Who Uses Which Order Type and Why?*

Sida Li
University of Illinois, Urbana-Champaign

Mao Ye
University of Illinois, Urbana-Champaign and NBER

Miles Zheng
University of Illinois, Urbana-Champaign

Abstract

Clientele effects explain the proliferation of order types on U.S. stock exchanges. Market and plain limit orders lose money, indicating that informed traders use more complex orders. The most complex order types refuse to trade with the national best bid and offer (NBBO) if the NBBO appears on another exchange. Fees provide one explanation for non-routable orders because Reg NMS might route orders to worse prices after adjusting for routing fees. Non-routable orders also win speed races to capture quick profits and contain short-term information. However, all order types containing long-term information are routable and often tailored to corporate events.

Keywords: Order Types, Regulation NMS, Information, Make/take fees, Speed Competition

* We thank Jim Angel, Dan Bernhardt, Ekkehart Boehmer, Colin Clark, Carole Comerton-Forde, Nicolas Crouzet, Ian Dew-Becker, Joel Hasbrouck, Edwin Hu, Pankaj Jain, Phil Mackintosh, Bruce Mizrach, Dermot Murphy, Maureen O'Hara, Steven Poser, Elvira Sojli, Jeremy Stein, Kumar Venkataraman, Sunil Wahal, and seminar participants at the Louisiana State University, New York Stock Exchange, University of Illinois at Urbana-Champaign, Rutgers University, and Microstructure Online Seminars Asia Pacific for helpful discussions and suggestions. We thank the New York Stock Exchange for providing us with their proprietary SOD data for this paper. Ye acknowledges support from National Science Foundation grant 1838183 and the Extreme Science and Engineering Discovery Environment (XSEDE). Send correspondence to Mao Ye, the University of Illinois at Urbana-Champaign, 340 Wohlers Hall, 1206 South 6th Street, Champaign, IL, 61820. Email: maoye@illinois.edu.

Any investor needs to select an order type when trading on stock exchanges. Most relevant studies, however, abstract from order types or build on simple binary order types between market and limit orders. Using proprietary data from the New York Stock Exchange (NYSE), we find that traditional market and limit orders account for less than 10% of observed trading volume and they both lose money. Market orders, which are orders that do not specify prices, lose 26 bps at the 30-day horizon. Limit orders that specify prices but offer no further instructions lose 49.2 bps at the 30-day horizon. Therefore, informed traders, whether they are short- or long-term traders, should use more complicated order types. Understanding the diversity of order types and their origins therefore is important for comprehending the trading ecosystem.

We contribute to the literature by providing the first anatomy of exchange order types. We uncover a three-tiered ecosystem of order types. 1) The simplest order types lose money, which implies that sophisticated traders manage their orders. The extent of such management depends, however, on traders' information horizons. 2) The most sophisticated order types earn small and quick profits but contain no long-term information. 3) All order types that result in long-term profits are neither too simple nor too complex. Therefore, for orders other than plain market or limit orders, sophistication in short-term trading and in long-term investments are substitutes but not complements.

The most sophisticated order types share a counterintuitive feature: refusing the best price. Stocks of U.S.-listed firms can be traded on 13 U.S. stock exchanges. The 2005 Regulation National Market System (Reg NMS) rules direct U.S. stock exchanges to consolidate quotes, establish the national best bid and offer (NBBO), and to route an incoming order to the exchange that displays the NBBO. Under Reg NMS, each stock exchange serves only as an entry point to the same destination—the best price. Surprisingly, 57% of the orders in our sample refuse Reg NMS routing to the NBBO. This fact raises two questions: Why do these orders reject the best price? Who uses these non-routable orders?

We find that exchange fees serve as one driver of exchange-routing refusal. We find that exchanges often have to route displayed limit orders to worse prices after adjusting for fees to comply with Rule 610 of Reg NMS, which prohibits a displayed order in one exchange to lock or cross an existing quote in another exchange. Suppose that the best NYSE ask price is \$10.00 while the best NYSE bid price is \$9.98. A trader who submits a sell limit order at \$9.99 would improve the NYSE best ask price by one tick. If NASDAQ has a bid price of \$9.99, even if the bid is established a small fraction of a second earlier, the NYSE limit sell order at $9.99 locks the NASDAQ bid at $9.99, and Rule 610 would require the NYSE to route the limit sell order at $9.99 to take liquidity from NASDAQ. Routing unlocks the market, but it leads to a worse price for the NYSE limit sell order because the NYSE offers a rebate of 0.13 cents per share for orders that make liquidity and charges a fee of 0.30 cents per share for routing orders outside the exchange. Therefore, we find that more than 50% of non-marketable orders are attached with "do-not-ship (DNS)" instructions, which ask the NYSE to cancel an order if it locks or crosses quotes in another exchange. We find that DNS limit orders earn a small and quick profit of 1.34 bps after collecting the rebate but would lose 0.38 bps if they paid the routing fee.

Speed is another driver of exchange-routing refusal. Exchange routing can lead to three types of latencies: geographic, consolidation, and transmission (SEC (2018)). Latencies create incentives for speed-sensitive traders to route their own orders. One order type, Intermarket Sweep Order (ISO), asks exchanges not to check prices on other exchanges. Instead, traders comply with Reg NMS by simultaneously sending ISOs to all other exchanges that offer better prices. We find that ISOs pay transaction costs of 3.62 bps upon execution but make profits of 0.54 bps one second after execution. Therefore, ISOs execute against stale quotes, and speed is the key to snipe those stale quotes (Budish, Cramton, and Shim (2015)). As ISOs often snipe stale quotes, they are informative on the short horizon. We find that ISOs contribute to more than 90% of price discovery on a trade-by-trade basis.

We also find that non-routable orders win speed races in making liquidity. First, DNS limit orders successfully cancel stale quotes 42% of the time, but routable limit orders rarely escape sniping. The adverse selection cost of DNS limit orders is 0.26 bps lower than routable limit orders. Second, DNS limit orders are more likely to establish time priority in liquidity provision. As orders at the front queue position are more profitable than orders at the end of the queue (Li, Wang, and Ye (2020)), we find that the return of DNS limit orders is 0.76 bps higher than routable limit orders. The ability to quickly add new quotes and cancel stale quotes makes DNS limit orders experience a 1.02 bps higher return than routable limit orders 30 seconds after execution.

Although DNS orders and ISOs predict tick-by-tick price movements and typically make profits in the short run, they would not earn any positive returns if they were to hold their positions beyond the end of the trading day. We find all order types that have positive returns beyond a day accept exchange routing. These results indicate the following clientele segmentation: traders who have short-term information use complicated non-routable orders to control their order flows; traders who have long-term information manage their orders to some extent, but they delegate the search for the best price to the exchange. One order type that contains long-term information is the buy-minus-zero-plus (BMZP) order, which directs the NYSE to buy a stock at a price that is not higher than the most recent market price. We find that BMZP orders realize a 30-day return of 7.06% and are associated with share repurchases by issuers. The NYSE states that it designs BMZP orders to comply with Rule 10b-18.[1] To prevent firms from inflating their share prices, Rule 10b-18 discourages firms from buying at prices that are higher than the highest independent bids or the most recent transaction prices. Therefore, BMZP orders have high returns because they cater to repurchasing firms, who are informed traders in their own stocks (Dittmar and Field (2015)). BMZP orders are routable, probably because firms and their brokers care more about filling their orders than about the fees or the speed of the fill.

Our most vivid head-to-head comparison involves order types that differ only in their routing decisions. Our sample includes two variations of limit orders that hide their sizes: regular reserve limit orders and DNS reserve limit orders. Regular reserve limit orders accept routing offered by the exchange, while DNS reserve limit orders refuse exchange routing. We find that DNS reserve limit orders generate a return that is 0.22 bps higher than those generated by regular reserve orders upon execution. In contrast, the return on regular reserve limit orders increases over time, rising to 40.40 bps at the 30-day horizon; DNS reserve limit orders, however, lose 8.23 bps at the 30-day horizon. Our results suggest that reserve orders are designed to hide long-term information and they accept exchange routing once another exchange offers an acceptable price. DNS reserve limit orders, however, are used by fee-sensitive traders who refuse the matching opportunity because they may fail to break even after paying routing fees.

Most studies focus on plain market and limit orders.[2] Bessembinder and Venkataraman (2004) and Bessembinder, Panayides, and Venkataraman (2009) analyze hidden orders before the

[1] SEC File No. SR-NYSE-2016-59, available at https://www.sec.gov/rules/sro/nyse/2016/34-78679.pdf.

[2] Harris and Hasbrouck (1996) compare limit order performance with market order performance on the NYSE. Biais, Hillion, and Spatt (1995) use Paris Bourse exchange data; Griffith et al. (2000) use Toronto exchange data; and Hollifield, Miller, and Sandas (2004) use Swiss exchange data to analyze order placement strategies for limit and market orders. Kelley and Tetlock (2013) analyze the informational roles of retail limit and retail market orders. Theoretical limit-order-book studies also focus on analyzing the tradeoff between market and limit orders and find that traders submit limit orders when they are patient (e.g., Kaniel and Liu (2006)), when pick-off risks are low (e.g., Hollifield, Miller, and Sandas (2004)), and when the market is thinner (e.g., Biais, Hillion, and Spatt (1995); Parlour (1998)). Parlour and Seppi (2008) provide an excellent survey of the limit order book literature.

era of algorithmic trading. We find a dramatic contrast between routable and non-routable reserve limit orders. Although the flow of regular reserve limit orders can predict long-term returns, DNS reserve limit orders cannot. This result indicates the importance of decomposing order instructions: two order types that differ in only one instruction can have completely different properties. Using publicly available TAQ data, Chakravarty et al. (2014) find that ISOs account for higher information shares than trades led by non-ISOs. Our paper uses more detailed order-level data and finds that the information contained in ISOs does not last long. All orders that contain long-term information are routable. Our results suggest that we should differentiate informed trading at the short-term execution horizon from informed trading at the long-term investment horizon.

Crouzet, Dew-Becker, and Nathanson (2020) model the interaction between traders at varying horizons. Because each trader has a fixed budget for learning information (Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016)), traders choose to separate in the frequency domain. Some traders acquire short-term information and become uninformed in the long term, whereas others acquire long-term information and become uninformed in the short term. Our results regarding order types provide the first empirical support to Crouzet, Dew-Becker, and Nathanson (2020). Non-routable order users are informed at the execution horizon but not at the investment horizon; however, long-term informed traders do not excel at the short horizon. Both short-term informed traders and long-term informed traders in Crouzet, Dew-Becker, and Nathanson (2020) make money because noise traders bear the cost. Our results show that users of the simplest order types may be such noise traders.

Our paper, combined with studies examining other rules in Reg NMS, provides explanations for the landscape of U.S. market structure following the implementation of Reg NMS. Chao, Yao, and Ye (2019) show that Rule 612 of Reg NMS, which imposes a minimum price variation of 1 cent per share, leads to the make/take fee pricing model and the proliferation of exchanges that charge varying fees. These authors show that the make and take fees should be different in equilibrium. Therefore, the net price to make or take liquidity should be different after adjusting for fees. Routing further widens the gap. When an exchange take liquidity on behalf its traders from another exchange, the former tends to add a mark up to the take fee charged by the latter. In this paper, we show that fee differences and Rule 610 generate order types that refuse routing. Chao, Yao, and Ye (2019) show that the tick size leads to multiple exchange operators, with each offering multiple exchanges. We find that the destination of an exchange routing tends to be another exchange with a common owner. Therefore, traders who are sophisticated in execution may prefer to route orders by themselves. Yao and Ye (2018) find that Rule 612 also leads to speed races to capture rents created by tick size, and we find that DNS limit orders tend to win such speed races. We find that routing that complies with Rule 610 leads to a worse price after adjusting for fees, but routing that complies with Rule 611 leads to price improvement. However, we show that ISO, an order type created by the Securities and Exchange Commission (SEC) to exempt from Rule 611, deviate from its original purpose. SEC created ISOs to provide institutional investors with immediate access to liquidity at multiple price levels, in multiple markets, to fill large block trades with parallel order submissions (Chakravarty et al. (2012)). Yet we find that the average size of an ISO is only 265.63 shares and only 0.53% of ISOs sweep multiple price levels, both of which are substantially lower than other liquidity taking orders. We also find that the main driver of ISOs is speed races for sniping stale quotes. Taken together, our paper and previous studies explain how Reg NMS generates four features in current market: high-frequency trading, market fragmentation, exchanges with heterogenous make/take fees, and proliferations of order types.

