

NBER WORKING PAPER SERIES

FINANCIAL REGULATION, CLIENTELE SEGMENTATION,
AND STOCK EXCHANGE ORDER TYPES

Sida Li
Mao Ye
Miles Zheng

Working Paper 28515
<http://www.nber.org/papers/w28515>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2021

We thank Jim Angel, Dan Bernhardt, Ekkehart Boehmer, Colin Clark, Carole Comerton-Forde, Nicolas Crouzet, Ian Dew-Becker, Thomas Ernst, Harry Feng, Joel Hasbrouck, Edwin Hu, Pankaj Jain, Phil Mackintosh, Bruce Mizrach, Dermot Murphy, Maureen O'Hara, Marios Panayides, Steven Poser, Jeffery Smith, Elvira Sojli, Chester Spatt, Jeremy Stein, Kumar Venkataraman, Sunil Wahal, and seminar participants at the Louisiana State University, the New York Stock Exchange, the University of Illinois at Urbana-Champaign, Rutgers University, AFA 2021, and Microstructure Online Seminars Asia Pacific for helpful discussions and suggestions. We thank the New York Stock Exchange for providing us with their proprietary data for this paper. Ye acknowledges support from National Science Foundation grant 1838183 and the Extreme Science and Engineering Discovery Environment (XSEDE). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Sida Li, Mao Ye, and Miles Zheng. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Financial Regulation, Clientele Segmentation, and Stock Exchange Order Types
Sida Li, Mao Ye, and Miles Zheng
NBER Working Paper No. 28515
February 2021
JEL No. G14,G18

ABSTRACT

Financial regulations and clientele segmentation explain the proliferation of order types on stock exchanges. Plain market and limit orders lose money, indicating that informed traders use complex orders. Fifty-seven percent of trading volume comes from non-routable orders, which are designed to bypass Reg NMS. Because Reg NMS routes orders based on the best gross prices, it often routes orders to worse net prices after adjusting for fees. Non-routable orders win speed races to capture short-term profits, but all order types containing long-term information are routable. An order type that complies with share repurchase regulations earns a 30-day return of 7%.

Sida Li
Department of Finance,
Gies School of Business
University of Illinois, Urbana-Champaign
1206 South 6th Street
Champaign, IL 61820
sidali3@illinois.edu

Miles Zheng
Department of Finance,
Gies School of Business
University of Illinois, Urbana-Champaign
1206 South 6th Street
Champaign, IL 61820
yzheng60@illinois.edu

Mao Ye
Department of Finance,
Gies School of Business
University of Illinois, Urbana-Champaign
343K Wohlers Hall
1206 South Sixth Street
Champaign, IL 61820
and NBER
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 choices between market and limit orders.¹ Using proprietary message-level data from the New York Stock Exchange (NYSE), we find that traditional market and limit orders account for less than 10% of trading volume, and both types lose money. Market orders, which are orders that do not specify prices, lose 24.48 bps at the 30-day horizon. Limit orders that specify prices but offer no further instructions lose 49.72 bps at the 30-day horizon. As the returns on all order types are zero-sum, informed traders must use more complicated order types than traditional market and limit orders. Understanding the diversity of order types and their origins is, therefore, important for comprehending the trading ecosystem.

Our paper makes two contributions to the literature. First, we provide the first anatomy of exchange order types. Our results identify 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 win speed races for small profits, but they contain no long-term information. 3) All order types that earn long-term profits are neither too simple nor too complex. Therefore, beyond plain market or limit orders, sophistication in short-term trading and sophistication in long-term investment are substitutes but not complements.

Second, we find that financial regulations and clientele segmentation are the main drivers of the proliferation of order types, as most complex order types are designed to help stock exchange clients comply with or bypass existing regulations. One main driver of exchange order types is Regulation National Market System (Reg NMS), which directs all 13 U.S. stock exchanges to consolidate their quotes, establish the national best bid and offer (NBBO), and route incoming

¹ 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. Parlour (1998)). Harris and Hasbrouck (1996) compare limit order performance with market order performance on the NYSE. Parlour and Seppi (2008) provide an excellent survey of the limit order book literature.

orders to the exchange that displays the NBBO. Surprisingly, we find that 57% of trading volume involves order types that enable traders to bypass Reg NMS. That is, they refuse Reg NMS to route them to the best price.

One driver of this counterintuitive finding is the definition of the best price under Reg NMS. Reg NMS defines the NBBO based on the gross price, but traders must also pay fees to exchanges. We find that Reg NMS often routes orders to worse net prices after adjusting for fees. Rule 610 of Reg NMS prohibits a displayed order on one exchange from locking or crossing an existing quote on 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, the NYSE limit sell order at \$9.99 locks the NASDAQ bid. NYSE then needs to comply with Rule 610 by routing the limit sell order at \$9.99 to take liquidity from NASDAQ. Such routing leads to the same gross price but worse net price 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. The worse net price explains why 55% of non-marketable orders come with “do-not-ship” (DNS) instructions, which comply with Reg NMS by canceling an order if it locks or crosses an existing quote on another exchange. We find that DNS limit orders earn a small profit of 1.57 bps after collecting the rebate but would lose 0.89 bps if they paid the routing fee.

Reg NMS creates Intermarket Sweep Orders (ISOs), an order type that asks an exchange not to check prices on other exchanges. Instead, traders comply with Reg NMS by sending ISOs simultaneously to all other exchanges that offer better gross prices. ISOs were initially designed to allow institutions to use large orders to access liquidity in multiple markets simultaneously at heterogeneous price levels.² Surprisingly, we find that ISOs have an average execution size of 244

² See the SEC’s FAQ for Rules 611 and 610 of Reg NMS, page 3: “*If an order-router wishes to execute a large order by sweeping both the protected quotations and depth-of-book quotations at one or more preferred trading centers, it can implement this strategy by routing large-sized IOC/ISOs with aggressive limit prices to the preferred trading centers, while routing additional IOC/ISOs, as necessary, to less preferred trading centers that are priced and sized to execute against only the protected quotations of those trading centers.*” The FAQ is available at: <https://www.sec.gov/divisions/marketreg/rule611faq.pdf>.

shares, which is the smallest across all liquidity-taking orders. Also, only 0.53% of ISOs are executed at multiple price levels, whereas 5.46% of market orders are so executed. Speed provides an explanation for why ISO clients refuse exchange routing because routing can lead to geographic, consolidation, and transmission latencies (SEC (2018)). 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, where speed is the key to sniping those stale quotes (Budish, Cramton, and Shim (2015)). As ISOs often snipe stale quotes, they are informative at the short-term horizon. We find that ISOs contribute more than 90% of price discovery on a trade-by-trade basis.

We also find that non-routable orders win speed races to make liquidity. First, DNS limit orders successfully cancel stale quotes 42.75% of the time, but routable limit orders rarely escape from sniping. The adverse selection cost of DNS limit orders is 0.26 bps lower than that of 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 (2021)), we find that the return on DNS limit orders is 0.76 bps higher than that on routable limit orders at the time of execution. The ability to add new quotes and cancel stale quotes quickly provides DNS limit orders with 1.02 bps higher returns than routable limit orders 30 seconds after execution.

We find that, although DNS limit orders and ISOs make profits in the short run, all order types that earn positive returns beyond a day accept exchange routing. The order type that is most informed over the long term is also a consequence of financial regulations. To prevent firms from inflating their share prices, SEC Rule 10b-18 discourages firms from buying at prices that are higher than the highest independent bids or the most recent transaction prices. The NYSE offers buy-minus-zero-plus (BMZP) orders to comply with Rule 10b-18.³ We find that BMZP orders are

³ Specifically, Rule 10b-18 is a “safe harbor” provision which reduces liability for firms’ share repurchases. The price condition of Rule 10b-18 requires that the repurchase price should not exceed the highest independent bid or the last independent transaction price, whichever is higher. In its SEC filing, the NYSE states: “*The BMZP instruction is designed to assist member organizations in their compliance with the “safe harbor” provisions of Rule 10b-18 under the Act (“Rule 10b-18”) for issuer repurchases.*” Available at <https://www.sec.gov/rules/sro/nyse/2016/34-78679.pdf>.

associated with share repurchases and realize 30-day returns of 706 bps, which is consistent with that repurchasing firms are informed traders in their own stocks (Dittmar and Field (2015)). BMZP orders are routable, which indicates that firms and their brokers care more about filling their orders than about the exchange fees or order execution speed that is measured in milliseconds.

