Refusing the Best Price? *

Sida Li

University of Illinois, Urbana-Champaign

Mao Ye

University of Illinois, Urbana-Champaign, and NBER

Miles Zheng

University of Illinois, Urbana-Champaign

Abstract

Regulation National Market System (Reg NMS) aims to link fragmented stock exchanges by routing orders to the exchange that displays the National Best Bid and Offer (NBBO). Surprisingly, we find that 57% of orders refuse Reg NMS routing. Fees provide one explanation, as Reg NMS routes 62% orders to worse net prices after adjusting for exchange fees. Speed provides another explanation, as non-routable orders win speed races to capture short-term profits. Our anatomy of order types shows a three-tiered world: the most and the least informed order types accept Reg NMS routing, whereas orders in between refuse Reg NMS routing.

Keywords: Regulation NMS, Order Types, Routing, Make/Take Fees, High-Frequency Trading

---

* 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, Albert Menkveld, 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, Citadel, Two Sigma Investments, AFA 2021, 3rd Future of Financial Information Conference, 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). Send correspondence to Mao Ye, the University of Illinois at Urbana-Champaign, 340 Wohlers Hall, 1206 South 6th Street, Champaign, IL, 61820. Email: maoye@illinois.edu

1. Introduction

One major function of markets is to help buyers and sellers to locate the best price. In reality, a series of frictions may prevent buyers (sellers) from transacting at the lowest (highest) price. They may not be aware of the best price, they may not have access to all markets, or the search cost for the best price is too high. Regulation National Market System (Reg NMS), which is passed in 2005, is probably one of the most aggressive efforts to remove these frictions. A firm listed in any U.S. stock exchange can be traded in more than ten other stock exchanges. Reg NMS establishes the national best bid and offer (NBBO) across these exchanges and designs rules that aim to route orders submitted to any exchange to the NBBO. Under Reg NMS rules, exchanges are interconnected, and each exchange serves only as a point of entry for the same destination − NBBO. Surprisingly, we find that 57% of the New York Stock Exchange (NYSE) trading volume comes from orders designed to refuse Reg NMS routing, and the incentive to circumvent Reg NMS is so high that it leads to the proliferation of order types.[1]

Why do the majority of orders refuse the Reg NMS routing to NBBO? The answer to this question is the key to understand the current U.S. market structure, as the establishment of NBBO and routing to NBBO are two cornerstones that differentiate the market structure of the U.S. from other counties.[2] However, due to data limitations, no existing studies ever evaluate the performance of routing to NBBO, let alone explaining why stock exchanges keep on designing new order types that refuse Reg NMS routing. By analyzing propriety *message-level* data from NYSE, our paper contributes to the literature by providing the first study on Reg NMS routing and the first anatomy of stock exchange order types.

[1] Mackintosh (2014) identified 133 order types across U.S. stock exchanges in 2014, the majority of which were designed after the implementation of Reg NMS with the purpose to refuse Reg NMS routing.
[2] For example, MiFID in Europe does not establish NBBO nor require an exchange to route orders to other exchanges (European Securities and Markets Authority 2019). Under MiFID, investment firms are responsible for best execution. Exchanges are not interconnected.

To understand the incentive for order types that refuse Reg NMS routing, the first step is to evaluate the performance of order types that accept Reg NMS routing. Surprisingly, we find that 62% of Reg NMS routings leads to worse net prices. We uncover two economic forces underlying this result. First, routing to the best price does not necessarily imply routing to a better price: Reg NMS routes 62% of orders to the same gross price. Second, Reg NMS defines NBBO based on gross price. As traders also need to pay routing fees, routing to the same gross price leads to a worse net price.