Reg NMS is an attempt to link fragmented markets, but it is extremely difficult to empirically examine the market microstructure linkages due to data limitations. Our data come from the NYSE, but routing from it allows us to provide the first glimpse to market linkages, which adds significant insight to the literature. Exchange linkage blurs key concepts in market microstructure, such as the distinctions between marketable and non-marketable orders. Reg NMS has two main goals: to promote the use of displayed non-marketable orders and to enable best execution price for marketable orders.[3] However, linked exchanges create a gray area: an order that displays liquidity in one exchange may become marketable in another exchange. To comply with SEC Rule 610, exchanges need to route these orders to a worse price after adjusting for fees. In practice, Rule 610 creates an incentive for traders to hide their quotes. Because orders that completely hide their quotes cannot lock the displayed market, we find that 99.95% of these orders result in a price improvement. Further, according to Rule 610, the exchange that locks the market is the exchange that displays a quote later. Therefore, exchange latency plays a role in defining which exchange locks the market. For example, we find that a reserve order can be routed to

[3] SEC File No. S7-10-04, available at https://www.sec.gov/rules/final/34-51808.pdf.

another exchange due to a small latency to refill its displayed quotes.[4] Therefore, a slow exchange tend to pay take fee to a fast exchange. Also, by routing an order out, a slow exchange increase the trading volume and market share of a fast exchange. Therefore, the presence of exchange linkages and SEC Rule 610 is one possible reason why exchanges compete on speed and adopt a price-time priority to reward traders who establish fast quotes, even though the SEC allows exchange to adopt alternative priority rule such as price-size priority and frequent batch auctions.[5]

We find that non-routable orders win three types of speed races modeled by Li, Wang, and Ye (2020): sniping stale quotes, canceling stale quotes, and achieving front queue position. DNS limit orders are more likely to win the race to achieve the front queue position when price competition is constrained by discrete pricing, and they are also more likely to win the race to cancel stale quotes when stock prices jump. ISOs tend to win the race to snipe stale quotes (Budish, Cramton, and Shim (2015); Li, Wang, and Ye (2020)). We find that it is rare for routable orders to participate in these speed races; when they do participate, they are unlikely to win.

Although we focus on the economics of NYSE order types in this paper, it also helps us to understand order types in other U.S. exchanges, as their order types respond to similar economic issues raised by market structure post Reg NMS. We provide a brief discussion on cross-exchange variation and time series evolution of order types in Section VI. The rest of the paper is organized as follows. In Section I, we describe the proprietary NYSE data. In Section II, we discuss the data and provide a taxonomy of order types based on their economic functions. In Section III, we present the three-tiered world of order types. In Section IV, we show that fees serve as one explanation for refusing exchange routing. In Section V, we show that speed provides another explanation of routing refusal. We conclude in Section VII.

I. Data and Taxonomy of Order Types

We use the proprietary NYSE System Order Database (SOD) to study order types. The sample period is from January 1, 2010 to March 1, 2011, and we have a stratified sample of 109 stocks.

[4] Consider a reserve order of 5,000 shares that display only 100 shares and refill another 100 shares after the execution of the first 100 shares. After the execution and before the refill of another 100 shares, a displayed order on the opposite side occurs in another exchange, so the refilled shares need to take liquidity from outside.

[5] See SEC memo, available at https://www.sec.gov/spotlight/emsac/memo-rule-611-regulation-nms.pdf. The NASDAQ PSX exchange tried price-size priority in 2010, but it attracted little volume and moved back to price-time priority in 2012. The PSX revived price-size priority in 2014 but once again attracted little volume. The NYSE has parity among traders who are not the first to establish a quote but reward the first one who sets the quote.

Panel A of Table 1 shows our sample selection criteria. Starting with all securities listed on the NYSE in December 2009, we apply standard filters (O'Hara and Ye (2011)) to remove non-common equities, dual-class shares, real estate investment trusts (REITs), and the common stocks of non-U.S. companies. We also exclude stocks with prices below $5.00. We then rank the remaining 1,086 stocks based on their trading volume and pick every tenth stock from the stock with the highest trading volume. In Panel B, we list the tickers of our final sample of 109 stocks. In Panel C, we summarize the characteristics of these 109 stocks.

[Insert Table 1 about here]

Our data contain 3.2 billion records of messages that include order additions, order updates, order cancelations, or order executions during regular trading hours, and these order instructions result in 152 million transaction records. Each message has a timestamp, order size, execution size, price, buy/sell/short-sell indicator, route-out indicator, route-out destination, make/take indicator, and order ID. If an order involves multiple trading counterparties or multiple executions, our data contain one observation for each counterparty and each execution. The data do not include trader IDs. Therefore, instead of asking whether a trader has information, we ask whether users of specific order types have information. Even if we cannot "see" the driver of a vehicle, we can obtain a noisy signal of the driver by observing the car.

We identify order types using their price condition, time-in-force (TIF), and special order instructions (SOIs). Each price condition, TIF or SOI adds a function to an order type. We start our introduction by describing these basic functions because they are the building blocks of an order type. A complex order may be designed to achieve multiple functions. The best way to understand the economic drivers of order types is to compare two order types that are otherwise identical but differ in only one function. This decomposition also helps us to understand order types in other exchanges. Other U.S. exchanges, such as the NASDAQ and BATS, offer different order types, but their order types offer similar functions because other U.S. exchanges aim to address similar economic issues raised by market structure post Reg NMS. Some order types of other exchanges are just a different mix and match for the same basic functions from NYSE order types, and some other order types are NYSE order types under different names. The decomposition also helps us to understand the evolution and proliferation of order types within the NYSE. Our economic analysis below shows that some order type may not provide the best solution for an economic problem, and some economic issues may not have an associated order type at all.

Therefore, order types evolve along two long lines: 1) a new order type supersedes an old order type if the new order type provides a better solution for an economic problem; and 2) exchanges may span the space of order types by designing new order types. We provide a brief overview of order types across exchanges and their evolution in Section VI.

A. Price condition

Based on price conditions, we can categorize orders into three types: market orders, limit orders, stop orders, and orders with trailing prices. A market order does not have a specified price and it is used to buy or sell a stated amount of a security that is to be traded at the best price obtainable. A limit order has a specified price, and it is used to buy or sell a stated amount of a security at a specified price or better. For example, a limit order to buy 100 shares with a limit price at $100 would buy a stock if its price is $100 or below. A STOP order, also known as a stop-loss order, is designed to buy or sell a stated amount of a security once the price of a stock becomes worse than the stop price. Investors generally use a stop order to limit a loss or to protect a profit on a stock that they own.

Orders with trailing prices can be further categorized based on their functions. A primary pegged order is a limit order whose working price dynamically follows the current market price, where buy orders peg to the bid, and sell orders peg to the ask. For example, when the protected bid increases, the NYSE will increase the price of a primary pegged order to buy to follow the market price trend, and vice versa. A primary pegged order will be rejected on arrival, or canceled when resting, if there is no reference price available. Primary pegged orders are available only to floor brokers.

Orders with trailing prices can be even more aggressive. A “buy-minus-zero-plus” order is a buy order, and its limit price is the higher of the current bid price *and* the last trade price. Thus, a BMZP order cannot be the initiator to push up the price by improving the best bid or taking liquidity from the ask side. However, if another trader takes liquidity from the ask size, the BMZP order will follow up, take the remaining liquidity at the ask, route out to take liquidity, and (if the BMZP order size is large enough) make liquidity at the new bid price. This behavior dovetails the requirements of Rule 10b-18 while maximizing the chance of order execution.

B. Time in force

TIF can take one of the following four values: IOC (immediate-or-cancel), DAY (expires at the end of the day, AUC (available only in the open/reopen/close auctions), or GTC (Good Til Canceled). An IOC instruction cancels an order if it fails to execute immediately. All IOC orders have a limit price, but they only accept the limit price and refuse to offer other traders the limit price. Therefore, IOC orders are more similar to market orders than to limit orders. DAY orders have expiration times at the end of trading days, and they often include additional SOIs to form more complex order types. We do not find any salient pattern for AUC orders, which is consistent with our clientele story. Because a market's open and close involves a diverse body of traders, it is hard to infer information from the AUC order type. For the sake of brevity, we exclude orders with AUC instructions from our analysis.

GTC orders have two variations. We describe GTC orders without any further instructions as plain limit orders because they do not impose an expiration time or any other further instructions. Such orders remain on the limit order book until execution or cancellation. The NYSE will help the GTC order submitter to adjust its limit price if the underlying firm pays dividends, splits, or reverse splits its stock. A Do-Not-Reduce order further instructs the NYSE not to adjust its limit price for dividends.

C. Routing decisions

Rules 610 and 611 of Reg NMS aim to link fragmented markets together through exchange routing. Two order instructions, DNS and ISO, are employed to reject exchange routing.

Rule 610 requires each national securities exchange and national securities association to avoid displaying quotations that lock or cross any protected quotation in a National Market System (NMS) stock. Suppose that NBO is $5.01 and the NBB is $5.00, while the best offer at the NYSE is $5.02 and the best bid at the NYSE is $5.00. A new limit buy order at $5.01 would improve the best bid of the NYSE to $5.01. Rule 610, however, forbids such improvement because the bid price of $5.01 would lock the NBO of $5.01. Rule 610 then would route the new limit buy order at $5.01 to the NBO to resolve the locked market. A trader can refuse such routing by canceling the quote. Rule 610 is the main driver of DNS limit orders in our sample.

Rule 611 prohibits the execution of trades on one venue at prices that are inferior to quotes displayed on another venue. Therefore, an exchange has the obligation to route an incoming marketable order to the exchange that offers the NBBO. A DNS tag on the order refuses exchange routing to the best price by canceling the order. If a trader attaches an ISO instruction to her order, she accepts the obligation to check protected quotes from all exchanges. The exchange to which she submits her order can then directly execute the order against *local* orders at any price equal to or better than the limit price without being obligated to check protected quotes on other exchanges. The sender of an ISO must simultaneously send ISOs to other exchanges to execute against the full displayed size of any protected quotes. Rule 611 is the main driver of two major order types in our sample: DNS IOC orders and ISOs.[6]

D. Display decisions

When a trader chooses a limit order,[7] she can use a reserve order to partially or fully hide her trading interests. When she hides part of the order size, say, 100 shares, the NYSE will automatically replenish the order with another 100 shares once the original 100 shares are executed. Rule 610 would route the order out if the displayed part locks or crosses the market. Surprisingly, we find that the displayed part can lock the market not only upon arrival but also during the refill process. Traders can add DNS instructions to a reserve order so the NYSE will cancel the order if it needs to be routed out.

We find that the NYSE seldom routes out a reserve order with zero displayed size when the order locks the market, because Rule 610 does not apply to hidden quotes. As a consequence, routing leads to price improvement for 99.95% orders that are completely hidden, because such routings are consequences of Rule 611, not Rule 610. Consider the following example. Suppose there is a NYSE hidden order to sell at $5.00, which establishes a hidden ask price of $5.00. If the best bid price at the away exchange is $5.01, the displayed bid price crosses the hidden ask price. Such a cross does not violate Rule 610 because the ask price is hidden. However, if a market order in the NYSE hits the hidden sell order, the hidden sell order executes at $5.00, which trades through the bid price in the away exchange of $5.01.