Our most telling 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 exchange routing, while DNS reserve limit orders refuse exchange routing. Upon execution, DNS reserve limit orders earn a return that is 0.27 bps higher than regular reserve orders earn. 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. Therefore, two order types that differ by only one instruction can have completely distinct clienteles and return patterns. Clients of routable reserve orders seem to hide long-term information and they accept exchange routing once another exchange offers an acceptable price. Clients of DNS reserve limit orders, however, refuse matching opportunities because they may fail to break even after paying routing fees. This result highlights the importance of decomposing orders into their detailed instructions. One main contribution of our paper lies in providing the first anatomy of complex order types.

Reg NMS is the fundamental regulatory framework that has shaped U.S. market structure since 2005 (Hasbrouck 2017). The core of Reg NMS, as suggested by its name (“National Market System”), is linking exchanges through routing. Our paper provides the first analysis of exchange routing, thereby extending the market microstructure literature to linked exchanges. Exchange linkage and routing add an entirely new dimension to the literature. For example, one key concept in market microstructure is the distinction between making and taking liquidity. Exchange linkage blurs this difference, because an order that displays liquidity on one exchange may become marketable on another exchange. We find that routing seldom improves prices for order types that are able to display liquidity on local exchanges. For example, 77.49% of limit orders are routed to the same gross price or a worse net price. On the other hand, routing improves prices for 100% of orders that cannot display liquidity, such as market and Immediate-Or-Cancel (IOC) orders. As

these orders refuse to display liquidity on any exchange, they cannot lock or cross the quote on another exchange and are exempted from Rule 610. Routing of these orders is directed by Rule 611 of Reg NMS, known as the “Order Protection Rule” or the “No Trade-Through Rule,” which restricts execution of trades at one venue at prices that are inferior to displayed quotations at another venue. Our results indicate that Rule 611 protects orders that refuse to provide liquidity. Once an order can display liquidity locally, it is subject to both Rule 610 and Rule 611. We find that Rule 610 dominates Rule 611 in exchange routing for these orders. As 77.49% of limit orders are routed to the same gross price or a worse net price, Rule 610 and exchange linkages incentivize traders to refuse routing.

Reserve orders provide a further contrast between Rule 610 and Rule 611. A completely hidden order is subject only to Rule 611, because it cannot lock displayed quotes. We find that routing improves prices for 99.95% of completely hidden orders. A partially hidden order is subject to both Rule 610 and Rule 611, because its displayed component can still lock displayed quotes. We find that only 18.50% of partially displayed orders have price improvements. Thus, although one goal of Reg NMS is to encourage traders to display quotes through Rule 611 (SEC (2005)), Rule 610 creates an incentive for traders to hide their quotes. Exchange linkage and Reg NMS results in the following tradeoff: displayed orders may execute with marketable orders from other exchanges because of Rule 611, but they may lock quotes on other exchanges because of Rule 610. This new trade-off under linked exchanges adds to the trade-off between hidden and displayed liquidity in isolated exchanges that is documented by Bessembinder and Venkataraman (2004), Bessembinder, Panayides, and Venkataraman (2009), and Chakrabarty et al. (2020).

Rule 610 and exchange linkages also provide an incentive for exchanges to compete for speed, because the exchange that locks the market is the exchange that displays a quote later. The cleanest evidence comes from the refill of reserve orders, because the refill process involves exchange latency but not trader latency. We find that a reserve order can be routed to another exchange

because of a small latency involved in refilling its displayed quotes.⁴ If a displayed order occurs on another exchange after the execution of the displayed part of a reserve order and before its refill, the refill may lock the market and NYSE has to route the reserve order out or cancel the order. The exchange's latency in refilling orders incurs costs for order submitters, because we find that 76.88% of routing led by refills execute at worse net prices. This latency also incurs costs for stock exchanges, because the slow exchange loses market share in trading volume and also pays take fees to the fast exchange. Therefore, the presence of Rule 610 and exchange linkages can explain why exchanges invest to reduce their own latency and adopt price-time priority to reward traders who establish fast quotes, even though the SEC allows the exchange to adopt alternative priority rules such as price-size priority and frequent batch auctions.⁵

Our paper is the first to explain the proliferation of order types. Combined with studies that examine other rules in Reg NMS, our paper also helps to explain the post-Reg NMS U.S. market structure. Spatt (2018) suggests that Rule 611 leads to market fragmentation because it protects only the top of the book. Chao, Yao, and Ye (2019) show that Rule 612, which imposes a minimum price variation of 1 cent per share, leads to the proliferation of exchanges that charge varying fees. Routing further widens the price gap between making and taking liquidity, because an exchange does not want to lose money when taking liquidity from another exchange. The NYSE's routing fee of 0.30 cents is exactly equal to the take fee cap imposed by Rule 610 on all exchanges. In this paper, we show that exchange fees 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 an exchange tends to route orders to another exchange with a common owner. Therefore, traders who are sophisticated in execution may prefer to route orders themselves. Yao and Ye (2018) find that Rule 612 leads to speed races to capture rents created by

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

⁵ See <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 reverted 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.

the tick size, and we find that DNS limit orders tend to win such speed races. Budish, Cramton, and Shim (2015) argue that market fragmentation leads to speed races to snipe stale quotes, and we find that ISOs tend to win speed races to snipe stale quotes while DNS limit orders are more likely to escape from sniping. In summary, Reg NMS leads to the proliferation of order types directly through Rules 610 and 611 and indirectly through its impact on exchange fees, market fragmentation, and speed races. Taken together, our paper and previous studies provide a roadmap that helps to explain how Reg NMS generates four features in the current market: high-frequency trading, market fragmentation, exchanges with heterogeneous make/take fees, and proliferation of order types.

Crouzet, Dew-Becker, and Nathanson (2020) model 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 over 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 of Crouzet, Dew-Becker, and Nathanson (2020). Traders who use non-routable orders are informed at the execution horizon but not at the investment horizon; however, long-term informed traders do not excel at the short horizon. In Crouzet, Dew-Becker, and Nathanson (2020), both short-term informed traders and long-term informed traders make money because noise traders bear the costs of trading. Our results show that clients using the simplest order types may be such noise traders.

Our paper focuses on order types on the NYSE, but the economic force revealed by our paper provides a roadmap to understanding order types on other U.S. stock exchanges. In Section VI, we briefly discuss cross-exchange variation and the time-series evolution of order types. Most new order types, such as slide orders, Hide Not Slide orders, and post-only orders, are designed to provide solutions to comply with Reg NMS. The remainder of the paper is organized as follows. In Section I, we describe our proprietary NYSE data. In Section II, we 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 Rule 610 and fees help to explain what motivates refusal

of 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 runs from January 1, 2010, to March 1, 2011, and we have a stratified sample of 109 stocks. 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 common stocks of non-U.S. companies. We also exclude stocks that, at the end of year 2009, are priced 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 clients 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’s information by observing the car.

We identify order types using price conditions, time-in-force (TIF), and special order instructions (SOIs). Each price condition, TIF, or SOI adds a function to an order type. The best way to understand the economic drivers of an order type is to decompose it into its basic functions. Such decomposition helps us generalize the economic intuition underlying our paper beyond the NYSE. Order types on other U.S. exchanges may have different names, but they offer similar

economic functions. The decomposition also helps us understand the evolution of order types on the NYSE. New order types supersede old ones by providing a better function for an economic problem. In Section VI, we provide a brief overview of order types across exchanges and their evolution.

A. Price conditions

Based on price conditions, we categorize orders into four types: market orders, limit orders, stop orders, and pegged orders. A market order does not have a specified price and it buys or sells at the best price obtainable. A limit order buys or sells a security at a specified price or better. For example, a limit order to buy 100 shares at a limit price of \$100 would buy a stock if its price is \$100 or below. A stop order, also known as a stop-loss order, buys or sells 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.

Pegged orders track reference prices, and they differ in the prices they track. NYSE primary pegged orders track the bid and ask prices, and traders may modify the reference price by choosing a specific number of ticks away from the bid and ask price, such as buying at the bid price minus one tick or buying at the bid price plus one tick. The reference price for the BMZP order is the higher of the current bid price *and* the last trade price. This reference price dovetails with the requirements of SEC Rule 10b-18, which discourages firms from repurchasing their shares at prices that are higher than the highest independent bids or the most recent transaction prices, whichever is higher. All BMZP orders are routable, which increases their fill rate. For example, if another trader takes liquidity from the ask side, the BMZP order will follow up, take the remaining liquidity at the same ask price on the NYSE, and the NYSE will route the remaining BMZP orders out if other exchanges can also offer liquidity at the last trade price. Pegged orders may also include a limit price so that the order no longer pegs the reference price once it moves outside the price limit.