Why do Reg NMS routings lead to the same gross price for the majority of orders? The key is to understand the difference between two Reg NMS rules that govern routings: Rule 610 and Rule 611. The well-known reason for Reg NMS routing is to comply with Rule 611, which prohibits the execution of marketable orders at prices that are inferior to displayed quotes at other exchanges (no trade-through). We find that Rule 611 routings always lead to better gross prices, but they account for the minority (38%) of routing. The rest 62% of routing aims to comply with Rule 610. The name of Rule 610, "access to quotations", seems unrelated to routing. Most textbook introductions on Rule 610 covers only its first three clauses, which grant fair access of quotes and cap the access fee to 0.3 cents per share (e.g., Hasbrouck 2007, Foucault, Pagano, and Roell 2014). However, we find that the overlooked Rule 610 clause (d), which prohibits displayed quotes in one exchange to lock quotes in other exchanges (no locked market), strictly dominates Reg NMS routing.[3] Rule 610 routing to unlock the market leads to the same gross price. For example, suppose that the best NYSE ask price is \$10.00, the best NYSE bid price is \$9.98 and the best NASDAQ bid is \$9.99. A sell limit order submitted to NYSE at \$9.99 improves the NYSE's best ask price by one tick, but it locks the NASDAQ bid of \$9.99 and would be routed to NASDAQ to unlock the market by taking liquidity at \$9.99 without price improvement. One key concept in market microstructure is the distinction between limit and marketable orders, but exchange linkage blurs this distinction. Reg NMS aims to protect displayed limit orders and to improve prices for

[3] Rule 610 also prohibits a displayed order on one exchange from crossing the quote from another exchange (no crossed market). For example, NYSE cannot establish a bid price of \$9.98 when another exchange displays an ask price of \$9.97. We focus on locked quotes because crossed quotes are extremely rare.

marketable orders (SEC 2015). However, Rule 610 creates a third possibility beyond the dichotomy: a limit order that aims to display on one exchange may become marketable on another exchange.

Why does the same gross price lead to a worse net price? The key is to understand the stock exchange fee structure. Like the major exchanges in the U.S., NYSE has three types of fee schedules: make rebates, take fees, and routing fees. For example, at the end of our sample period, an order that makes liquidity at NYSE collects a make rebate of 0.15 cents per share, and an order that takes liquidity at NYSE pays a take fee of 0.23 cents per share. When NYSE routes an order outside, it takes liquidity from another exchange and charges a routing fee of 0.3 cents per share, which is equal to the take fee cap mandated by Reg NMS. Therefore, routing to unlock the market leads to a worse net price of 0.45 cents (0.15-(-0.3)).

To evaluate Reg NMS routing performance under Rule 610 and Rule 611, we first compare the routing performance between order types that can display liquidity and order types that cannot display liquidity. Because Rule 610 only governs orders that aim to display liquidity, an order that refuses to display liquidity cannot lock the market and thus are not subject to Rule 610. Therefore, we find Reg NMS routings always improve price for market order and Immediate-Or-Cancel (IOC) orders because their routings are subject only to Rule 611. In contrast, Reg NMS routing leads to worse net prices for 78% of routable limit orders because their routings are subject to both Rule 610 and Rule 611. Therefore, 55% of limit orders choose "do-not-ship" (DNS) instructions, which cancel the order if it locks the market. Refusing Reg NMS routing is essential to make profits from displaying liquidity, as we find DNS limit orders earn a small profit of 0.99 bps after collecting the rebate but would lose 1.80 bps if they paid the routing fee. In summary, although Reg NMS aims to encourage displaying liquidity through Rule 611, Rule 610 creates incentives to cancel liquidity.

Reserve orders, which are orders partially or fully hide their trading interests, enable further comparisons between Rule 610 routing and Rule 611 routing. Fully hidden orders are only subject to Rule 611 because they cannot lock displayed quotes. We find that routing improves prices for 99.95% of fully hidden orders. In comparison, partially hidden orders are subject to both Rule 611

and Rule 610 because their displayed part can lock displayed quotes. We find routings lead to worse net prices for 81.50% of partially hidden orders. In summary, although one goal of Reg NMS aims to provide incentives for traders to display quotes (SEC 2005), Rule 610 and exchange fees create incentives for traders to hide their quotes.[4]