II. Summary Statistics

Table 2 provides the summary statistics of order types in our sample. Some combinations of order

[6] DNS IOC orders are also known as NMS IOC, SEC IOC, or SOC orders.
[7] Again, a limit order must have a TIF equaling to a DAY or GTC order, because an IOC is not designed to provide liquidity.

instructions do not exist in the data. We do not find any SOIs for market orders, probably because traders who do not care about the execution price also do not care about other execution conditions. We also do not find any further instructions for stop orders and BMZP orders. Pegged orders have two versions: one is routable and the other has a DNS tag. We combine these two types because both types of pegged orders show similar patterns. Indeed, because pegged orders in our sample aim to peg the best bid or ask price, it is rare for a pegged order to be routed out even if it does not include a DNS tag. We find that only 0.02% of routable pegged orders execute in an exchange other than the NYSE.

[Insert Table 2 about here]

Our sample includes three types of IOC orders: plain IOC orders, ISOs, and DNS IOC orders. Plain IOC orders can take liquidity in any exchange, while ISOs and DNS IOC orders take liquidity only in the NYSE. The orders that participate in open and close auctions contribute to about 18% of share volume, but we exclude them in our analysis because we do not find any salient patterns for auction orders. A DNS limit order adds a do-not-ship instruction to a DAY limit order. A reserve limit order adds a non-display instruction to a DAY limit order. A DNS reserve limit order adds both do-not-ship and non-display instructions to a DAY limit order. We omit the DAY label for a DNS limit order, a reserve limit order, and a DNS reserve limit order for brevity.

A. Market share of order types

Table 2 presents the market share of each order type. As each trade has two sides, we present the market share based on double-counted volume so that the market share from all order types sums to 100%. We find that non-routable orders dominate share volume across four types of orders that only take liquidity. ISOs account for 15.00% of total share volume and DNS IOC orders account for 10.87% of share volume. Market orders and plain IOC orders, however, account for only 4.29% and 2.72% of share volume. Non-routable orders also dominate orders that aim to provide liquidity. DNS limit orders have a market share that is similar to DAY limit orders (19.50% vs. 23.10%).

In the Glosten and Milgrom (1985) and Kyle (1985) frameworks, all traders excepting market makers use only market orders. We find that market orders account for less than 5% of the share volume, and the number is still as low as 7% even if we consider plain IOC orders as market orders. The majority of the share volume comes from order types that are outside the reach of retail traders, as stop orders only account for 0.12% of share volume, and DAY limit orders account for

23.10% of share volume. Surprisingly, the two types of pegged orders account for less than 0.2% of share volume.

B. Execution condition for order types

One way to understand an order type is to examine the condition under which the order is executed. An order executes under three possibilities: providing liquidity on the NYSE, taking liquidity from the NYSE, or taking liquidity from other exchanges. These three types of execution incur different fees, and Rule 610 plays a central role in the differences between these fees. Rule 610 mandates a cap for a take fee of 0.3 cents per share. In our sample period, some exchanges, such as the NASDAQ and BATS X, charge a take fee of 0.3 cents per share. When the NYSE routes an order to another exchange, it takes the liquidity from the exchange; the NYSE charges a routing fee of 0.3 cents per share so that it does not lose money by routing an order outside. The NYSE charges a take fee of 0.21 cents per share; the take fee imposes a cap for the size of the rebate to liquidity-making orders. The size of the rebate to make liquidity in the NYSE is 0.13 cents; the difference between the take fee and the make rebate, 0.08 cents per share, is the NYSE's profit. The difference between the make rebate and routing fee in NYSE is 0.43 cents. As exchanges that charge a 0.30 cent take fee are able to give more generous rebates to liquidity makers, the differences between the make and routing fees among all exchanges are close to twice the size of the take fee cap.

Table 2 provides the market share of three types of execution: take local liquidity, route, and make liquidity. Pegged orders and BMZP orders include all three types of executions, although most of the execution comes from making liquidity at the NYSE. Particularly, pegged orders execute 99.28% of their share volume from making liquidity at the NYSE. Plain limit orders, DAY limit orders, reserve limit orders, and Do-Not-Reduce orders also include three types of executions. DAY limit orders provide liquidity for 55.94% of their executed volume and take liquidity from the NYSE for 29.78% of their executed volume, while the NYSE routing contributes to 14.28% of the volume of DAY limit orders. Reserve limit orders are slightly more likely to take liquidity from the NYSE (31.09%) and less likely to take liquidity outside the NYSE (16.30%), and they make liquidity for 52.62% of their executed volumes.

DNS limit orders and DNS reserve limit orders are much more likely than routable limit orders to make liquidity, not only because they refuse to take liquidity from outside but also because they take less liquidity within the NYSE. For example, DNS limit orders take liquidity only for 11.43% of their executed volume, and the percentage is as low as 5.00% for DNS reserve limit orders. These results suggest that one factor driving exchange-routing refusal should be the preference for making but not to taking liquidity.

Market orders exhibit the highest routing rate among all 13 order types (33.64%). STOP has similar route-out rate (31.65%) because STOP orders, once triggered, becomes a market order. The routing rate is as low as 1.93% for IOC orders. By default, 100% of the executed volume for DNS IOC orders and ISOs originate in taking liquidity from the NYSE because they refuse to make liquidity on the NYSE and refuse to take liquidity from other exchanges.

For each routing destination, Table 3 shows its market share in routed volume and its market share of the total trading volume for 109 stocks.[8] The top destination for NYSE routing is NYSE Arca. Across all volume routed out, NYSE Arca accounts for 53.9% of the market share. This number is almost double NYSE Arca's market share of the total trading volume (27.7%). Two economic forces can explain the NYSE's preference for routing to NYSE Arca. First, NYSE Arca has the same owner as the NYSE. Second, NYSE Arca and the NYSE use the same data center in Mahwah, New Jersey. Routing to NYSE Arca, therefore, has the lowest latency.

[Insert Table 3 about here]

III. The Three-tiered Ecosystem of Order Types

As the first paper to study the diversity of exchange order types, it is impossible to provide pairwise comparison between tens of order types in one paper. Our initial step is to provide some categorization of order types. One way to categorize order types is through economic reasoning, and the other way is based on empirical analysis on the property of order types. Both methods point out the same three-tiered ecosystem of order types. We begin this section with a brief discussion on categorization based on economic reasoning, and then provide the empirical support for the three-tiered ecosystem of order types.

The simplest order types are plain market orders, plain limit orders, and stop orders. These are order types in a regular brokerage account and in academic textbooks. Most theoretical and empirical work on order types discuss the choice between market and limit orders (See Seppi and Parlour (2008) for the survey) and Easley and O'Hara (1991) analyze stop orders. For other order types, we distinguish between orders that are non-routable and routable. Non-routable orders must

[8] We compute the total trading volume in each exchange from the TAQ data.

come from traders with the most sophisticated execution technology. For example, users of ISOs take over the compliance for Reg NMS from the exchange, and they should have the same or lower latency than the exchanges. We call order types that are not the simplest three order types but are routable orders of medium complexity. These orders manage their execution to some extent but defer the search for best price to the NYSE.[9]

Next, we present the evidence for the three-tiered ecosystem of order types based on our analysis of the return of order types. Because we do not have trader IDs, our results contain two type of noise. First, each order type can be used by multiple types of traders. For example, both retail and institutional traders can user DAY limit orders. We categorize DAY limit orders as having medium complexity because they include expiration time and retail traders cannot account for the large share volume of DAY limit orders.[10] This simplification tends to bias the estimate of return differences between order types toward zero, just like studies using the NASDAQ HFT dataset (Brogaard, Hendershott, and Riordan (2014); Yao and Ye (2018)), where the dataset includes some HFTs in the non-HFTs group. Therefore, the significant differences in returns between order types show the robustness of our results. Second, we measure returns based on a fixed horizon. As traders may close their positions before or after that horizon, we may find an order type has no return even if users earn positive returns. This problem becomes more serious as the horizon increases, or when traders have more flexibility to choose when they close their position. Again, we find certain order types have significant returns over a long-term horizon despite such noises, which supports the hypothesis that users of certain order types have long-term information.

A. Short-term return

Figure 1 presents the average returns for each order type for horizons less than five minutes after execution. The benchmark for measuring an order's return is the bid–ask midpoint of the NBBO at each horizon. Panel A presents the results for four liquidity-taking order types. Horizon 0 indicates that the benchmark price is the midpoint at the time of execution. All four types of orders

---

[9] Our conversation with an industry expert also supports the categorization based on routing decisions. Non-routable orders usually come from computer algorithms that refuse routing, and these algorithms tend to be faster proprietary algorithms. Routable orders usually come from humans or computer algorithms that accept routing, and these algorithms tend to be slower agency algorithms.

[10] Kelley and Tetlock (2013) find that retail orders only account for ~6.9% of the total market share.

have negative returns at time 0 because they pay the bid–ask spread upon execution. Yet market orders pay the highest average transaction cost of 4.51 bps, implying that market order users probably deploy the least sophisticated execution strategy. The routable IOC orders have a lower average transaction cost of 4.32 bps, whereas non-routable IOC orders have the lowest average transaction costs, 3.62 bps for ISOs and 3.72 bps for DNS IOC orders.

[Insert Figure 1 about here]

Figure 1 shows that ISOs and DNS IOC orders break even the next second after execution, whereas the returns on plain IOC orders and market orders remain negative. Therefore, the midpoint prices immediately cross the execution prices within one second after the execution of ISOs and DNS IOC orders (i.e., the midpoint prices increase above the execution prices after a buy ISO or DNS IOC order or the midpoint prices decrease below the execution price after a sell ISO or DNS IOC order). Therefore, ISOs and DNS IOC orders execute against stale quotes.

Panel B of Figure 1 shows that non-routable liquidity-making orders earn higher returns than routable orders in the short term, although the difference is less dramatic than the difference between non-routable and routable liquidity-taking orders. There are three reasons for why the difference is less dramatic. First, liquidity-taking orders can immediately seize profit opportunities, whereas liquidity-making orders need to wait for other liquidity takers to realize these opportunities. Second, we show below that the differences within the same stock is much larger, because liquid stocks (stocks with lower spreads) tend to have more non-routable orders as a consequence of more intense competition in liquidity provision. Third, Figure 1 presents the results before fees. Making liquidity on the NYSE has a cost that is 0.43 cents lower than taking liquidity from an outside exchange, while taking liquidity from the NYSE costs 0.09 cents less (0.30 routing fee – 0.21 take fee) than taking liquidity from outside. As fees lead to a greater cost reduction for liquidity-making non-routable orders than for liquidity-taking non-routable orders, it is natural that the difference between routable order returns and non-routable order returns is smaller before fees for liquidity-making orders.

As can be seen in Panel C of Figure 1, all limit orders without a specified expiry date (i.e., plain limit, Do-Not-Reduce, and STOP orders) continue to lose money after execution. This is not surprising because these order types are all naïve in order execution. Panel D shows that the orders with a trailing limit price (i.e., Pegged and BMZP orders) earn returns of around 0 within the 5-minute horizon.

B, Long-term returns

Although non-routable orders contain short-term information, Figure 2 shows that all orders that contain long-term information are routable. We present the results based on raw returns, and the returns after adjusting factors are similar and available upon request. DAY limit orders lose money until the end of the day but realize a positive return of 1.59 bps three days after execution and 3.50 bps 30 days after execution. The return on DNS limit orders, however, are close to zero and sometimes negative. We find that reserve limit orders realize a return of 3.57 bps three days after execution, a number that gradually increases to 40.40 bps 30 days after execution. The return on DNS reserve limit orders gradually decreases in three days, becoming negative in 30 days. By comparing non-routable orders with their routable cousins, we find that the routable versions contain more long-term information even though they lose money in the short run.

[Insert Figure 2 about here]

The champion for long-term returns is BMZP orders, which are routable. BMZP orders realize the highest 30-day return (706 bps). Note that our results do not mean that any trader can profit from a BMZP order but that traders who use a BMZP order are more likely to be informed. Because BMZP orders comply with Rule 10b-18, we find that they are associated with share repurchases by firms. For firm-months that have BMZP orders, the firm has an ongoing repurchase program 81% of the time, but for firm-months without BMZP orders, the figure is only 35%. For the remaining 19% of firm-months in which no BMZP orders come from repurchase programs, BMZP orders lose 0.51% at the monthly horizon. These return results are driven by the users of an order type, not by the order type itself. The order type provides inference for the order users, albeit a noisy one.