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 open/reopen/close auctions), or GTC (Good-Till-Cancel). An IOC instruction cancels an order if it fails to execute immediately. All IOC orders have limit prices, but they accept only the limit price and refuse to offer other traders the limit price. Therefore, IOC orders are more like market orders than limit orders. DAY orders have expiration times at the ends of regular trading hours, and they often include additional SOIs to form more complex order types. We do not find any salient return pattern of 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 main analysis.

GTC orders come in two variations. We describe GTC orders without any further instructions as plain limit orders because they do not impose expiration times or any further instructions. Such orders remain on the limit order book until execution or cancellation. The NYSE will help a GTC order submitter 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

Reg NMS were designed 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 an NMS stock. Suppose that the NBO is \$5.01 and the NBB is \$5.00, while the best offer on the NYSE is \$5.02 and the best bid on 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 an improvement because the bid price of \$5.01 would lock the NBO at \$5.01. Rule 610 would then 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 at one venue at prices that are inferior to quotes displayed at another venue. Therefore, an exchange is obligated 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 on which she submits her order can then directly execute the order against *local* orders at any price that is 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, 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 were to lock or cross the market. Surprisingly, we find that the displayed component 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.

II. Summary Statistics

Table 2 provides summary statistics for order types in our sample. Some combinations of order 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 or BMZP orders. Primary pegged orders come in two versions: one is routable and the other has a DNS tag. We combine these two

⁶ DNS IOC orders are also known as NMS IOC, SEC IOC, or SOC orders.

types because both types of primary pegged orders show similar patterns. Indeed, because primary pegged orders in our sample are designed 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 0.01% of routable pegged orders have been routed out.

[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 from any exchange, while ISOs and DNS IOC orders take liquidity only from the NYSE. Orders that participate in open and close auctions contribute to 13.70% of share volume, but we exclude them from our analysis because we do not find any salient patterns for auction orders. A DNS limit order adds a DNS 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 DNS and non-display instructions to a DAY limit order. For the sake of brevity, we omit the DAY label for DNS limit orders, reserve limit orders, and DNS reserve limit orders.

A. Market shares 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 15.66% of total share volume and DNS IOC orders account for 11.35% of share volume. Market orders and plain IOC orders, however, account for only 4.48% and 2.84% of share volume, respectively. Non-routable orders are also important for orders that are designed to provide liquidity. DNS limit orders take a market share that is similar to that of DAY limit orders (20.36% vs. 24.12%). DNS reserve orders take a market share of 1.85% and reserve orders take a market share of 4.84%. In total, non-routable orders account for 57% $\left(\frac{15.66+11.35+20.36+1.85}{100-9.82-2.02-0.73-1.13}\right)$ of continuous (non-auction) trading volume. Stop orders account for only 0.13% of share volume. The two types of pegged orders account for 0.61% of share volume.

B. Execution conditions for order types

One way to understand an order type is to examine the condition under which the order is executed. An order can be executed in three ways: providing liquidity on the NYSE, taking liquidity from the NYSE, or taking liquidity from other exchanges. These three types of execution incur differing exchange fees, and Rule 610 plays a central role in determining the differences between these fees. Rule 610 mandates a cap for a take fee of 0.30 cents per share. In our sample period, some exchanges, such as NASDAQ, NYSE Arca, and BYX, charge a take fee of 0.30 cents per share, which equals to the fee cap. When the NYSE routes an order to another exchange, it takes liquidity from the exchange; the NYSE charges a routing fee of 0.30 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 on the size of the rebate for liquidity-making orders. The rebate for making liquidity on the NYSE is 0.13 cents per share, and NYSE earn $0.21 - 0.13 = 0.08$ cents per share. For a trader, the fee difference between making liquidity on the NYSE and taking liquidity from another exchange is 0.43 ($=0.13+0.30$) cents per share. As exchanges that charge a 0.30 cent take fee can give more generous rebates to liquidity makers, the differences between the make and routing fees across all exchanges are close to twice the size of the take fee cap.

Table 2 provides the market shares of three types of execution: take local liquidity, route, and make liquidity. Primary pegged orders and BMZP orders include all three types of executions, although most executions involve making liquidity on the NYSE. In particular, primary pegged orders execute 99.28% of their share volume from making liquidity on 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 NYSE routing contributes to 14.28% of the volume of DAY limit orders. Reserve limit orders are slightly more likely to take liquidity on 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 volume.

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 on the NYSE. For example, DNS limit orders take liquidity for only 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 is the preference for making but not taking liquidity.

Market orders exhibit the highest routing rate among all 13 order types (33.64%). Stop orders have a similar route-out rate (31.65%) because stop orders, once triggered, become market orders. 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 on the NYSE because they refuse to make liquidity on the NYSE and refuse to take liquidity on 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.⁷ The top destination for NYSE routing is NYSE Arca. Across all routed-out volume, NYSE Arca accounts for 53.86% of the market share. This number is almost double the NYSE Arca's market share of the total trading volume (27.74%). 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. Routing to NYSE Arca, therefore, involves the lowest latency.

[Insert Table 3 about here]

III. The Three-tiered Ecosystem of Order Types

The paper is the first to study the diversity of stock exchange order types. It is therefore impossible to provide pairwise comparisons between all order types in one paper. Our initial step is to provide some categorization of order types. We can categorize order types through economic reasoning or based on empirical analysis of order types. Both methods point out the same three-tiered ecosystem of order types. We begin this section with a brief discussion of categorization based on economic reasoning and then provide empirical support for the three-tiered ecosystem of order types.

⁷ We compute the total trading volume in each exchange from TAQ data.

The simplest order types are plain market orders, plain limit orders, and stop orders. These order types can be found in regular brokerage accounts and academic textbooks. Most theoretical and empirical work on order types discusses the choice between market and limit orders (see Parlour and Seppi (2008) for a survey), and Easley and O’Hara (1991) analyze stop orders. We then divide more complex orders into two groups: non-routable orders and routable orders. Non-routable orders must come from traders with the most sophisticated execution technologies. For example, ISO clients should have low latency, otherwise they cannot take over the compliance obligation for Reg NMS from exchanges. Orders of medium complexity comprise order types that are routable but not among the simplest three types. These orders manage their execution to some extent but defer the search for the best price to the NYSE.⁸

Next, we present evidence for the three-tiered ecosystem of order types based on return patterns of order types. Because we do not have trader IDs, our results contain two types of noise. First, each order type can be used by multiple types of traders. For example, both retail and institutional traders can use DAY limit orders. We categorize DAY limit orders as of medium complexity because they include expiration times and retail traders cannot account for the large share volume of DAY limit orders.⁹ This simplification tends to bias estimates of return differences between order types toward zero. Therefore, the significant differences in returns between order types show the robustness of clientele effects. Second, we measure returns based on fixed horizons. As traders may close their positions before or after that horizon, we may find that a given order type earns no return even if users earn positive returns. This problem becomes more serious as the horizon stretches or when traders have more flexibility to choose when they close their positions. These two biases would tend to attenuate our results. Still, we find that certain order types generate significant returns at a long-term horizon despite such noise, which supports the hypothesis that users of certain order types have long-term information.

⁸ Our conversation with an industry expert also supports 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.

⁹ Kelley and Tetlock (2013) find that retail orders account for only ~6.9% of the total trading volume.

A. Short-term return

Figure 1 visualizes the average returns on each order type for horizons of less than five minutes after execution. We tabulate the results and their statistical inferences in Appendix A2. 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 have negative returns at time 0 because they pay bid-ask spreads upon execution. Market orders pay the highest average transaction cost of 4.51 bps. The routable IOC orders incur a lower average transaction cost of 4.32 bps, whereas non-routable IOC orders incur 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 by 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 or the midpoint prices decrease below the execution price after a sell). 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 that the difference is less dramatic for liquidity-making orders. 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 in Section V that differences within the same stock are much larger, because non-routable orders take a large market share for liquid stocks (stocks with lower return of making liquidity). Third, Figure 1 presents the return results calculated before exchange fees. Making liquidity on the NYSE has a cost that is 0.43 cents lower than that of taking liquidity from an outside exchange, while taking liquidity from the NYSE costs

0.09 cents less (by subtraction—0.30 routing fee – 0.21 take fee) than taking liquidity from outside. As fees reduce the cost of non-routable liquidity-making orders more than the cost of non-routable liquidity-taking 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 specified expiry dates (i.e. plain limit, Do-Not-Reduce, and stop orders) lose money in the short horizon. This is not surprising because these order types do not actively manage their executions. Panel D shows that primary pegged and BMZP orders earn returns of around 0 within the 5-minute horizon.