Next, we analyze the impact of exchange fees on refusing Reg NMS routing. We first exploit cross-sectional variations of traders' incentives to refuse Reg NMS routing based on the bid-ask spread. As NYSE charges flat fees for all stocks in our sample, routing to worse net price reduces the liquidity-providing revenue more for stocks with lower bid-ask spreads.[5] Therefore, traders have higher incentives to use DNS limit orders for stocks with lower spreads. Indeed, we find that DNS limit orders have around 5% higher market share for stocks whose bid-ask spread is at the lowest tercile than for stocks whose bid-ask spread is at the highest tercile. Next, we exploit two exogenous shocks on exchange fees in our sample period. NYSE increased make rebates from 0.10 cents per share to 0.13 cents per share on May 1, 2010, and then to 0.15 cents per share on January 3, 2021, while keeping the routing fee constant at 0.3 cents per share. Therefore, the cost difference between making liquidity in NYSE and taking liquidity outside increased from 0.40 cents to 0.43 cents and then to 0.45 cents. We find that the market share of DNS limit orders on average increases by 2.03 (1.77) percentage points after the first (second) fee change. Finally, we conduct difference-in-difference analyses to test the heterogeneous treatment effects of fee changes. As the fee changes are relatively more important for stocks with lower spreads, we find that after the first (second) fee change, the market share of DNS limit orders increases 1.40 (1.21) percentage points more for the stocks whose bid-ask spread is at the lowest tercile than for stocks whose bid-ask spread is at the highest tercile.

[4] Exchange linkage and Reg NMS results in the following tradeoff between displayed and hidden orders: displayed orders have more opportunity to execute with orders from other exchanges, but they may be routed out to worse net prices. This new trade-off under linked exchanges adds to the trade-off between displayed and hidden liquidity in isolated exchanges documented by Bessembinder, Panayides, and Venkataraman (2009) and Chakrabarty et al. (2020).

[5] Penny stocks (stocks priced below $1) have lower fees. We do not have penny stocks in our sample.

After showing that Rule 610 and exchange fees create significant incentives for orders that aim to provide liquidity to refuse Reg NMS routing, we then show speed create significant incentives for orders that aim to take liquidity to refuse Reg NMS routing. Exchange routing can incur geographic, consolidation, and transmission latencies (SEC 2018). Traders can reduce these latencies through self-routing. For example, an Intermarket Sweep Order (ISO) asks an exchange to execute the order immediately without checking prices on other exchanges. Instead, ISO submitters take the obligation to comply with Reg NMS by sending ISOs simultaneously to all exchanges that offer better gross prices. Traders can also control their routings using NMS IOC orders, which ask an exchange to cancel the order if it needs to be routed out. We find ISOs (NMS IOCs) pay an average transaction cost of 3.62 (3.72) bps upon execution but make a profit of 0.54 (0.84) bps one second after execution. The results suggest that ISOs and NMS IOCs execute against stale quotes, and speed is the key to snipe stale quotes (Budish, Cramton, and Shim 2015). Indeed, we find that these two types of non-routable marketable orders win 94% of races in sniping stale quotes in NYSE.

Our data do not include the ID of submitters of orders, but we can provide an inference of the motivation and main clientele behind each order type by examining their empirical patterns. The motivation to design ISO, 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). We find, however, only 0.53% of ISOs sweep multiple price levels, which is a much lower number than IOC orders (1.32%) or market orders (5.46%). 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. The evidence that they win speed races to snipe stale quotes indicates that ISOs are more likely to come from high-frequency traders (HFTs) rather than from institutional investors.

For liquidity-making orders, the previous literature suggests two possible clientele who aim to collect rebates and avoid fees. The first type is HFTs who profit from market making (Brogaard,

Hendershott, and Riordan 2014). The second type is brokers who aim to maximize the rebates when they execute orders for their clients because brokers can pocket the rebates without passing them to their clients (Battalio, Corwin, and Jennings 2016). Our evidence suggests that HFT market-making is the main driver for refuse Rule 610 routing, as DNS limit orders tend to win speed races for making liquidity. DNS limit orders win 72% of races to establish time priority in liquidity provision, whereas routable limit orders only win 23% of such races. DNS limit orders successfully cancel stale quotes 43% of the time, but routable limit orders escape sniping only 9% of the time. The fact that DNS limit orders try to cancel also suggests their main users are not brokers, as brokers who aim to pocket the rebates may have less incentive to cancel orders.