BMZP orders are all routable, probably because their main users, repurchasing firms, care more about filling the orders and less about in which exchange the order is filled. One indication is the high fill rate of BMZP orders. BMZP orders in our sample have a fill rate as high as 45.88%, second only to market orders and pegged orders. Note that firms face constraints in using market orders (Ye, Zheng, and Li (2020)), because market orders execute at the ask price, and a market order repurchase execute at ask price to lose the safe harbor from Rule 10-b-18 if its price is higher than the last sale price.

The other type of long-term informed order also relates to corporate payouts. We find Do

Not Reduce orders have a 30-day return of 156.67 bps. Both Do Not Reduce and BMZP orders do not have a large share volume, but their higher returns mean that they are economically significant. Traders with fundamental information might be the minority in the population, but they achieve high return probably because they are the minority.

Figure 1 in the previous section shows that GTC leads to poor execution, as stop orders, plain limit orders and Do Not Reduce orders with GTC instructions all lose money at execution horizon. Figure 2 shows that stop orders and plain limit orders lose about 50 bps in 30-day horizon but Do Not Reduce orders have strong positive returns. Therefore, it is possible that both long-term informed traders and naïve traders do not deploy sophisticated execution strategy. One possibility is that naïve traders are not able to deploy sophisticated execution strategies, whereas some long-term informed traders have less incentive to deploy sophisticated execution strategies.

IV. The First Driver of NBBO Refusal: Fees

As the majority of orders refuse the NBBO, our understanding of the trading ecosystem would be highly incomplete without analyzing the economic drivers behind the incentive to refuse the best price. In this section, we show that Rule 610 is one driver to exchange-routing refusal. In Subsection IV.A., we show that Rule 610 leads to routing without price improvement. In order to prevent a locked or crossed market, an exchange often routes non-marketable orders to the same price on another exchange. Considering that the routing fee is higher than the make rebate, we show that routing to comply with Rule 610 often leads to worse prices. In Subsection V.B, we show that the 0.43-cent difference in fees makes a significant difference: DNS limit orders earn positive returns after collecting rebates but lose money after paying routing fees.

A. Paying fees but receiving no price improvement

The fee difference between taking liquidity outside and making liquidity on the NYSE is less than a price improvement. Why do traders avoid paying 0.43 cents if routing can lead to a price improvement? Panel A of Table 4 shows that for non-marketable orders, more than 77.49% of DAY limit orders are routed to an outside exchange without a price improvement, and the figure is as high as 79.51% for reserve limit orders.[11]

[Insert Table 4 about here]

As DNS limit orders and DNS reserve limit orders are never routed outside, we do not know their price improvements. However, Rule 610 implies that an order that improves the current price by one tick will see little if any price improvement when it is routed out. Suppose that the NYSE best ask is $10.00 and the best bid is $9.97. A buy limit order can improve the NYSE best bid to $9.98. However, under Rule 610, the exchange will route the order outside if an away exchange has an ask price of $9.98 or lower because the buy limit order submitted at $9.98 is marketable at the national best ask price. The buy limit order can obtain price improvement only if the ask price at other exchanges is $9.97 or lower. Such an ask price should not exist because Rule 610 would require the ask price of $9.97 or lower to transact with the NYSE bid of $9.97 to resolve the cross or locked market. Therefore, when the buy limit order is routed outside, it most likely takes liquidity at $9.98, experiences no price improvement, yet pays the routing fee and loses potential rebates from making liquidity on the NYSE.

In Panel B of Table 4, we show that DNS limit orders are more likely to improve the best bid and ask on the NYSE than DAY limit orders (24.03% vs. 6.1%), and DNS reserve limit orders are more likely to improve the quotes than reserve limit orders (17.90% vs. 5.49%). Our summary statistics in Table 2 show that orders with DNS tags are less likely to be marketable upon arrival. Therefore, 77.49% and 79.51% should be very generous lower bounds for the market share with zero price improvement for DNS orders if they could be routed outside.

We find that IOC orders always obtain a price improvement, because they cannot lock the market. Therefore, the exchange routing of IOC orders follows from Rule 611, not Rule 610. Table 4 shows that we should not consider routing by Reg NMS as homogeneous. Although routing by Rule 611 tends to offer price improvement, routing by Rule 610 does not.

Another way to compare routing led by Rule 610 and routing led by Rule 611 is to compare the reserve orders with *zero* display size to the reserve orders with *positive* display size. The results are in Table 5. When a reserve order completely hides its size, it can lock the market because Rule 610 applies only to displayed quotes. We find that 99.95% of reserve orders with a hidden quote

[11] We define price improvement in two cases. For all order types except market orders, the price improvement is defined as the difference between the execution price and the limit price. Market orders do not have limit prices, and their price improvement is defined as the difference between the execution price and the best available price at NYSE.

get a price improvement. Once the displayed size is not zero, the case for zero price improvement rises dramatically to 81.50%. Therefore, although one aim of Reg NMS is to encourage traders to display their quotes, Rule 610 creates an incentive for traders to completely hide their quotes.[12]

[Insert Table 5 about here]

Surprisingly, we find that a reserve order with a positive display size can also lock the market during the refill process. For example, on July 1, 2010, a reserve buy limit order of 10,000 shares for ELY (Callaway Golf Co) established an ask price at $6.04 using its display part of PUBQTY = 100 shares at 9:43:10:100. The order was executed and refilled with no problem until 9:50:47.752 when a series of buy orders arrived, which consumes the 100 displayed shares. At 9:50:47:755, we find that the reserve order is routed to NASDAQ and Boston Stock Exchange and takes liquidity at an ask price of $6.04. Note that the ask prices should occur only after the display of the 100-share bid at $6.04 is executed; otherwise, the NASDAQ bid would lock the NYSE ask. Therefore, we conjecture that the reserve sell limit order is routed out due to the latency in refilling the shares. We find that 73.29% of the routing orders led by refilling incur no price improvement. A routing during a refill is a rare event, and only affects 0.052% of the volume, but it still illustrates the importance of reducing exchange latency, because slow exchanges tend to lock the market of fast exchanges. As the exchange that locks the market needs to pay a take fee to other exchange or cancel the order, Rule 610 makes slow exchanges subsidize fast exchanges and becomes one possible incentive for the speed of the arms race across exchanges. To the best of our knowledge, all studies on speed competition focus on traders. Under this framework, the exchange latency does not play a role as long as all traders experience the same delay.[13] Our results suggest that Rule 610 may incentivize exchanges to be fast.

In summary, for non-marketable orders, Rule 610 creates two incentives for exchange routing refusal. First, it blurs the definition between making liquidity and taking liquidity by forbidding a locked market. Under Rule 610, an order that aims to make liquidity in one exchange may end up with taking liquidity in another to resolve locked market. Second, Rule 610 is one

---

[12] Although a reserve order with zero displayed size enjoys the benefit of locking the market, the order loses the privilege of attracting liquidity demanders from other exchanges through Rule 611. Therefore, Rules 610 and 611 create a trade-off between partially hiding quotes and completely hiding quotes.

[13] The only exception is the model by Pagnotta and Philippon (2018), who argue that fast exchanges subsidize slow exchanges. Their model focuses on Rule 611 but not 610.

driver for the differences between make, take, and routing fees.[14] As the current fee cap is 0.3 cents per share, many exchanges choose to charge a fee close to this price and give a generous rebate to liquidity makers. The difference in fees between the make liquidity in one exchange and the take liquidity in another exchange is close to twice the size of the market fee cap.

B. Returns before and after fees

After showing that Reg NMS is unlikely to provide price improvements for DNS limit orders, we compute their returns before and after fees. The 0.43-cent difference in fees would drive the returns on DNS orders from positive to negative. DNS limit orders earn a return of 0.82 bps (-0.44 bps) at the 1-second (5-minute) horizon. The number becomes positive at 1.57 bps (0.31 bps), after it collects the rebate of 0.13 cents by making liquidity at the NYSE. The returns become negative at any horizon for an order that pays a routing fee of 0.30 cents (-0.89 bps for 1-second and -2.15 bps for 5-minutes). The negative profit led by routing fees gives DNS limit order users a powerful incentive to cancel orders if they need to take liquidity from another exchange.

V. The Second Driver of NBBO Refusal: Speed

As marketable orders cannot lock or cross markets, such routings are due to Rule 611, Rule 610. Table 4 shows that exchange routing following Rule 611 leads to a price improvement of one cent or above for more than 73% of the executed market orders, and the figure is 100.00% for plain IOC orders. As a price improvement of one cent or above overwhelms the differences in fees, we search for additional motivation for exchange routing refusal. The motivation to design ISOs, according to the SEC, is to provide institutional investors with immediate access to liquidity at multiple price levels, in multiple markets, to fill large block trades with parallel order submissions (Chakravarty et al. (2012)). The results in Table 2 show, however, that the average size of an ISO is only 244.30 shares, which is smaller than both plain IOC orders (268.63 shares) and market orders (278.96 shares). The results in Panel B of Table 4 show that only 0.53% of ISOs sweep multiple price levels, which is a much lower percentage than for IOC orders (1.32%) and market orders (5.46%). These results indicate that the main driver of ISO usage diverges from the SEC's purpose to design ISOs.

[14] SEC memo. Available at: https://www.sec.gov/spotlight/emsac/emaac-regulation-nms-subcommittee-discussion-framework-040317.pdf.

In this section, we show that speed serves as an incentive for refusing exchange routing. In Subsection IV.A., we show that ISOs tend to win the speed races to take liquidity, that is, they are more likely to the first one to snipe stale quotes. In Subsection IV.B., we show that non-routable orders tend to win two types of speed races to make liquidity. First, non-routable orders are more likely to cancel stale quotes. Second, non-routable orders are more likely to win the race to add liquidity at favorable queue positions.

A. Speed races to take liquidity: Snipe stale quotes

In this section, we analyze speed races undertaken to snipe stale quotes. Aquilina, Budish, and O'Neill (2020) define races based on successful sniping, because their goal is to quantify the economic costs of stale-quote sniping. Yet we find that liquidity providers successfully cancel stale quotes 32% of the time before any sniper arrives. To study speed races in general, we extend the definition of a speed race proposed in Aquilina, Budish, and O'Neill (2020) as follows. Like these authors, we require a race to include two or more IOC orders arriving at the same stock, side, and price level within 0.1 seconds. At least one IOC order fails to fill (the "loser"). We then require some other orders to display liquidity 0.1 seconds before the first IOC order arrives, and these liquidity-providing orders are the target of sniping.[15] These requirements lead to two subcases.

1. At least one IOC snipes some stale quotes successfully; in this case, we call the IOC the winner. If $q > 1$ orders win the race, each order is counted as winning $1/q$ of the race. This subcase is similar to that discussed in Aquilina, Budish, and O'Neill (2020), in which at least one sniper wins.
2. All limit orders escape from sniping and no IOC orders win the race.

In this subsection, we consider case 1, in which at least one liquidity taker wins the race. We defer the case in which no liquidity takers win to the next subsection.

The results in Table 6 show that there are 244.76 cases where, on an average stock day, at least one IOC order wins. We find that plain IOC orders are barely aware of such speed races, participating in only 10% of them. To obtain the number, notice that plain IOC orders win 13.42 races and lose 17.37 races. We then exclude double-counting where a plain IOC order wins the

[15] We exclude reserve orders for two reasons. First, reserve orders are unlikely to win the speed race in providing liquidity, which we study in the next subsection. Second, if a reserve order completely hides its quotes, liquidity-taking orders may take extra steps to detect the order, a strategy outside the scope of this paper. Including reserve orders will increase the number of races from 244 to 258, which does not change the results substantially.

speed race over another plain IOC order. The participating rate of a plain IOC is then $\frac{13.42+17.37-5.96}{244.76}$, or about 10%. The participation rate for ISOs is as high as 77% ($\frac{134.51+140.71-86.35}{244.76}$), and the figure is 58% for DNS IOC orders.

