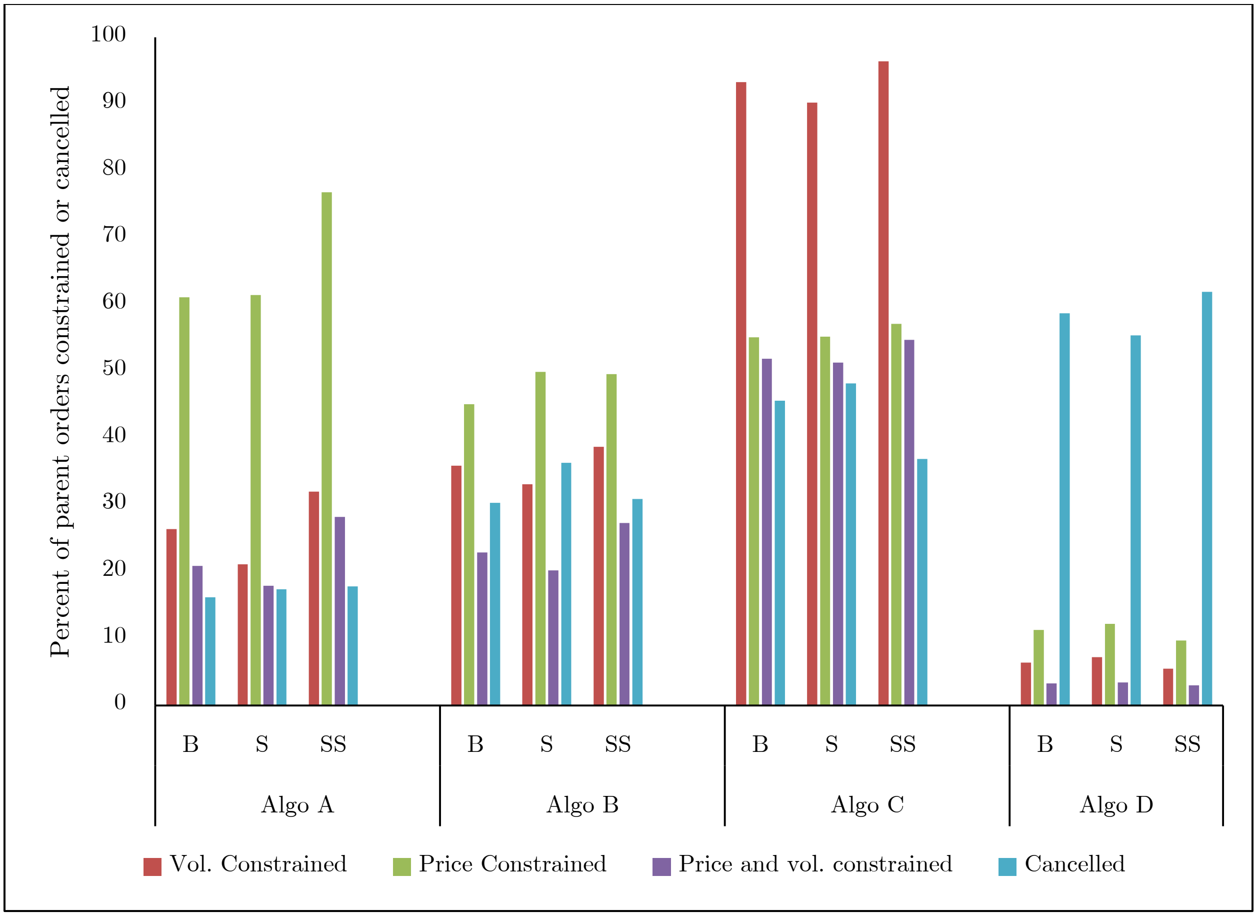

Figure 2: The chart shows the percent of parent orders constrained by price limits, volume limitations, price and volume constraints, and cancelled before completion. The statistics are shown separately for buys, sells, and short sales for each algorithm.