B. Long-term returns

Figure 2 plots the average returns on each order type for horizons of longer than end of day after execution.¹⁰ Although the previous subsection shows that non-routable orders contain short-term information, Figure 2 shows that all orders that contain long-term information are routable. 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. Long-term returns on DNS limit orders, however, are close to zero and sometimes negative. We find that reserve limit orders earn a positive 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 do not make money in the short run.

[Insert Figure 2 about here]

The champions of long-term returns are BMZP orders, which are routable. BMZP orders earn the highest 30-day return of 706 bps. Note that our results do not imply that any trader can profit from a BMZP order but imply that traders who use BMZP orders are more likely to be informed. Because BMZP orders comply with SEC Rule 10b-18, we find that they are associated with share

¹⁰ We present the results based on raw returns; our results of returns after adjusting for Fama French three or five factors are similar and available upon request.

repurchases by firms: in firm-months that include BMZP orders, firms conduct ongoing repurchase programs 81% of the time, but for firm-months that do not include 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. Therefore, the returns on any particular order type are driven by clients of that order type, not by the order type itself. The return of an order type provides us a signal for the clients of the order type, albeit a noisy one.

BMZP orders are all routable, probably because their main users, repurchasing firms, care more about filling orders and less about the exchange on which an order is filled. The high fill rate of BMZP orders is consistent with this hypothesis. The fill rate of BMZP orders is as high as 45.88% (Table 2), a percentage that trails only those of market orders and primary pegged orders. Note that firms face constraints in using market orders (Ye, Zheng, and Li (2020)) because market orders execute at the ask price may lose the safe harbor provided by Rule 10b-18 if its price is also higher than the most recent transaction price.

The other type of long-term informed order also relates to corporate payouts. Do-Not-Reduce orders earn a 30-day return of 156.67 bps. Although the market share for Do-Not-Reduce and BMZP orders are both small, their high long-term returns mean that they are economically significant. Traders with fundamental information might be the minority in the population.

Figure 1 from 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 the execution horizon. Figure 2 shows that stop orders and plain limit orders lose about 50 bps at the 30-day horizon but Do-Not-Reduce orders earn strong positive returns. One possibility is that naïve traders cannot deploy sophisticated execution strategies, whereas some long-term informed traders have weaker incentive to deploy sophisticated execution strategies because their returns are much higher than transaction costs.

IV. The First Driver of NBBO Refusal: Routing to Worse Net Prices

In this section we show that Reg NMS routes most non-marketable orders to worse net prices. Rule 610 and exchange fees therefore help to explain NBBO Refusal. To prevent a locked market, Rule

610 asks an exchange to route non-marketable orders to the same gross price, which leads to a worse net price after fee adjustments. We show the impact of Rule 610 and exchange fees in three steps. In subsection IV.A., we show that routing to worse net prices dominates routing of non-marketable orders. We then show the impact of Rule 610 on routing to worse net prices by comparing orders that are subject to Rule 610 and orders that are not. In subsection IV.B., we compare the price aggressiveness of routable and non-routable orders. Our analysis suggests that non-routable orders are more likely to lock the displayed market. In Subsection V.C, we show that DNS limit orders earn positive returns after collecting rebates but would have lost money had they paid routing fees. Therefore, the 0.43-cent difference in making and routing fees makes a significant difference.

A. Paying fees but receiving no price improvement

The fee difference between taking liquidity outside and making liquidity on the NYSE is less than one tick. Why do traders avoid paying 0.43 cents if routing can improve prices? The first step in addressing this puzzle is to examine orders that are routed to other exchanges. Therefore, the sample in this subsection contains only routable orders. The study of routable orders serves two purposes. First, examining the execution of routable orders provides the intuition regarding why non-routable orders refuse exchange routing. Second, our analysis of routing provides the literature with the first glimpse of exchange linkages created by Reg NMS.

Reg NMS was designed to achieve two main goals: providing the best price for liquidity-taking orders and encouraging liquidity-making orders on multiple markets to display and compete (SEC 2005). To examine the first goal, we first examine whether routing improves prices. We define a price improvement as the difference between an order execution price and its limit price, adjusted for the order direction.¹¹

The results reported in Panel A of Table 4 show that more than 77.49% of DAY limit orders

¹¹ As market orders do not have limit prices, we use the best available NYSE limit price at the time of routing as the benchmark for price improvement.

are routed to outside exchanges without price improvements, and the figure is as high as 79.51% for reserve limit orders. Therefore, routing does not improve prices for most non-marketable orders. To see that Rule 610 serves as the main driver of routing without price improvement, notice that routing always improves prices for IOC orders. As IOC orders cannot lock the market, the NYSE routes IOC orders based on Rule 611 but not Rule 610. The results reported in Table 4 show that routing by Reg NMS is not homogeneous. Although routing by Rule 611 improves prices, routing by Rule 610 does not.

[Insert Table 4 about here]

Reserve orders of *zero* display size and reserve orders of *positive* display size also enable us to compare routing led by Rule 610 and Rule 611. The results are reported in Panel B of Table 4. Because Rule 610 applies only to displayed quotes, reserve orders of *zero* display size cannot lock the market, and we find that they receive price improvements 99.95% of the time. Once the displayed size is non-zero, the case for zero price improvement rises dramatically from 0.05% to 81.50%. Therefore, although the second goal of Reg NMS is to encourage traders to display their quotes, Rule 610 creates an incentive for traders to hide their quotes completely to avoid routing with zero price improvement. Certainly, a completely hidden order loses the privilege of attracting liquidity demanders from other exchanges through Rule 611. Therefore, Rules 610 and 611 create a trade-off between displaying quotes and hiding quotes.

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 component of PUBQTY = 100 shares at 9:43:10:100. The order was executed and refilled on the NYSE until the displayed 100 shares were consumed at 9:50:47.752. Three milliseconds later, the NYSE routed the order to NASDAQ and the Boston Stock Exchange to take liquidity at \$6.04 instead of refilling the order. Note that the NASDAQ and Boston bid prices must occur only within the 3-millisecond window; otherwise, the NASDAQ bid would lock the NYSE ask. Therefore, the reserve sell limit order is routed out as a result of the latency in refilling the shares. We find that 76.88% of the routing orders led by refilling experience no price improvement.

Routings during refills provide a clean environment in which to examine the impact of exchange latency. Holding all other things equal, Rule 610 implies that slow exchanges tend to lock the market for fast exchanges. It is extremely difficult to test this hypothesis in other setups, unless 1) we can obtain data from two exchanges with perfectly synchronized timestamps and 2) we can separate trader latency from exchange latency. Routing during refills helps us surmount these two challenges because 1) refills involve only exchange latency but not trader latency and 2) we know that exchange latency causes the routing because a refilled order cannot be routed out if exchange latency is zero. Although routings during refills are rare events, affecting only 0.052% of the volume, we believe these clean examples can be generalized to the impact of exchange latency in other scenarios. Under Rule 610, exchange latency costs traders, because we find that 76.88% of orders are routing to worse net prices because of exchange latency. Exchange latency also costs slow exchanges, because they need to pay the take fee to fast exchanges and also lose trading volume.

Therefore, Rule 610 provides an incentive for arms races on speed across exchanges. Existing studies of speed competition focus on traders.¹² Within this framework, exchange latency does not play a role as long as all traders experience the same delay. Our results suggest that Rule 610 incentivizes stock exchanges to be fast.