[Insert Table 6 about here]

Conditional on participation, we also find that non-routable orders are more likely to win the speed race over plain IOC orders. For example, ISOs win 7.03 races when they compete with plain IOC orders, while ISOs lose 4.47 races. The odds for ISOs of winning this race are $\frac{7.03}{7.03+4.47}=$ 61%. DNS IOC orders win 4.38 cases more than plain IOCs and lose 2.99 cases. The odds of winning are 59%. Overall, ISOs win 57.26% of speed races, DNS IOC orders win 37.28%, and the figure is as low as 5.46% for plain IOC orders.

Both ISOs and DNS IOC orders are faster than plain IOC orders. The difference between the ISOs and DNS IOCs is the compliance cost. To comply with Reg NMS, users of ISOs need to have fast connections to all exchanges, including the smallest one. Users of DNS IOCs do not take over the compliance and they need fast connections with only a subset of exchanges. We find that ISOs win slightly more races in head-to-head competition with DNS IOCs ($\frac{47.33}{47.33+43.69}$) = 52%.

As non-routable orders tend to win the speed race to snipe stale quotes, they are more informed on a trade-by-trade basis. To identify which order type moves prices, we follow Barclay and Warner (1993) to calculate the weighted price contribution of each order type. The sum of the weighted price contribution from all order types is equal to 1. The Appendix provides the details on how we calculate the weighted price contribution.

Table 7 shows that ISO contributes more than 90% of intraday price discovery. ISOs are also the most informative on a per-share basis. They account for 35.69% of the volume [16] but 90.32% of price discovery. DNS IOC orders rank second (27.10%) in aggregate price discovery, and their contribution to price discovery is slightly larger than their share of trading volume (25.87%). Plain IOC orders only contribute to 3.86% of price discovery, and their price contribution is much less than their share of liquidity taking volume (6.36%). Plain market orders contribute negatively to price discovery. In summary, for liquidity-demanding orders, orders with

[16] The number is larger than that in Table 2 because weighted price contribution considers only orders that take liquidity.

non-routing instructions are more informative than orders without non-routing instructions.

[Insert Table 7 about here]

B. Speed races for liquidity making orders

Speed races involving liquidity taking lead to immediate profit, but the link between speed races and liquidity making is indirect. In this subsection, we discover two types of speed races for liquidity making orders. In Subsection V.B.1., we show that DNS limit orders win the speed race to cancel stale quotes. In Subsection V.B.2., we show that DNS limit orders win the speed race to secure front queue positions. In Subsection V.B.3., we provide evidence on the economic benefits of winning these two types of speed races. Again, the main driver of the results is clientele effects. It is not because DNS limit orders are faster than DAY orders, but because DNS limit order users are faster than DAY limit order users. Our comparison is based on displayed orders as reserve orders are less likely to be involved in speed races. Even if a reserve order arrives before a displayed order, its non-displayed part has a lower execution priority than the displayed order.

B.1. Speed races to cancel stale quotes

In the speed races illustrated in Subsection V.A, IOC orders aim to snipe stale quotes, but orders that establish quotes can cancel their stale quotes before being sniped. If all stale quotes are canceled, speed races end up with no trades. To obtain a complete picture of speed races, we broaden the definitions of races proposed in Aquilina, Budish, and O'Neill (2020) by considering cases in which all or some liquidity providers win races. To ensure that stale quotes exist before the races, we require that at least one limit order exists at the price of an IOC order 0.1 seconds before the first IOC order arrives. We then track all limit orders at the price 0.1 seconds before the first IOC order arrives. These orders would execute if they did not cancel. A limit order wins the speed race if it successfully cancels.[17] Table 8 shows that 380.37 orders are targets of sniping. Among them, 253.25 are sniped.[18] Table 8 also shows that 33.42% of orders are not sniped during the races. Therefore, we consider broader samples of races than in the previous section or in

[17] We require at least two IOC orders to participate in a race, which helps rule out other drivers of cancellation. For example, HFTs can cancel their orders because they want to back-run (Yang and Zhu (2020)) or front-run (Baldauf and Moller (2020)) order flows of institutional traders. It is unlikely, however, that two institutional traders arrive at the same time and submit IOC orders at the same price.

[18] This number is larger than the total number of races documented in the previous subsection (244.76) because sometimes more than one order is sniped.

Aquilina, Budish, and O'Neill (2020).

[Insert Table 8 about here]

Table 8 also shows that DNS limit orders are more likely to be targets of stale-quote sniping: 274.37 DNS limit orders are on the limit-order book as the race starts. However, 117.29 DNS limit orders successfully cancel, leading to an escape rate of 42.75% and DAY limit orders escape only 9.27% of the time. Combining this result with the results reported in Panel B of Table 4, we see that DNS limit orders not only are more likely to improve BBO quotes but also are more likely to avoiding sniping.

B.2. Speed races to establish front queue positions

The quick cancellation of stale quotes can reduce losses for DNS limit orders. However, to profit from making liquidity, DNS limit orders should also win the speed race to add liquidity. Li, Wang, and Ye (2020) model the speed race to provide liquidity. Because Rule 612 of Reg NMS imposes a one-cent tick size (a minimal price variation) for any stock with a price above $1, the discrete price constrains price competition and creates rents for liquidity provision. Speed becomes essential in capturing such rents because orders at the front queue position have higher execution priority (Yao and Ye (2018)).

We capture such speed races following Li, Wang, and Ye (2020). Their model shows that traders would continue to add more shares to the queue until the marginal profit equals zero. A new profitable queue position opens when a marketable order moves the existing queue forward. Therefore, we examine which order type is more likely to react after a marketable order moves the queue forward. We define the first responder to a trade as the limit order that (1) is submitted within 0.1 seconds after a transaction executes in the opposite direction, (2) is not marketable, (3) has a limit price identical to the trade price, and (4) no race to take liquidity occurs in the following 0.1 seconds. If multiple orders satisfy (1) – (4), we take the first arriving order as the first responder.

The results in Table 9 show that first responders are disproportionally DNS limit orders. On average, 31.7% of liquidity-taking orders are followed by orders to refill at the same prices within 0.1 seconds, indicating that the market is resilient even at sub-second horizons. DNS limit orders win the race 71.5% of the time, suggesting that DNS limit orders are more likely to achieve favorable queue positions.

[Insert Table 9 about here]

B.3. Economic benefit to win speed race

Because DNS limit orders are more likely to establish front queue position and to escape from sniping, we expect that they will have higher returns than DAY limit orders. Although it is hard to quantify the economic gain from wining these two types of speed races, in this subsection we provide some preliminary evidence for the reward of winning speed.

To control for the cross-sectional and time series variations, we present the results controlling for stock day fixed effect. In Table 10, we compare the return differences for routable and non-routable orders for the same stock for the same day. Column (1) shows that DNS limit orders earn an effective spread 0.76 bps higher than DAY limit orders.[19] Li, Wang, and Ye (2020) provide one interpretation for why faster orders earn a higher spread. Because fast traders are more likely to achieve front queue positions, slow traders have to undercut them to achieve execution, often at a bid or ask price that loses money. Indeed, we find DAY limit orders lose 0.19 bps one second after execution, whereas DNS limit orders make 0.82 bps 1 second after execution.

[Insert Table 10 about here]

The front queue position, however, is only one of the drivers of higher profits of DNS limit orders one second after execution. We find that their profit margin widens to 1.01 bps one second after execution. The 0.25-bps difference (1.01 - 0.76) implies that DAY limit orders are more likely to be adversely selected than DNS limit orders. This result is consistent with results in Subsection V.B.1, where we find that DNS limit orders are more likely to escape from sniping.

VI. Discussion on Cross-exchange Variation and Time Series Evolution of Order Types

The economic intuition we reveal provides interpretations for order types in other exchanges, as well as the evolution of order types. Exchanges constantly revise their order types or introduce new ones to cater to their clienteles. For example, users of DNS limit orders aim to make liquidity; however, DNS limit orders have two limitations. Exchanges later introduce new order types to address these two limitations.

First, we find that 11% DNS limit orders take liquidity from the NYSE. The difference

[19] The difference is higher than the difference in Figure 1 because DNS limit orders take a larger market share for high volume stocks, for which the bid-ask spread is smaller. The stock day fixed effect adjusts such bias and leads to a larger return difference between DNS limit orders and DAY limit orders for the same stock at the same day.

between the take fee and the make rebate, 0.34 cents, is lower than the difference between the routing fee and the make rebate, 0.43 cents. Still, traders who want to avoid routing fees may also want to avoid take fees. Therefore, the NYSE introduced add-liquidity-only (ALOs) orders in 2014. ALO orders refuse to take liquidity at their limit price in all exchanges, including the NYSE. Thus, an ALO order better fits the needs of fee-sensitive liquidity makers. The first generation of ALO orders refuses to take liquidity at any price. The NYSE later updated ALO orders so that they take liquidity if they receive a price improvement at least one tick from its limit price. The introduction of ALO orders and their later update follows two economic mechanisms revealed by our paper. First, liquidity making orders refuse to take liquidity at their limit price due to higher cum fee costs. Second, a true price improvement overwhelms the fee.

Second, DNS instructions cancel an order when an order locks or crosses a market. The NYSE resolved potential unnecessary cancelations by introducing slide orders ("non-routable limit" orders) in 2016 and day ISOs in 2014. The slide instruction reprices an order to comply with Rule 610. For example, suppose that the current NBO is $5.01 and the NBB is $5.00. A limit order that aims to buy at $5.01 would lock the market. After a DNS instruction cancels the order, the DNS order submitter needs to resubmit the order at a different price, or continuously monitor the market to wait for an opportunity to submit a bid to buy at $5.01 later. The slide instruction implements these two functions automatically. The slide instruction first reprices the limit order to $5.00 so that it does not lock the market. If the market ticks up to an NBO of $5.02 and an NBB of $5.01, the slide instruction would reprice the order from buy at $5.00 to buy at $5.01. Besides NYSE, NASDAQ and CBOE also offer slide orders.

In Subsection V.B., we show winners of the speed race to secure the front queue positions have higher profits, and in Subsection IV.A, we show that orders that are completely hidden can lock the market. The Direct Edge Exchange introduced Hide Not Slide orders by combining these two features. The Hide Not Slide instruction hides the order if it locks the market. Once the market unlocks, this order lights up. These orders enjoy time priority over the slide order, because their timestamp is the time of the entry. These orders attract huge media attention and regulatory scrutiny. A 2012 *Wall Street Journal* article claim these orders can jump ahead of the queue.[20] On January 12, 2015, the SEC announced that "two exchanges formerly owned by Direct Edge

[20] *The Wall Street Journal*. For superfast stock traders, a way to jump ahead in line. September 19, 2012. Available at https://www.wsj.com/articles/SB10000872396390443989204577599243693561670.

Holdings and since acquired by BATS Global Markets have agreed to pay a $14 million penalty to settle charges that their rules failed to accurately describe the order types being used on the exchanges. The penalty is the SEC's largest against a national securities exchange."[21] Interestingly, the SEC penalty is not based on the Hide Not Slide order but on the inaccurate disclosure of the functions of the order. Direct Edge describes a single "price sliding" process, but they actually "offered three variations of 'price sliding' order types. The exchanges' rules did not completely and accurately describe the prices at which those orders would be ranked and executable in certain circumstances, and they also failed to describe the execution priority of the three order types relative to each other and other order types." In addition, Direct Edge "provided complete and accurate information about the order types to only some members, including certain high-frequency trading firms that Direct Edge also solicited for input about how the Hide Not Slide order type should operate. Direct Edge originally developed this order type following a request from one of the firms." This penalty provides a unique case that summarizes three implications of our paper. First, two orders that differ in one instruction can have dramatically different economic functions and outcomes. Second, exchanges design order types to cater to their clientele. Certainly, designing new order types without accurate disclosure leads to regulatory actions. Third, Rule 610 is a major driver of the proliferation of order types, as many order types are designed to provide liquidity without violating Rule 610. Fourth, queue position is important in liquidity provision.