Finally, we compare order informativeness of non-routable and routable orders. We show a three-tiered world of order types on order informativeness. Order types that refuse Reg NMS routing are short-term informed, while the most uninformed and the most long-term informed order types accept Reg NMS routing. We find order types without any further instruction lose money in both the short and long term. For example, at 5-minute (30-day horizon), plain market orders lose 0.48 (24.48) bps, plain limit orders lose 6.86 (49.72) bps, and plain stop orders lose 2.33 (53.43) bps. As the returns on all order types are zero-sum, the loss from plain order types indicates that informed traders must use orders with further instructions.

We find that non-routable orders are more informed at the short execution horizon: non-routable orders in aggregate contribute to 117% of price discovery, whereas the contribution to price discovery from routable orders is -17%.[6] Non-routable orders also earn higher returns within 5 minutes after execution compared to their routable counterparts.

However, all order types that have long-term positive returns are routable. The most vivid head-to-head comparison is between routable and non-routable reserve orders. At the short-term execution horizon, non-routable reserve orders earn returns 0.27 bps higher than routable reserve orders. However, at the 30-day horizon, non-routable reserve orders on average lose 8.23 bps,

[6] We use the weighted price contribution (WPC) approach (Barclay and Warner 1993). The negative price discovery implies that of routable orders push prices systematically in the "wrong" direction.

whereas routable reserve orders on average have positive returns of 40.40 bps. Therefore, routable reserve orders seem to hide long-term information, and they accept exchange routing once another exchange offers an acceptable price. Non-routable reserve orders, however, refuse matching opportunities because they may fail to break even after paying routing fees.

The order types that have the highest long-term returns are both related to corporate payouts. Buy-minus-zero-plus (BMZP) orders are related to corporate share repurchases, they execute only when their buy price does not exceed the highest independent bid or the last independent transaction price. BMZP orders assist firms to repurchase shares by dovetailing the price condition of SEC Rule 10b-18, which aims to prevent repurchasing firms from inflating their share price by aggressively consume liquidity at the ask price.[7] BMZP realizes a 30-day return of 706 bps, which is consistent with that repurchasing firms are informed in their stocks (Dittmar and Field 2015). Do-not-reduce orders are related to dividend payouts, they ask NYSE not to adjust the limit price for dividends, and they realize a 30-day return of 157 bps. Both BMZP and do-not-reduce orders accept Reg NMS routing, indicating their users care more about filling their orders than exchange fees or execution speed measured in milliseconds.

Our paper provides the first analysis of Reg NMS routing, thereby extending the market microstructure literature to linked exchanges. Besides evaluating the performance of Reg NMS routing and explaining the incentive to refuse Reg NMS routing, our paper also explains several puzzles. For example, Rule 610 explains why exchanges compete in speed. Existing studies of speed competition focus on traders' speed.[8] Within this framework, exchange latency plays no role as long as a slow exchange sequentially processes its orders. The literature on exchange speed documents the benefits of slow exchanges. Budish, Cramton, and Shim (2015), Menkveld and

[7] 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, 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.

[8] Menkveld (2016) provides an excellent survey for the literature on speed competition among traders.

Zoican (2017), and Baldauf and Mollner (2020) show that slow exchanges can reduce sniping risk. Pagnotta and Philippon (2018) find that Rule 611 benefits slow exchange because the no-trade through rule makes fast exchange subsidize slow exchange. Yet all major U.S. exchanges choose to be fast: they invest in reducing exchange latency and adopt price-time priority to reward traders who establish fast quotes, even though the SEC allows alternative priority rules such as price-size priority and frequent batch auctions (SEC 2015).[9] Our paper shows that Rule 610 incentives exchange speed competition by implicitly introducing time priority *across exchanges* because slower exchange tends to lock the faster exchange. The refills of partially displayed reserve orders provide an ideal laboratory to show the cost of exchange latency because the refill process does not involve trader latency. We find that a partially displayed reserve order can be routed to another exchange because it locks the market during the small latency in refilling its displayed part.[10] Exchange latency incurs costs for order submitters because the majority of routing led by exchange latency execute at worse net prices. Exchange latency also incurs costs for stock exchanges because the slow exchange loses market share in trading volume and pays fees to the fast exchange. As a slow exchange tends to lock the quotes from a fast exchange, Rule 610 explains why no major U.S. exchanges have incentives to slow down liquidity-making orders.[11]

Our paper contributes to the literature on order types. All investors need to select an order type when they trade on stock exchanges. Most relevant studies, however, abstract from order

[9] 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 always rewards the trader who sets the quote with higher priority over other traders, no matter whether the other trader is a designated market maker, floor broker or a participant of the limit order book.