In summary, Rule 610 creates two incentives for non-marketable orders to refuse exchange routing. First, it blurs the definition between making liquidity and taking liquidity by precluding a locked market. Under Rule 610, an order that aims to make liquidity on one exchange may end up taking liquidity from another to avoid locking the market. Second, Rule 610 mandates a take fee cap of 0.30 cents. Many exchanges choose take fees that are exactly equal to the cap and provide generous rebates to liquidity makers close to the cap (Spatt 2020). The difference in fees between making liquidity on one exchange and taking liquidity from another exchange is close to twice the size of the take fee cap. Two policy proposals are designed to amend Rule 610 to narrow the differences in net prices: (1) reducing the take fee cap under Rule 610 (SEC 2018); (2) removing

¹² The only exception is Pagnotta and Philippon (2018), who argue that fast exchanges subsidize slow exchanges. Their model focuses on Rule 611 but not Rule 610.

the rebate (Harris 2018).

B. Price aggressiveness for routable and non-routable orders

In this subsection, we find evidence that DNS limit orders and DNS reserve limit orders are more likely than their routable counterparties to lock or cross the market. In Table 5, we show the price aggressiveness of non-routable orders relative to their routable counterparties. The most aggressive orders walk the book and take liquidity at multiple prices, followed by orders that take liquidity at one price. The most aggressive non-marketable orders aim to improve the NYSE BBO followed by orders at the BBO and then orders away from the BBO. Notice that marketable orders cannot lock the market.¹³ An order that is placed at the BBO or behind the BBO also cannot lock or cross the market. An order can lock the market only when it aims to improve the BBO on the NYSE. That is, it can lock the market when it is more aggressive than existing limit orders on the NYSE but not aggressive enough to take liquidity from the NYSE.

The results we report in Table 5 show that DNS limit orders are more likely to improve the best bid and ask on the NYSE than DAY limit orders (9.79%+2.08% vs. 1.76%+0.93%), and DNS reserve limit orders are more likely to improve the quotes than reserve limit orders (5.19%+3.73% vs. 1.79%+0.98%). These comparisons indicate that non-routable orders are more likely to lock the market.

[Insert Table 5 about here]

Among DNS limit orders that improve prices, we find that 82% ($= \frac{9.79}{2.08+9.79}$) of them improve the price by only one tick. This choice is natural not only because one tick is the least costly way to improve the quotes but also because the average spread in our sample is two ticks, leaving little room for improving quotes by two ticks. Finally, the following economic reasoning shows that routing cannot improve the price for an order that improves a local quote by one tick.

Suppose that the NYSE best ask is \$9.99 and the best bid is \$9.97. A buy limit order can improve

¹³ If an order is partially marketable, it can lock the market only when the non-marketable component improves the BBO on the NYSE, which follows the same mechanism described in the next few sentences.

the NYSE best bid to \$9.98. 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 on 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 crossed or locked market. Therefore, when the buy limit order is routed outside, it will take liquidity at \$9.98, experience no price improvement, but yet pay the routing fee and lose potential rebates for making liquidity on the NYSE.

C. Returns before and after fees

After showing that Reg NMS is unlikely to improve prices 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. For an average stock-day, DNS limit orders earn a return of 1.57 bps (0.31 bps) at the 1-second (5-minute) horizon after collecting the rebate of 0.13 cents, but lose -0.89 bps (-2.15 bps) if they pay routing fees of 0.30 cents. The negative profit led by routing fees gives DNS limit order clients incentives 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, their routings are driven by Rule 611 but not Rule 610. Table 4 shows that exchange routing following Rule 611 leads to a price improvement of one cent or higher for all executed market and plain IOC orders. As a price improvement of one cent or higher overwhelms the differences in fees, we search for additional motivations for refusing exchange routing. 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 (SEC (2005), Chakravarty et al. (2012)). Our results reported in Table 5 show, however, that only 0.53% of ISOs sweep multiple price levels, which is a much lower percentage than for IOC orders (1.32%) or market orders (5.46%). Also,

the results reported in Table 2 show that the average size of an ISO is only 244.30 shares, which makes these orders smaller than either plain IOC orders (268.63 shares) or market orders (278.96 shares). ISO's lower price aggressiveness and smaller size indicate that the main driver of ISO usage diverges from the SEC's purpose for designing ISOs.

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 speed races to take liquidity, that is, they are more likely to be the first to snipe stale quotes. In subsection IV.B., we show that non-routable limit orders tend to win two types of speed races to make liquidity. First, non-routable limit orders are more likely to cancel stale quotes. Second, non-routable limit orders are more likely to win races 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 33% of the time before any sniper arrives. To include speed races in which liquidity providers win, we extend the definition of 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 targets of sniping.¹⁴ 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

¹⁴ We exclude reserve orders from speed races for two reasons. First, reserve orders, which we study in the next subsection, are unlikely to win speed races to provide liquidity. Second, if a reserve order completely hides its quotes, liquidity-taking orders may take extra steps to detect the order, a strategy that lies 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.

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 discussion of the case in which no liquidity takers win to the next subsection.

The results reported 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. With respect to this number, notice that plain IOC orders win 13.42 races and lose 17.37 races. We exclude double counting where a plain IOC order wins the speed race over another plain IOC order. The participation rate of a plain IOC is about 10% ($\frac{13.42+17.37-5.96}{244.76}$). The participation rate for ISOs is as high as 77% ($\frac{134.51+140.71-86.35}{244.76}$), and the figure is 58% ($\frac{93.88+2.99-43.56}{244.76}$) 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 than plain IOC orders. For example, ISOs win 7.03 races and lose 4.47 races when they compete with plain IOC orders. The odds that ISOs win this race are $\frac{7.03}{7.03+4.47} = 61\%$. DNS IOC orders win 4.38 cases over plain IOCs but lose 2.99 cases. The odds of winning are $\frac{4.38}{4.38+2.99} = 59\%$. 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 lies in compliance requirement. To comply with Reg NMS, clients of ISOs need to have fast connections to all exchanges, including the smallest one. Clients of DNS IOCs do not take over compliance and they do not need fast connections with all 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 speed races 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 A1 provides the details on how we calculate the weighted price contribution.

Results reported in Table 7 show that ISOs are most informative both in the aggregate and on a per-share basis: ISOs account for 35.69% of the volume but 90.32% of price discovery.¹⁵ 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 contribute only 3.86% of price discovery, and their price contribution is much smaller than their share of liquidity-taking volume (6.36%). Plain market orders contribute negatively to price discovery. In summary, for liquidity-taking orders, orders with non-routing instructions are more informative on a trade-by-trade basis.

[Insert Table 7 about here]

B. Speed races for liquidity-making orders

Speed races involving liquidity taking lead to immediate profits, but the link between speed races and liquidity making is indirect. In this subsection, we analyze two types of speed races for liquidity-making orders. In Subsection V.B.1., we show that DNS limit orders win speed races to cancel stale quotes. In Subsection V.B.2., we show that DNS limit orders win speed races to secure front queue positions. In Subsection V.B.3., we provide evidence of the economic benefits that accrue to winners of these two types of speed races. Again, we maintain that the main driver of the results is not because DNS limit orders are faster than DAY orders but because DNS limit order clients are faster than DAY limit order clients. 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 component has lower execution priority than the displayed order.

¹⁵ The trading volume percentage in Table 7 is higher than that reported in Table 2 because the weighted price contribution considers only orders that take liquidity.

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 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.¹⁶ Results reported in Panel A of Table 8 show that 380.37 orders are targets of sniping on an average stock-day. Among them, 253.25 are sniped.¹⁷ The results in Panel A also show 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 Aquilina, Budish, and O’Neill (2020).

[Insert Table 8 about here]

Results reported in Panel A of Table 8 show that DNS limit orders are more likely to be targets of stale-quote sniping: when the race starts, 274.37 DNS limit orders are on the limit order book. However, 117.29 DNS limit orders successfully cancel, leading to an escape rate of 42.75%. The escape rate for DAY limit orders is only 9.27%. Combining this result with the results reported in Table 5, we see that DNS limit orders are not only more likely to improve BBO quotes but also are more likely to avoid sniping.

B.2. Speed races to establish front queue position

The rapid cancellation of stale quotes can reduce losses for DNS limit orders. To profit from making liquidity, though, DNS limit orders should also win speed races to add liquidity. Li, Wang, and Ye (2021) model speed races to provide liquidity. Because Rule 612 of Reg NMS imposes a

¹⁶ 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 from institutional traders. It is unlikely, however, that two institutional traders arrive at the same time and submit IOC orders at the same price.

¹⁷ This number is higher than the total number of races documented in the previous subsection (244.76) because sometimes more than one order is sniped.