In our sample, there are two types of non-routable IOCs: ISO-IOC and DNS IOC, but only one type of non-routable limit orders: DNS limit. The NYSE and other exchanges introduced Day ISOs to span the spectrum of order types. Day ISOs resolve the locked or crossed market by sweeping the quotes that lock or cross the market. A Day ISO, if marketable on arrival, will be immediately traded with contra-side interest in the NYSE book up to its full size and limit price. Any untraded quantity of a Day ISO will be displayed at its limit price. The sender of the Day ISO complies with Rule 610 by sending Day ISOs to other exchanges to clear the locked or crossed quotes. NYSE also introduced ALO Day ISOs. This order type cancels if it needs to take liquidity from the NYSE. If the order can establish a quote in the NYSE, the submitter complies with Rule 610 by sending ISOs to other exchanges to clear lock or crossed markets. Both Day ISOs and ALO Day ISOs make liquidity first by taking liquidity, but ALO Day ISO takes less liquidity. For

[21] The SEC charges Direct Edge Exchanges with failing to properly describe order types. January 12, 2015. Available at: https://www.sec.gov/news/pressrelease/2015-2.html.

example, Day ISOs may need to execute against a large hidden order within the NYSE, whereas an ALO Day ISO cancels the order. As of 2020, the market share of ALO Day ISOs is more than ten times higher than Day ISOs, which indicates that ALO Day ISOs aim to provide liquidity by taking as little liquidity as possible in the first step. More broadly, this evolution is consistent with our finding that Rule 610 is one of the main drivers of the proliferations of order types.

As a combination of these three new order types (ALO, slide, day ISO) provides better solutions for the problem that DNS limit orders aimed to address, DNS limit orders disappeared in 2016. Other exchanges also adopted these three solutions: to make liquidity only, to slide, or to sweep, although they may provide these three functions with order types that bear different names.

The order types in our sample mostly cater to short-term traders. Exchanges gradually fill in the spectrum of order types by designing new order types that serve long-term investors. We show that IOC orders, particularly non-routable IOC orders, are often used to snipe stale quotes. The NASDAQ recently designed the midpoint extended life order (M-ELO). This order type refuses to interact with IOC orders, although they welcome trades with resting limit orders as well as other M-ELOs. M-ELOs thus aim to avoid being sniped by restricting the counterparties of a trade.

We find that plain market orders and plain limit orders lose money, which provides one explanation for why the NYSE no longer offers these two order types. The NYSE requires that all orders have a limit price or peg to a price, and have an expiration time. Retail traders can still see these order types in their brokerage accounts, but it is either because their brokers internalize these orders and never bring them to the NYSE, or because brokers synthetically create obsolete order types for people who believe that they still exist. Brokers may assign a working price for a market order before they enter them in an exchange, and they will reenter a GTC order to an exchange at the market open after the exchange cancels them at the previous close. The market structure in the U.S. evolves so fast that order types that most people believe to be dominant are actually obsolete, whereas the true dominant order types are outside the radar of most people. We hope that our paper provides some foundation and structure for the order types and reveals their economic principles.

VII. Conclusion

Any stock-trading strategy must choose an order type to interact with the market. As the first comprehensive study of exchange order types, we identify a three-tiered world of order types.

Plain market and limit orders and stop orders lose money and are most likely to come from naïve traders. Sophisticated traders manage their order types but to varying degrees. Users of the most complex order types do not allow the NYSE to route their orders to other exchanges. These non-routable orders bring small profits at the intraday level, but they are not profitable beyond the trading day. Traders who delegate exchanges to find the best price incur higher transaction costs and earn negative returns at short horizons, but all order types that contain long-term information are routable. Our results suggest that, at the order level, short-term execution strategies and long-term investment strategies are substitutes but not complements.

Fees provide one incentive for orders to refuse exchange routing. SEC Rule 610 prohibits a locked market. One way to unlock the market is to route an order to take liquidity from another exchange. Such routing does not lead to price improvement and leads worse price after adjusting for fees as the order collects a rebate of 0.13 cents if it makes liquidity at the NYSE but pays a routing fee of 0.3 cents if it takes liquidity from outside. These 0.43 cent fee differences are large enough to switch returns on DNS limit orders from positive to negative. Therefore, refusing exchange routing is the key to realizing small profits.

As routing may increase latency for traders, speed is another factor that drives traders to refuse the best prices. We find that non-routable order users are more likely to win three types of speed races. First, non-routable orders are more likely to be the first to snipe stale quotes. Second, non-routable orders are more likely to win time priority for liquidity provision. Third, non-routable orders are more likely to cancel before adverse price movements. Therefore, refusing exchange routing helps traders capture quick profits.

As 57% of trading volume originates in orders that aim to avoid small fees or achieve speed advantages, we provide insights into a proposed New Jersey State Government transaction tax of 0.25 cents per share. We find that stock trading is highly sensitive to small differences in fees at the same magnitude of the transaction tax. Therefore, our results provide justification for the aggressive response of the NYSE and NASDAQ to the proposed transaction tax. Both exchanges activated their backup site in Chicago, Illinois to prove their ability to pull their business out of New Jersey.[22] Meanwhile, although major exchanges are currently operating at separate locations,

[22] Matthew Leising, Bloomberg, 2020. Leaving N.J. for Chicago gives no easy tax fix to Nasdaq, NYSE. Available at: https://www.bloomberg.com/news/articles/2020-09-24/leaving-n-j-for-chicago-gives-no-easy-tax-fix-to-nasdaq-nyse.

they will move to the same location once they activate their backup site in Illinois. The move of exchanges to the same location also affects speed. For example, the current latency difference between cable and microwave would disappear. Therefore, a seemingly small transaction tax would fundamentally change the landscape of trading in the U.S., in terms of both fees and speed.

References

Aquilina, Matteo, Eric B. Budish, and Peter O'Neill, 2020, Quantifying the high-frequency trading "arms race": A simple new methodology and estimates, Chicago Booth Research Paper.

Baldauf, Markus, and Joshua Mollner, 2020, High-Frequency Trading and Market Performance, *Journal of Finance* 75, 1495-1526.

Barclay, Michael J., and Jerold B. Warner, 1993, Stealth trading and volatility: Which trades move prices? *Journal of Financial Economics* 34: 281-305.

Bessembinder, Hendrik, and Kumar Venkataraman, 2004, Does an electronic stock exchange need an upstairs market? *Journal of Financial Economics* 73: 3-36.

Bessembinder, Hendrik, Marios Panayides, and Kumar Venkataraman, 2009, Hidden liquidity: an analysis of order exposure strategies in electronic stock markets, *Journal of Financial Economics* 94: 361-383.

Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50: 1655-1689.

Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High-frequency trading and price discovery, *Review of Financial Studies* 27, 2267-2306.

Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics* 130: 1547-1621.

Chao, Yong, Chen Yao, and Mao Ye, 2019, Why discrete price fragments U.S. stock exchanges and disperses their fee structures, *Review of Financial Studies* 32: 1068-1101.

Chakravarty, Sugato, Pankaj Jain, James Upson, and Robert Wood, 2012, Clean sweep: Informed trading through intermarket sweep orders, *Journal of Financial and Quantitative Analysis*, 415-435.

Crouzet, Nicolas, Ian Dew-Becker, and Charles G. Nathanson, 2020, On the effects of restricting short-term investment, *Review of Financial Studies* 33, 1-43.

Dittmar, Amy, and Laura Casares Field, 2015, Can managers time the market? Evidence using repurchase price data, *Journal of Financial Economics*, 115: 261-282.

Easley, David, and Maureen O'Hara, 1991, Order Form and Information in Securities Markets, *Journal of Finance* 46, 905-27.

Glosten, Lawrence R., and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.

Griffiths, Mark D., Brian F. Smith, D. Alasdair S. Turnbull, and Robert W. White, 2000, The costs and determinants of order aggressiveness, *Journal of Financial Economics* 56: 65-88.

Harris, Lawrence, and Joel Hasbrouck, 1996, market vs. limit orders: the SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31: 213-231.

Hollifield, Burton, Robert A. Miller, and Patrik Sandas, 2004, Empirical analysis of limit order markets, *Review of Economic Studies* 71: 1027-1

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, A rational theory of mutual funds' attention allocation, *Econometrica* 84, 571-626.

Kaniel, Ron, and Hong Liu, 2006, So what orders do informed traders use? *Journal of Business* 79: 1867-1913.

Kelley, Eric K., and Paul C. Tetlock, 2013, How wise are crowds? Insights from retail orders and stock returns, *Journal of Finance* 68: 1229-1265.

Kyle, Albert S, 1985, Continuous auctions and insider trading, *Econometrica: Journal of the Econometric Society*, 1315-1335.

Li, Sida, Xin Wang, and Mao Ye, 2020, Who provides liquidity and when? *Journal of Financial Economics, forthcoming*.

O'Hara, Maureen, and Mao Ye, 2011, Is market fragmentation harming market quality?, *Journal of Financial Economics* 100, 459-474.

O'Hara, Maureen, Chen Yao, and Mao Ye, 2014, What's not there: Odd lots and market data, The *Journal of Finance* 69, 2199-2236.

Pagnotta, Emiliano S., and Thomas Philippon, 2018, Competing on speed, *Econometrica* 86, 1067-1115.

Parlour, Christine A, 1998, Price dynamics in limit order markets, *Review of Financial Studies* 11: 789-816.

Parlour, Christine A., and Duane J. Seppi, 2008, Limit order markets: A survey. *Handbook of Financial Intermediation and Banking,* 5: 63-95.

U.S. Securities and Exchange Commission, 2018, Market Data Infrastructure, File No. S7-03-20

Yang, Liyan, and Haoxiang Zhu, 2020, Back-running: Seeking and hiding fundamental information in order flows, *Review of Financial Studies* 33, 1484-1533.

Yao, Chen, and Mao Ye, 2018, Why trading speed matters: A tale of queue rationing under price controls, *Review of Financial Studies* 31, 2157-2183.

Ye Mao, Miles Zheng, and Xiongshi Li, 2020, Price Ceiling, Market Structure, and Payout Policies, Working paper, University of Illinois at Urbana-Champaign and Guangxi University of Finance and Economics.

Table 1
Sample selection, sample stocks, and summary statistics

Panel A: Sample Selection Criteria	
All NYSE securities	2413
Non-common stock equities (American Depository Receipts, units, certificates, and shares of beneficial interest)	-565
Common stocks of non-U.S. companies, close-end funds, Real Estate Investment Trusts, and Americus Trust components, and exchange-traded funds	-517
Dual class stock	-152
Price (December 31, 2009) < 5	-93
Universe sample	1086
Final sample by selecting from our universe sample every tenth stock (December 2009)	109

Panel B: Tickers of final sample stocks
ACI, AGL, AGP, AIG, AMD, AME, AOS, ARM, ATI, B, BAC, BBT, BGG, BW, BYI, CGA, CMN, CNS, CPN, CPO, CPX, CSS, CSX, CYH, DGI, DPL, DVN, ELY, EMR, ENZ, ES, ESE, EXP, FCN, FMR, FOR, GAS, GCI, GD, GEO, HL, HLX, HNZ, HRS, HW, HXL, IEX, IFF, IPI, IWA, JBL, JLL, JMP, JWN, KCI, KSU, MCD, MCS, MCY, MDC, MHP, MMC, MOS, MTG, MTZ, N, NOV, NPK, NWN, OKE, OSK, PCP, PPD, PVR, RAI, RGA, RGR, RGS, ROK, RRC, RRI, SBX, SHW, SMG, SON, SUG, SUR, SY, TC, TE, THC, TIN, TLB, TRC, TUP, USB, VMI, VQ, VSH, VVI, WFR, WGL, WGO, WMS, WPP, WWW, XCO, XEC, Y

Panel C: Summary statistics

	N	Mean	SD	5%	25%	50%	75%	95%
Market Cap	109	6.3	15.28	0.23	0.79	2.17	3.73	26.29
Share Volume	109	1.36	8	0.02	0.09	0.23	0.61	2.41
Turnover	109	3.34	2.1	0.79	1.89	2.87	4.24	8.48
Share Price	109	31.63	31.19	7.54	13.62	26.54	41.14	68.17

In Panel A, we report our sample selection criteria. In Panel B, we report the tickers of our 109 sample stocks. In Panel C, we present summary statistics (at the beginning of fiscal year 2010). *Market Cap* is market capitalization ($billion). *Share Volume* is annual shares traded. Share *Turnover* is shares traded divided by shares outstanding. *Share Price* is the nominal share price.