[10] Consider a reserve order of 5,000 shares that displays only 100 shares at ask price of $10. Once the 100 share is executed, NYSE need to fill another 100 shares. During the small latency of the refill, another exchange may display a bid price of $10. NYSE then need to route the reserve order out to take liquidity because the refilled ask price of $10 lock the market.

[11] Chicago Stock Exchange, NASDAQ OMX PSX, and CBOE EDGA all proposed asymmetric speed bump that delays only liquidity taking orders but not liquidity making orders. IEX is the only U.S. stock exchange that implements symmetric speed bump that delays all orders but has not attracted much displayed trading volume. In the period of January-June 2020, only 13.4% of IEX's trading volume is from displayed orders and 86.6% of IEX's trading volume is from non-displayed orders (Aquilina, Budish, and O'Neill 2021). NYSE American implemented symmetric speed bump in 2017 but later removed it in 2019.

types or build on simple binary choices between market and limit orders[12] or between displayed and hidden orders.[13] Such dichotomy misses lots of important information. For example, exchange direct feed data such as NASDAQ ITCH or consolidate data such as NYSE Trade and Quote Data (TAQ) may show an order displays 100 shares at $10. However, we know its current displayed price but not the underlying price function that generates the displayed price, we know its displayed size but not total size, we know when it enters but not when it expires, we know an order adds liquidity but not whether the liquidity comes from a BMZP order, a routable limit order, a non-routable limit order, or from a refill of a partially displayed order. Our paper fills these gaps by providing the first anatomy of U.S. exchange order types. We find that order types contain important information on stock returns, which explains why exchanges choose to mute order type information in publicly available datasets and why traders deploy machine learning techniques to infer order types. For example, some traders aim to infer BMZP orders because they contain information on firm repurchases.[14] Our paper focuses on order types on NYSE, but the economic forces revealed by our paper provide economic intuition to understand order types in all U.S. exchanges and their evolutions. In section 8, we briefly discuss cross-exchange variations and time-series evolutions of order types. One main driver behind the innovations of order types is to

[12] 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.

[13] See Bessembinder and Venkataraman (2004), Bessembinder, Panayides, and Venkataraman (2009), and Chakrabarty et al. (2020).

[14] For example, the global head of trading at a large asset manager commented: "*When it comes to handling the corporate buyback, what's painfully obvious to us is that the corporate buyback is probably the most gameable order in the marketplace. If you pursue liquidity in a corporate buyback algorithm, other participants can easily sense how the algorithm is going to react and try to trade in front of it.*" See https://www.sec.gov/rules/petitions/2018/petn4-722.pdf.

provide solutions to unlock the market, such as repricing the order, hiding the order, refusing to take liquidity even in the local market, or making liquidity first by taking liquidity.

Combined with studies that examine other rules in Reg NMS, our paper also helps to explain the post-Reg NMS U.S. market structure. Rule 612 imposes a minimum price variation (tick size) of 1 cent. Tick size and fees explain why gross quotes frequently lock. When the gross quotes of two exchanges lock, their net quotes do not lock. Instead, the true bid-ask spread is equal to the sum of rebates in these two exchanges. Rule 610 and Rule 612 then implicitly impose a minimum net spread of 1 tick plus the sum of rebates. This paper finds that quotes for stocks with lower bid-ask spread are more likely to lock, and non-routable orders take a large market share for low bid-ask spread stocks. These two results indicate that their true bid-ask spread is often below 1 tick. The differences between the true and imposed minimum bid-ask spread then lead to speed races to capture rents (Yao and Ye 2018), and we find DNS limit orders win such speed races. Chao, Yao, and Ye (2019) show that Rule 612 leads to the proliferation of exchanges that charge varying fees. In this paper, we show that exchange fees and Rule 610 generate order types that aim to avoid fees. Budish, Cramton, and Shim (2015) argue that market fragmentation is one driver for speed races to snipe stale quotes, and we find that ISOs and NMS IOCs win speed races to snipe stale quotes while DNS limit orders are more likely to escape from sniping. Taken together, our paper and previous studies help to explain how Rule 610, 611, and 612 generate four features in the current market: high-frequency trading, market fragmentation, exchanges with heterogeneous make/take fees, and proliferation of order types.