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 (2021). Their model shows that traders 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) is such that 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 reported in Panel B of Table 8 show that first responders are disproportionately DNS limit orders. On average, $\frac{495}{1652} = 30\%$ 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 races $\frac{374}{495} = 76\%$ of the time, while DAY limit orders win races $\frac{121}{495} = 24\%$ of the time, suggesting that DNS limit orders are more likely to achieve favorable queue positions.

B.3. The economic benefits of winning speed races

Because DNS limit orders are more likely to establish a front queue position and escape from being sniped, we expect that they earn higher returns than DAY limit orders. In this subsection we provide some preliminary evidence for the reward for winning speed races.

To control for cross-sectional and time-series variations, we compare the return differences for routable and non-routable orders for the same stock at the same day and draw statistical inferences at the stock-day level. Results reported in column (1) of Table 9 show that DNS limit orders earn

an effective spread that is 0.76 bps larger than that of DAY limit orders.¹⁸ Li, Wang, and Ye (2021) provide one interpretation that can explain why faster orders earn a wider 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 that DAY limit orders lose 0.19 bps one second after execution, whereas DNS limit orders earn 0.82 bps 1 second after execution.

[Insert Table 9 about here]

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

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

Our sample comes from NYSE from 2010 to 2011, but the economic mechanisms we reveal provide intuition to understand order types on other exchanges as well as the evolution of order types, because the regulations we discuss apply to all U.S. exchanges that compete for the same clientele. NYSE and other U.S. stock exchanges constantly revise their order types or introduce new ones to cater to their clientele, especially by providing more innovative solutions to Reg NMS. For example, clients of DNS limit orders aim to make liquidity; however, DNS limit orders suffer from two limitations, and stock exchanges later introduce new order types to address these two limitations.

First, we find that 11.43% of DNS limit orders take liquidity from the NYSE. The difference between the take fee and the make rebate, 0.34 cents, is smaller than the difference between the

¹⁸ The difference is wider than the difference shown in Figure 1 because DNS limit orders take a larger market share for high volume stocks, for which the bid-ask spread is narrower.

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

Second, DNS instructions cancel an order when the order locks or crosses a market. The NYSE resolved unnecessary cancelations by introducing day ISOs in 2014 and slide orders in 2016.¹⁹ A 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. As 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. 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. In addition to the NYSE, NASDAQ and the CBOE also offer slide orders.

In subsection V.B. we show that winners of the speed race to secure the front queue positions earn higher profits, and in subsection IV.A. we show that Reg NMS allows completely hidden orders to lock the market. The Direct Edge Exchange introduced Hide Not Slide orders by combining these two features. The Hide Not Slide instruction hides an order if it locks the market. Once the market unlocks, this order lights up. These orders enjoy time priority over slide orders

¹⁹ The official name of the order type is “non-routable limit orders,” but DNS limit orders are also technically non-routable limit orders. Therefore, we rename this order type based on its economic function, which is to slide the quote if the quote locks the market.

because their timestamps are the times of entries. A 2012 *Wall Street Journal* article claims that these orders can jump ahead of the queue.²⁰ On January 12, 2015, the SEC announced that “two exchanges formerly owned by Direct Edge 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 based not 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.”²² 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.”²³ This penalty provides a unique case that summarizes four implications of our paper. First, two orders that differ by one instruction can have dramatically different economic functions and outcomes. Second, exchanges design order types to cater to their clienteles. 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 crucial in liquidity provision.

In our sample, there are two types of non-routable IOCs—ISO and DNS IOC orders—but only one type of non-routable limit orders—DNS limit orders. The NYSE and other exchanges

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>.

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>.

²² *Ibid.*

²³ *Ibid.*

introduced DAY ISOs to span the spectrum of order types. DAY ISOs resolve locked or crossed markets by sweeping quotes that lock or cross a 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 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 on the NYSE, the submitter complies with Rule 610 by sending ISOs to other exchanges to clear locked or crossed markets. Both DAY ISOs and ALO DAY ISOs make liquidity first by taking liquidity, but an ALO DAY ISO takes less liquidity. For example, DAY ISOs may need to execute against a large hidden order on the NYSE, whereas an ALO DAY ISO cancels the order. As of 2021, ALO DAY ISOs' market share is more than ten times higher than that of 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 order-type proliferation.

As a combination of these three new order types (ALO, slide, DAY ISO) provides better solutions for complying with Rule 610 than DNS limit orders do, 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 cater mostly to traders with the most sophisticated execution strategies. Exchanges gradually fill in the spectrum of order types by designing new order types that can serve investors with less sophisticated execution strategies. We show that IOC orders, particularly non-routable IOC orders, are often used to snipe stale quotes. The NASDAQ introduced the midpoint extended life order (M-ELO) in 2018. This order type refuses to interact with IOC orders, although it welcomes 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. Although M-ELOs cannot restrict counterparties based on their identities, the order types used by the counterparties convey information about their identities and trade incentives, which is the main intuition we aim to convey in this paper. Another order type that may serve less sophisticated

traders is the Discretionary limit (D-limit) order introduced by the Investors Exchange (IEX). The IEX runs an algorithm designed to dynamically forecast adverse selection risk. When adverse selection risk is low, a D-limit order behaves like a regular limit order. When adverse selection risk is high, the IEX automatically reprices a D-limit order to one tick worse than the NBBO. D-limits thus aim to reduce sniping risk by dynamically adjusting their limit prices.

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 now requires that all orders have a limit price or peg to a price and have an expiration time. Retail traders can still see these obsolete order types in their brokerage accounts, 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. A broker may assign a working price to a market order before entering it on an exchange, and the broker will reenter a GTC order to an exchange at the market open after the exchange cancels it at the previous close. The market structure in the U.S. evolves so quickly that order types that most people believe to be dominant are actually obsolete, whereas the true dominant order types operate outside most participants' radar. We hope that our paper provides a starting point for understanding complex order types and reveals their economic drivers.

VII. Conclusion

Any stock-trading strategy must choose an order type to interact with a market. In this first comprehensive study of exchange order types, we identify a three-tiered world of order types. Plain market, limit, and stop orders lose money and are most likely to come from naïve traders. Sophisticated traders manage their order types but to varying degrees. Traders who use 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 (gross) prices incur higher transaction costs and do not earn positive 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.

Rule 610 and exchange fees provide one incentive for orders to refuse exchange routing. 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 leads to the same gross price and worse net price because the order collects a rebate of 0.13 cents if it makes liquidity on the NYSE but pays a routing fee of 0.30 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.

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 win speed races to capture small profits.

As 57% of trading volume originates in orders designed 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 as 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.²⁴ Meanwhile, although major exchanges are currently operating at separate locations, 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 trading landscape in the U.S., in terms of both fees and speed.

²⁴ 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>.

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.
- 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.
- 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.
- Hasbrouck, Joel, 2017, Securities Trading: Principles and Procedures, Working Paper, New York

University.

Harris, Lawrence, 2018, Comment letter on transaction fee pilot for NMS stocks.

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.

Li, Sida, Xin Wang, and Mao Ye, 2021, 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, *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.

Spatt, Chester, 2018, The new realities of market structure and liquidity: Where have we been? Where are we going? in *Public Policy and Financial Economics: Essays in Honor of Professor George G. Kaufman for His Lifelong Contributions to the Profession*, edited by Douglas Evanoff, Tassos Malliaris, and George Kaufman, 247-263.

Spatt, Chester, 2020, Is equity market exchange structure anti-competitive? Working paper, Carnegie Mellon University.

U.S. Securities and Exchange Commission, 2005, Final rule: Regulation NMS, File No. S7-10-04.

- U.S. Securities and Exchange Commission, 2018, Transaction fee pilot for NMS stock, File No. S7-05-18.
- 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, closed-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								
	<i>N</i>	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 the annual shares of a stock traded (in billions). *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%
Pegged	Primary pegged	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, the total share volume, the market share of the share volume, average order size, and average fill rates. We leave the fill rates for the last three types of orders blank because they do not have expiration times. The average order size includes the unfilled portion of an order. We also present the percentage of the executed volume that takes local liquidity, routes to away markets, or makes 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 the 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 the consolidated volume of all 13 exchanges.