Table 2
Summary statistics for order types

Classification	Order Type	Trades	Executed Volume	Average Order Size	Market Share (Executed Volume)	Fill Rate	Take Local Liquidity	Route	Make Liquidity
Liquidity-taking	Plain market	6,364,008	1,775,278,400	278.96	4.48%	100.00%	66.34%	33.64%	0.00%
	Plain IOC	4,196,808	1,127,373,900	268.63	2.84%	19.48%	98.07%	1.93%	0.00%
	ISO	25,410,600	6,207,890,700	244.30	15.66%	22.48%	100.00%	0.00%	0.00%
	DNS IOC	14,652,500	4,499,703,400	307.09	11.35%	31.86%	100.00%	0.00%	0.00%
Liquidity-making	DAY limit	43,360,100	9,560,971,500	220.50	24.12%	2.71%	29.78%	14.28%	55.94%
	DNS limit	44,243,300	8,069,151,300	182.38	20.36%	3.72%	11.43%	0.00%	88.57%
	Reserve	6,660,529	1,917,804,300	287.94	4.84%	2.09%	31.09%	16.30%	52.62%
	DNS Reserve	4,077,848	732,919,700	179.73	1.85%	1.56%	5.00%	0.00%	95.00%
Trailing Price	Peg	1,256,004	237,375,000	188.99	0.60%	78.12%	0.71%	0.01%	99.28%
	Buy Minus Zero Plus	13,263	2,330,600	175.72	0.01%	45.88%	7.71%	12.91%	79.38%
Auction	Market-on-close	897,083	3,890,610,000	4336.96	9.82%	100.00%			
	Market-on-open	346,417	798,674,000	2305.53	2.02%	100.00%			
	Limit-on-open	373,525	289,195,000	774.23	0.73%	1.96%			
	Limit-on-close	212,533	448,966,000	2112.45	1.13%	15.30%			
No Expiry Time	STOP	150,566	50,090,400	332.68	0.13%	-	68.28%	31.65%	0.07%
	Good-Till-Cancel (Plain Limit)	36,612	20,495,800	559.81	0.05%	-	11.30%	4.28%	84.42%
	Do Not Reduce	6,297	6,119,500	971.81	0.02%	-	12.57%	5.30%	82.13%

In this table, we report the summary statistics for each order type, including the total number of trades, total share volume, market share of share volume, average order size, and average fill rates. The fill rates for the last three types of orders are blank because they do not have expiration time. The average order size includes the unfilled portion of an order. We calculate breakdown of the executed volume based on its execution destination, i.e., the percentage of the executed volume that take local liquidity, route to away markets, or make liquidity locally.

Table 3
Market share of routing volume and market share of executed volume

Destination	Number of routed trades	Number of routed shares	Market share (routed volume / total routed volume)	Operator (2010)	Market share (exchange trading volume / consolidated volume)
NYSE Arca	4,408,234	1,015,980,000	53.86%	NYSE	27.74%
NASDAQ	2,877,660	493,862,000	26.18%	NASDAQ OMX	31.97%
BZX	1,426,802	254,230,000	13.48%	Bats Global	16.97%
EDGX	225,823	49,951,000	2.65%	Bats Global	5.52%
National Stock Exchange	281,237	49,503,100	2.62%	NSX	1.42%
Boston Stock Exchange	73,749	11,315,900	0.60%	NASDAQ OMX	7.83%
Chicago Stock Exchange	15,414	4,634,200	0.25%	CHX Holdings	0.59%
EDGA	30,654	3,769,500	0.20%	Bats Global	6.17%
Philadelphia Stock Exchange	8,900	1,706,100	0.09%	NASDAQ OMX	0.58%
BYX	7,296	992,100	0.05%	Bats Global	0.99%
CBOE	3,350	546,000	0.03%	CBOE	0.24%
AMEX	0	0	0.00%	NYSE	0.00%

In this table, we report the routing destinations from the NYSE (sorted by number of shares routed), the number of routed trades, the number of routed shares, the exchanges' corresponding operators in 2010, their market shares of routing volume over total routed volume, and their market shares of trading volume over consolidated volume of all 13 exchanges.

Table 4
Order types and limit order book dynamics

Panel A: Price improvements for routed out orders

Improvement	Plain Market	Plain IOC	DAY Limit	Reserve
0	26.99%	0.00%	77.49%	79.51%
0.01	48.58%	70.49%	12.93%	9.43%
0.02	13.87%	12.00%	2.74%	3.52%
0.03	5.11%	4.52%	1.35%	1.68%
0.04	2.09%	2.46%	0.96%	1.32%
0.05	0.98%	1.59%	0.95%	1.33%
0.06	0.52%	1.00%	0.34%	0.63%
0.07	0.32%	0.81%	0.26%	0.36%
0.08	0.22%	0.69%	0.21%	0.21%
0.09	0.15%	1.13%	0.16%	0.15%
0.1	0.04%	0.66%	0.07%	0.05%
0.10+	1.12%	4.65%	2.56%	1.82%

Panel B: Orderbook impact by messages

Category	Plain Market	Plain IOC	ISO	DNS IOC	DAY Limit	DNS Limit	Reserve	DNS Reserve
Trade at multiple prices	5.46%	1.32%	0.53%	0.25%	0.04%	0.00%	0.06%	0.00%
Trade at one price	94.54%	98.68%	99.47%	99.75%	4.69%	4.30%	2.10%	3.38%
Improving Order	0.00%	0.00%	0.00%	0.00%	6.10%	24.03%	5.49%	17.90%
Order Placement at BBO	0.00%	0.00%	0.00%	0.00%	3.45%	12.07%	2.77%	9.14%
Order 1 tick from BBO	0.00%	0.00%	0.00%	0.00%	4.71%	10.39%	2.81%	6.79%
Order > 1 tick from BBO	0.00%	0.00%	0.00%	0.00%	81.02%	49.20%	86.78%	62.80%

In this table, we report price improvements and order book impact for routable orders. Panel A presents the realized price improvements by order type. For all order types except market orders, the price improvement is defined as the difference between the execution price and the limit price. Market orders do not have limit prices, and their price improvement is defined as the difference between the execution price and the best available price at NYSE. In Panel B we report limit-order-book impacts by order types. Market orders, Plain IOCs, NNS IOC orders, and ISOs can affect the limit order book only by trading (either at one price level or at multiple price levels). DAY limit, DNS limit, reserve orders, and DNS reserve limit orders can either be marketable or add liquidity to the limit-order book.

Table 5
Price Improvement for displayed and non-displayed reserve orders

Improvement	Reserve Orders	
	PUBQTY=0	PUBQTY>0
0	0.05%	81.50%
0.01	20.10%	9.16%
0.02	6.63%	3.45%
0.03	3.79%	1.63%
0.04	17.77%	0.90%
0.05	28.23%	0.66%
0.06	12.91%	0.32%
0.07	5.60%	0.23%
0.08	2.19%	0.16%
0.09	1.07%	0.12%
0.1	0.20%	0.05%
0.10+	1.44%	1.83%

In this table, we report price improvements for reserve orders categorized by whether or not they can lock the market. Reserve orders with no published quantity (PUBQTY = 0) do not lock the displayed market, while reserve orders that are partially published (PUBQTY > 0) may lock the displayed market. The price improvement is defined as the difference between an order's execution price and its limit price.

Table 6
Speed races to snipe stale quotes

Winner \ Loser	Plain IOC	ISO	DNS IOC	Total winning races	Winning rate
Plain IOC	5.96	4.47	2.99	13.42	5.46%
ISO	7.03	86.35	47.33	140.71	57.26%
DNS IOC	4.38	43.69	43.56	91.63	37.28%
Total Losing races	17.37	134.51	93.88	244.76	

In this table, we report the pairwise speed race counts among three types of IOC orders, where we define the speed races in Section IV.A. The order types in the rows are race winners and the order types in the columns are race losers. If $q > 1$ orders win (lose) the race, each order is counted as winning (losing) $1/q$ of the race. All numbers are averaged at the stock–day level.

Table 7
Weighted price contribution by order types

Order Type	Observation	Total shares (Taking)	Mean size (Taking)	Share of taking volume	WPC
Plain market	3,801,767	1,177,671,600	309.77	6.77%	-4.56%
Plain IOC	4,065,849	1,105,589,800	271.92	6.36%	3.86%
ISO	25,410,520	6,207,867,200	244.30	35.69%	90.32%
DNS IOC	14,652,458	4,499,691,500	307.09	25.87%	27.10%
DAY limit	11,366,776	2,846,871,200	250.46	16.37%	-13.78%
DNS limit	4,448,268	922,261,400	207.33	5.30%	-2.34%
Reserve limit	945,753	596,177,100	630.37	3.43%	-0.95%
Reserve DNS limit	148,111	36,655,900	247.49	0.21%	0.35%

In this table, we report the price discovery decomposition of order types, measured by the weighted price contribution following Barclay and Warner (1993). The sum of the weighted price contribution from all order types is equal to 1. We describe how we calculate the weighted price contribution in the Appendix.

Table 8
Speed races to cancel stale quotes

	DAY limit	DNS limit	Total
Number of outstanding limit orders 0.1 seconds before races	106.00	274.37	380.37
Executed in races	96.17	157.08	253.25
Escaped immediately before races	9.83	117.29	127.12
Escape rate	9.27%	42.75%	33.42%

In this table, we report the outcomes of speed races run to cancel stale quotes among limit orders and subsequent arriving IOC snipers. We track the outcomes of all outstanding limit orders 0.1 seconds before the first IOC snipers arrive. The escape rate is the proportion of limit orders that are canceled before the IOC snipers arrive. All numbers are averaged at the stock–day level.

Table 9
Speed races to establish front queue positions

	Average trades per day	First responder count per day		
		Total	DAY limit	DNS limit
Full Sample	1652	495	121	374
Breakdown			24.4%	75.6%
Large	3118	1108	258	850
Breakdown			23.3%	76.7%
Medium	1065	263	75	188
Breakdown			28.5%	71.5%
Small	680	124	36	88
Breakdown			29.0%	71.0%

In this table, we count the first responders to liquidity-taking orders. The first responders are non-marketable limit orders that refill liquidity at the same price where the liquidity is consumed. To ensure that the responding limit order is triggered by the liquidity-taking order, we require that the refill order be submitted within 0.1 seconds. The average number of first responders per day is reported in the second column, and the following columns report the average numbers of first responders per day by order types. The market shares of responding orders are reported under the counts.

Table 10
Pairwise return comparison for liquidity-making orders

	DAY limit vs. DNS limit				
	Returns				
	0	1sec	30sec	1min	5min
DAY limit	6.23	-0.19	-0.95	-1.01	-1.15
DNS limit	6.99	0.82	0.07	-0.09	-0.44
Diff.	**-0.76**	**-1.01**	**-1.02**	**-0.92**	**-0.70**
t-stat.	-8.66	-7.11	-7.98	-8.13	-8.42

In this table, we report the stock-day level pairwise return comparison for liquidity-making orders between DAY limit and DNS limit. The units are in bps. Returns are measured as differences between execution prices and midpoints at each time horizon after execution, adjusted by trading direction. The return at horizon 0 is also known as the effective half spread, and the return at horizon 5 min is also known as the realized half spread. We control for any stock-day specific order type usage difference by: (1) aggregating the share-weighted profits for each stock-day in our sample, and (2) pair the stock-day observations and conduct a paired *t*-test between order types. Standard errors are clustered at both the firm and day levels.