Table 4
Price improvements by routing

Panel A: Price improvements for routed out orders

Improvement	Plain Market	Plain IOC	DAY Limit	Reserve
0	0.00%	0.00%	77.49%	79.51%
0.01	62.28%	70.49%	12.93%	9.43%
0.02	15.84%	12.00%	2.74%	3.52%
0.03	6.24%	4.52%	1.35%	1.68%
0.04	2.86%	2.46%	0.96%	1.32%
0.05	1.52%	1.59%	0.95%	1.33%
0.05+	11.27%	8.94%	3.59%	3.22%

Panel B: Price improvements for routings led by Rule 610 and Rule 611

Improvement	Reserve Orders		
	PUBQTY = 0	PUBQTY > 0 AND Route Out Upon Arrival	PUBQTY > 0 AND Route Out After Posted
0	0.05%	81.50%	76.88%
0.01	20.10%	9.16%	11.61%
0.02	6.63%	3.45%	3.37%
0.03	3.79%	1.63%	0.89%
0.04	17.77%	0.90%	0.60%
0.05	28.23%	0.66%	0.10%
0.05+	23.42%	2.71%	6.55%

In this table, we report price improvements for routable orders. Panel A presents realized price improvements by order type. For all order types except market orders, price improvement is defined as the difference between the execution price and the limit price. Market orders do not have limit prices, and for them price improvement is defined as the difference between the execution price and the best available price on the NYSE. In Panel B, we report price improvements for reserve orders categorized by whether they can lock the market and whether they are routed out upon arrival. Reserve orders that are partially published (PUBQTY > 0) may lock the displayed market, so their routing is subject to Rule 610 and Rule 611. Reserve orders with no published quantity (PUBQTY = 0) do not lock the displayed market, so their routing is subject only to Rule 611.

Table 5
Limit order book dynamics

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 BBO by > 1 tick	0.00%	0.00%	0.00%	0.00%	0.93%	2.08%	0.98%	3.73%
Improving BBO by 1 tick	0.00%	0.00%	0.00%	0.00%	1.76%	9.79%	1.79%	5.19%
Order Placement at BBO	0.00%	0.00%	0.00%	0.00%	6.86%	23.95%	5.49%	18.12%
Order 1 tick from BBO	0.00%	0.00%	0.00%	0.00%	4.71%	10.44%	2.81%	6.79%
Order > 1 tick from BBO	0.00%	0.00%	0.00%	0.00%	81.02%	49.46%	86.78%	62.80%

In this table, we report limit order book impacts by order types. Market orders, Plain IOCs, DNS IOC orders, and ISOs can affect the limit order book only by trading (either at one price level or at multiple price levels upon arrival). DAY limit, DNS limit, reserve orders, and DNS reserve limit orders can either be marketable or add liquidity to the limit order book.

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; we define the speed races in section IV.A. The order types displayed in the rows are race winners while the order types displayed in the columns are race losers. If $q > 1$ orders win (lose) a 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 equals 1. We describe how we calculate the weighted price contribution in Appendix A1.

Table 8
Speed races to cancel stale quotes and establish front queue positions

Panel A: 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%

Panel B: Establish front queue positions

Daily number of trades	1,652		
	DAY limit	DNS limit	Total
Number of refills within 0.1 seconds after a trade	121	374	495

In this table in Panel A we report the outcomes of speed races 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 IOC snipers arrive. All numbers are averaged at the stock-day level. For Panel B 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 a responding limit order is triggered by a 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 in the following columns we report the average numbers of first responders per day by order types. The market shares of responding orders are reported under the counts.

Table 9
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
<i>t</i> -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 reported 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 specific stock-day order type usage difference by (1) aggregating the share-weighted profits for each stock-day in our sample, and (2) pairing the stock-day observations and conducting a paired *t*-test between order types. Standard errors are clustered at both the firm and day levels.

Figure 1
Short-term return dynamics of order types

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 expiration dates, i.e. plain limit orders, Do-Not-Reduce orders, and Stop orders. Panel D shows the results for order types with limit prices that are conditioned on the prevailing market prices, i.e. BMZP and primary 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 the effective half-spread.