Figure 1
Short-term return dynamics

These figures display the dynamics of average returns for various order types in basis points. Panel A shows the results for plain IOC orders, ISOs, and DNS IOC orders. Panel B shows the results for day limit orders, DNS limit orders, reserve limit orders, and DNS reserve limit orders. Panel C shows the results for all order types without an expiration date, i.e., plain limit order, Do-Not-Reduce order, and STOP order. Panel D shows the results for order types with limit prices that are conditioned on the prevailing market prices, i.e., BMZP and Pegged orders. Returns are measured as the differences between execution prices and the midpoints at each time horizon after execution, adjusted by trading direction. The profit at horizon 0 is also known as effective half-spread.

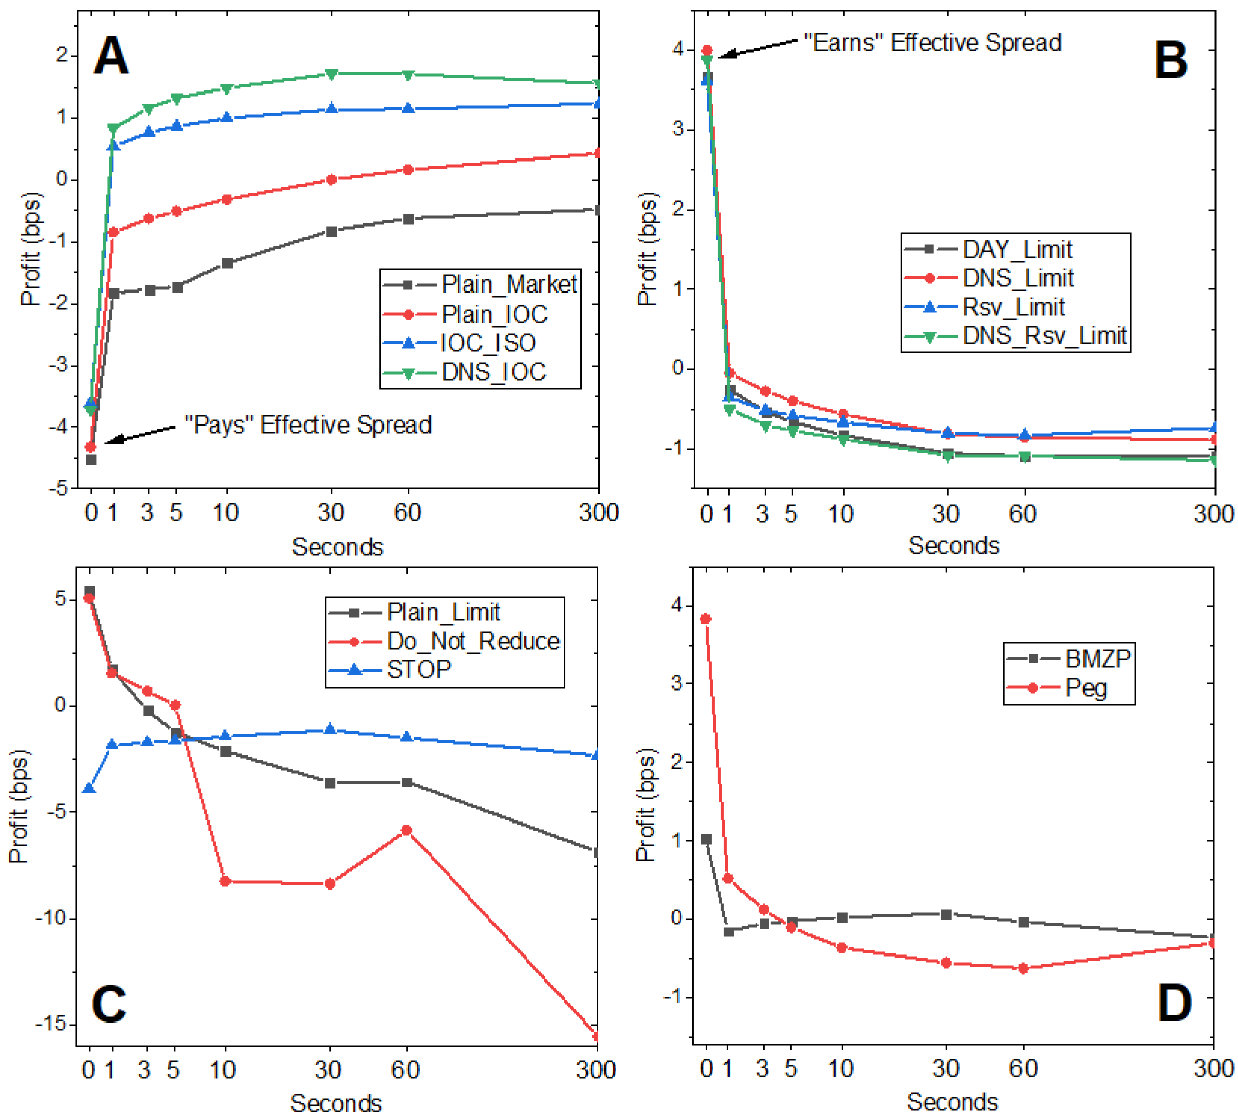

Figure 2
Long-term return dynamics

This figure shows the dynamics of average returns in basis points by order types. Panel A shows the results for plain IOC orders, ISOs, and DNS IOC orders. Panel B shows the results for DAY limit orders, DNS limit orders, reserve limit orders, and DNS reserve limit orders. Panel C shows the results for all order types without an expiration date, i.e., plain limit order, Do-Not-Reduce order, and STOP order. Panel D shows the results for order types with limit prices that are conditioned on the prevailing market prices, i.e., BMZP and Peg orders. Returns are measured as the differences between execution prices and the closing price each day after the execution, adjusted by trading direction.

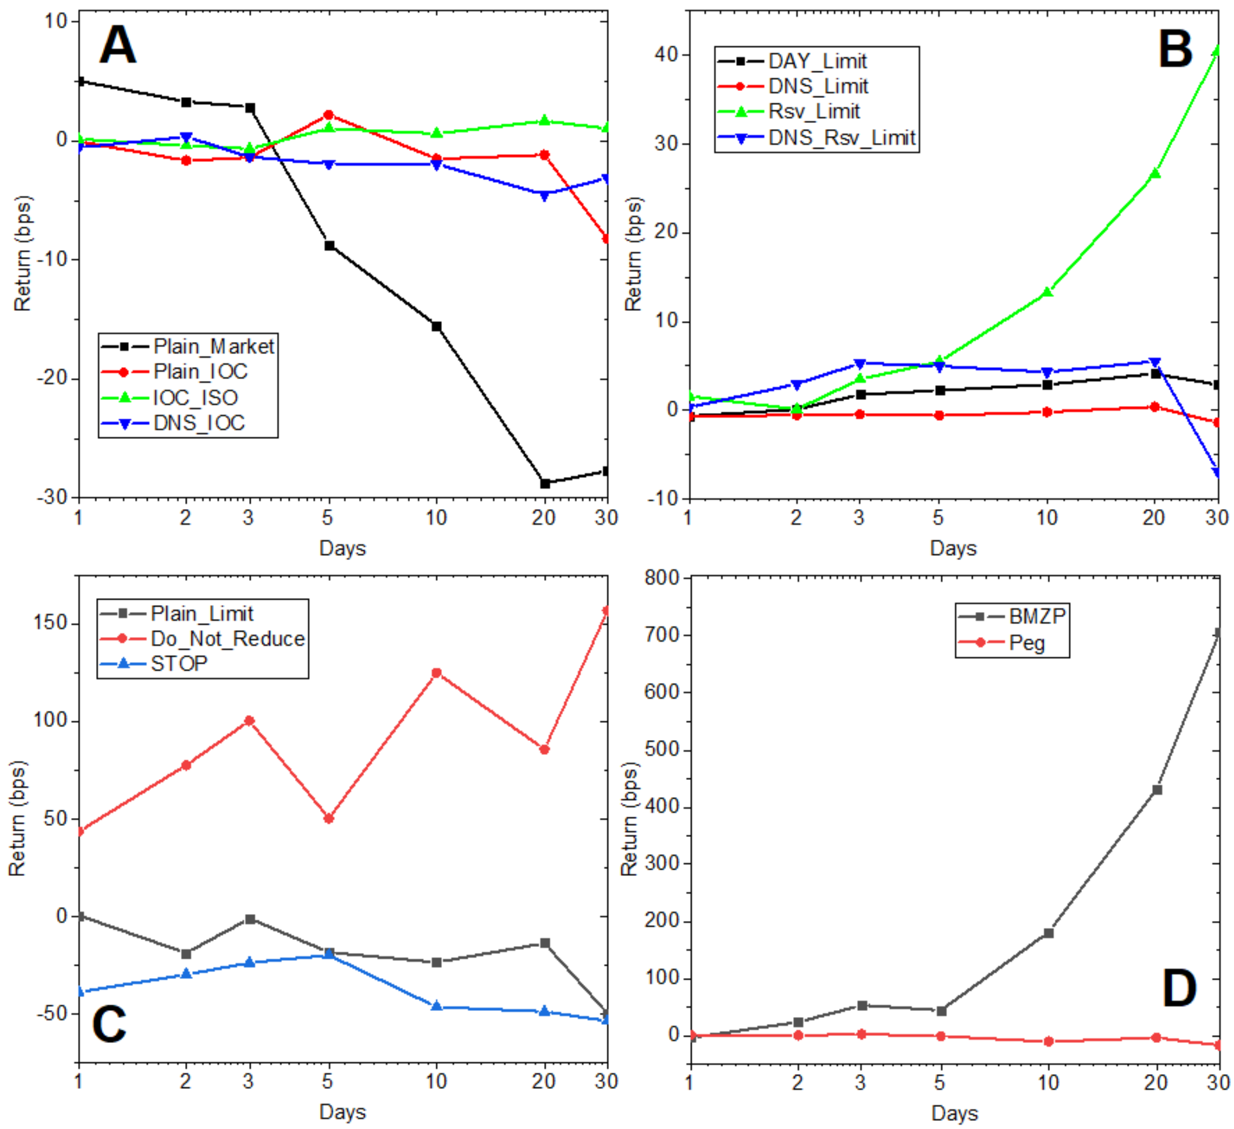

Appendix: Weighted Price Contribution

We define $r_n^{i,t}$ as the price difference between trade n and trade $n-1$, which measures the price contribution of the active side of an order. For example, $r_1^{i,t}$ is the difference between the first order and the price of the open auction. Define $\delta_{n,j} = 1$ if the order *n* belongs to order type j, and zero otherwise. The price contribution for order type j of stock i on day t is defined as:

$$PC_j^{i,t} = \frac{\sum_{n=1}^{N} \delta_{n,j} r_n^{i,t}}{\sum_{n=1}^{N} r_n^{i,t}} \tag{A.1}$$

The sum of the price contribution for all order types is 1. A trade that moves in the same direction as the daily movement positively contributes to price discovery. A trade that moves in the opposite direction negatively contributes to price discovery.

Next, we weight $PC_j^{i,t}$ across all stocks for each day. The weight of stock i on day t, $w^{i,t}$, is the ratio of the stock's absolute cumulative price change to the sum of all stocks' absolute cumulative price changes on day t. This weight mitigates the problem of heteroskedasticity, which may be severe for firms with small cumulative changes (O'Hara, Yao, and Ye (2014)).

$$w^{s,t} = \frac{\left|\sum_{n=1}^{N} r_n^{i,t}\right|}{\sum_{i=1}^{I} \left|\sum_{n=1}^{N} r_n^{i,t}\right|}. \tag{A.2}$$

We define the weighted price contribution of trades in size category *j* on day *t* as:

$$WPC_j^t = \sum_{s=1}^{S} \left(w^{s,t} PC_j^{s,t}\right). \tag{A.3}$$

Finally, we average WPC_j^t across dates, following O'Hara, Yao, and Ye (2014)

$$WPC_j = \sum_{t=1}^{T} WPC_j^t / T. \tag{A.4}$$