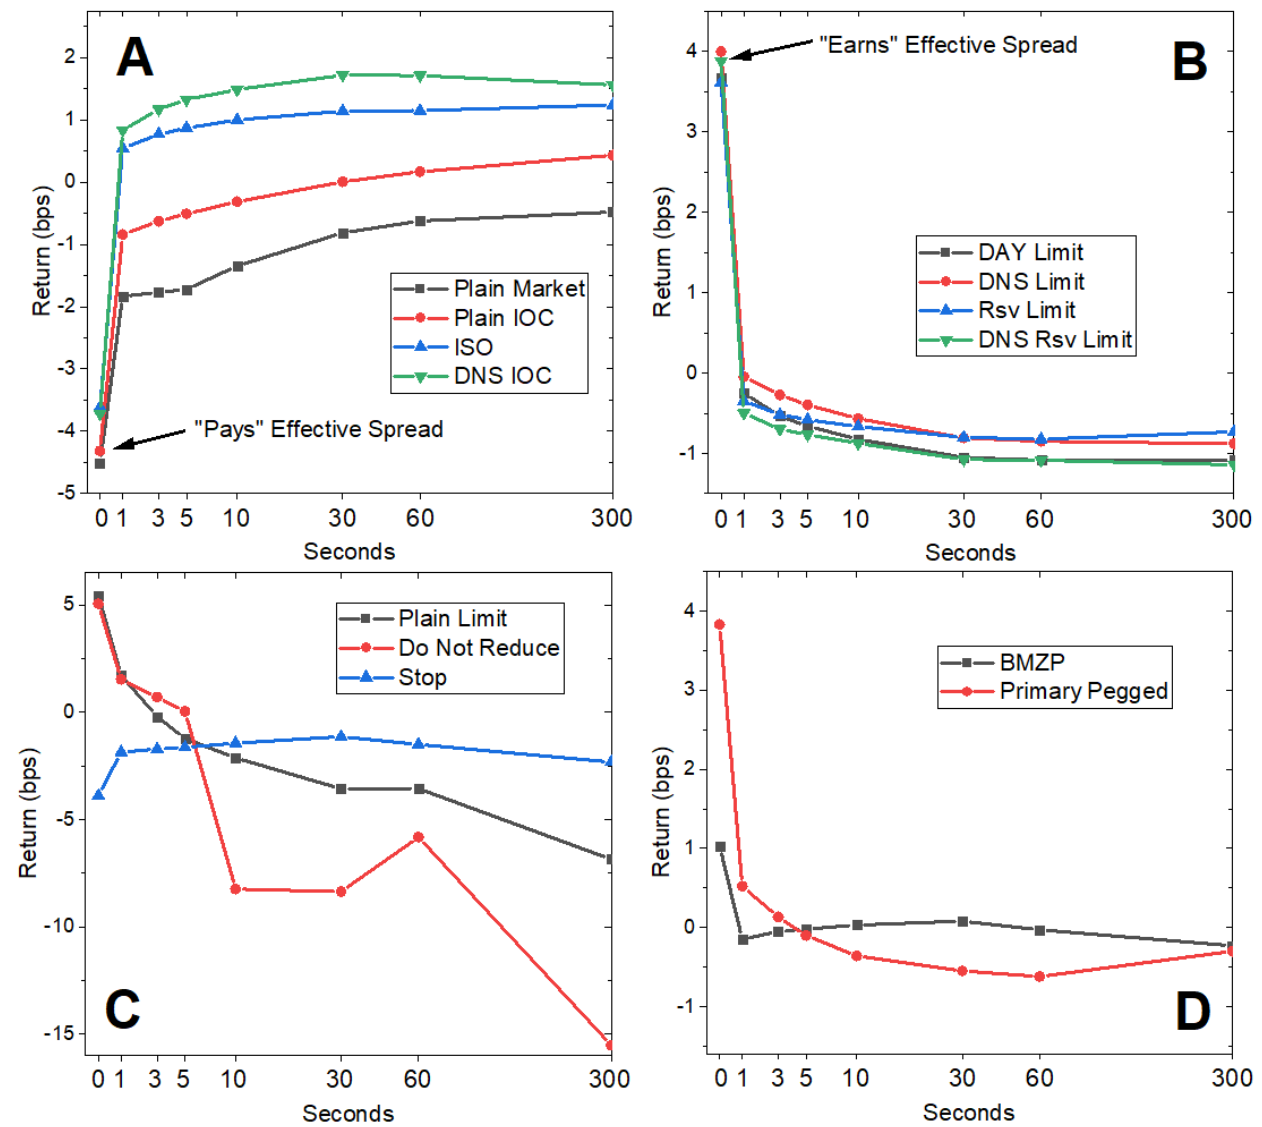
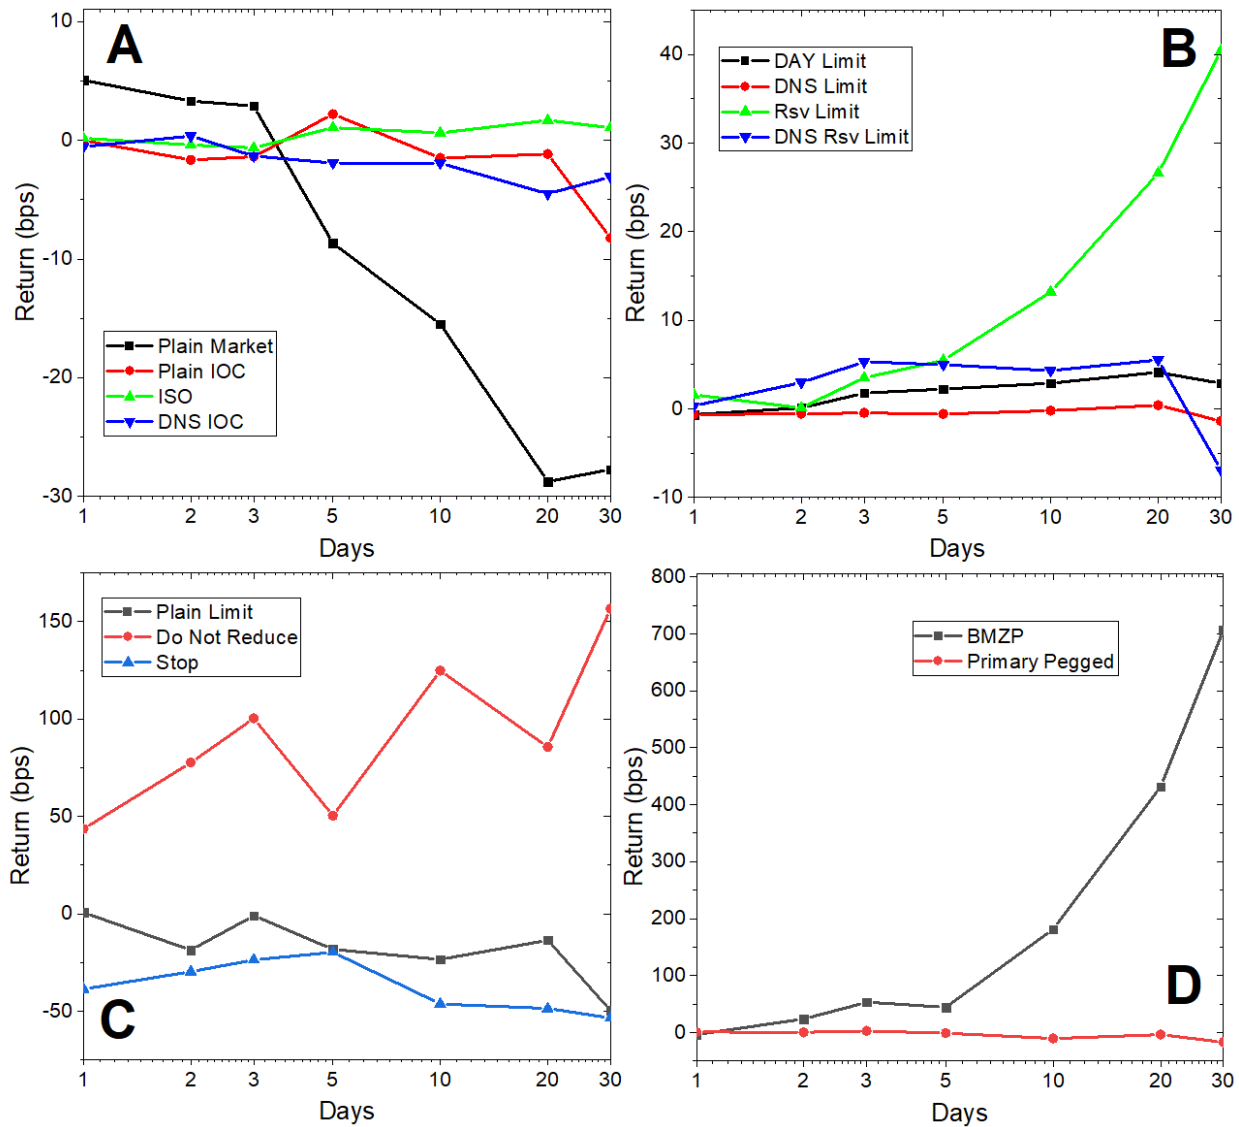


Figure 2
Long-term return dynamics of order types

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 expiration dates, i.e. plain limit orders, Do-Not-Reduce orders, and Stop orders. Panel D shows the results for order types with limit prices that are conditioned on the prevailing market prices, i.e. BMZP and Primary Pegged orders. Returns are measured as the differences between execution prices and the closing price each day after the execution, adjusted by trading direction.



Appendix A1: 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 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:

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

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

$$w^{s,t} = \frac{|\sum_{n=1}^N r_n^{i,t}|}{\sum_{i=1}^I |\sum_{n=1}^N r_n^{i,t}|}. \quad (\text{A.2})$$

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

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

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

$$WPC_j = \sum_{t=1}^T WPC_j^t / T. \quad (\text{A.4})$$

Appendix A2: Return dynamics of order types and statistical inferences

Order Type	0	1sec	3sec	5sec	10sec	30sec	1min	5min	EOD	2days	3days	5days	10days	20days	30days
Plain market	-4.51	-1.83	-1.77	-1.72	-1.35	-0.82	-0.62	-0.48	5.12	3.31	4.13	-7.34	-13.79	-25.99	-24.48
Plain IOC	-4.32	-0.84	-0.63	-0.51	-0.31	0.01	0.17	0.43	0.44	-1.65	-0.81	2.87	-0.42	-2.05	-10.08
ISO	-3.62	0.54	0.77	0.87	1.00	1.14	1.15	1.24	0.31	-0.40	-0.60	1.08	0.76	2.32	1.61
DNS IOC	-3.72	0.84	1.17	1.33	1.49	1.72	1.71	1.56	-0.43	0.35	-1.31	-1.81	-1.77	-3.98	-3.18
DAY limit	3.68	-0.25	-0.54	-0.66	-0.83	-1.05	-1.08	-1.08	-0.76	0.09	1.59	1.77	2.41	3.93	3.50
DNS limit	4.00	-0.04	-0.27	-0.40	-0.57	-0.81	-0.85	-0.88	-0.74	-0.56	-0.52	-0.66	-0.47	-0.21	-2.45
Reserve	3.61	-0.35	-0.52	-0.58	-0.67	-0.80	-0.82	-0.73	1.22	0.10	3.57	6.13	13.50	26.08	40.40
DNS Reserve	3.88	-0.50	-0.70	-0.77	-0.88	-1.08	-1.09	-1.14	-0.02	2.96	5.27	5.09	3.52	3.16	-8.23
Peg	3.83	0.52	0.13	-0.10	-0.36	-0.55	-0.62	-0.30	0.75	0.09	2.75	-1.28	-10.67	-3.78	-17.04
Buy Minus Zero Plus	1.03	-0.15	-0.05	-0.02	0.03	0.08	-0.03	-0.23	-3.60	23.64	53.24	44.22	181.55	431.85	706.19
Stop	-3.89	-1.86	-1.72	-1.64	-1.44	-1.14	-1.50	-2.33	-38.71	-29.77	-23.71	-19.73	-46.28	-48.68	-53.43
Good-Till-Cancel (Plain Limit)	5.44	1.72	-0.21	-1.23	-2.12	-3.58	-3.56	-6.86	0.58	-18.60	-0.99	-18.24	-23.51	-13.55	-49.72
Do-Not-Reduce	5.05	1.54	0.70	0.04	-8.24	-8.37	-5.82	-15.53	43.54	77.53	100.25	50.26	124.83	85.62	156.67

In this table, we present the return numbers underlying Figure 1 and Figure 2 and the statistical inference of returns. Returns are measured as differences between execution prices and midpoints, adjusted by trading direction. Numbers in bold are statistically significantly different from zero at the 1% level.