2. Data and taxonomy of order types

We study Reg NMS routing and stock exchange order types using proprietary message-level data from NYSE: System Order Database (SOD). To the best of our knowledge, this paper is the first study using message-level data in the U.S. stock market. The message-level data includes all messages sent by market participants to NYSE, which includes two types of information not in

other U.S. datasets. First, existing data on U.S. stock markets provide information only on actions that change the status of the LOB or leads to trade, whereas the message-level data also include actions that neither change the status of the LOB nor leads to trade. Therefore, existing U.S. datasets truncate the message-level data in terms of observations. Second, existing data on U.S. stock markets provide a simple dichotomy on whether an order provides or demands liquidity, whereas message-level data provide detailed information on the type of the orders. Therefore, existing data on U.S. stock markets truncate the message-level data in terms of variables in each observation. Our conversation with NYSE suggests that no traders have access to message-level data, but sophisticated traders may deploy sophisticated machine learning techniques to get noisy estimates of the important information from public data.

2.1. Recovering truncations in observations

Existing data from U.S. stock markets contain observations only when it affects the displayed limit order book or leads to a trade. Aquilina, Budish, and O'Neill (2021) use message-level data from the London Stock Exchange to address one type of truncation: failed attempts to take liquidity or to cancel liquidity. Our SOD data come from the U.S., and we analyze three other types of truncations.

First, SOD data track orders routed outside NYSE. They are missing from other datasets because these route-out orders do not leave a footprint in the local market. Second, SOD data include failed attempts to add liquidity. This new type of failed attempt is a consequence of exchange linkage and Rule 610. One example of the failed attempt of adding liquidity is that a DNS limit order is canceled if it locks or crosses the quotes of another exchange. Third, SOD data contain detailed information on reserve orders. Because the purpose of such orders is to hide trading interests, they leave no or limited traces in other datasets. SOD data include the price, total size, the displayed size, and the refilled size if the displayed part is consumed.

2.2. Recovering truncations in variables and taxonomy of order types

SOD data include information on order types. In contrast, existing data from U.S. stock markets provide only binary information on whether an order provides or takes liquidity. Under this simplification, we know the displayed price of a liquidity-providing order but not the price function underlying the order; we know the displayed size but not the actual size of the order; we know when an order enters but not when the order expires; we know whether an order is willing to provide or take liquidity but not what it refuses to do. The SOD data allow us to recover the important information from four variables in the data: price conditions, displayed size, time-in-force (TIF), and special order instructions. The permutation of these four variables leads to different order types.

2.2.1. 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 a better price. 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. Stop orders do not affect displayed limit orders unless their stop prices are triggered. If they are triggered, exchange direct feed data treat them as market orders.

The prices of pegged orders are contingent on market conditions, and pegged orders differ in the prices they track. NYSE primary pegged orders track the bid or 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 price of the current bid price *and* the last trade price. This reference price dovetails with the requirements of SEC Rule 10b-18, which aims to prevent price manipulation by discouraging firms from repurchasing their shares at prices that are higher than

the highest independent bids or the most recent transaction prices, whichever is higher. In its SEC filing, 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.*" 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, but its limit price can be much less relevant than its reference price. For example, BMZP orders in our sample have limit prices of 18 bps above the NYSE ask price upon entry, which indicates they would be immediately marketable without the constraints from the reference price. The willingness to accept a price much higher than the reference price explains why all BMZP orders are routable. Submitters of BMZP orders should care more about filling the order than the relatively small cost of routing fees (2 bps in our sample). Therefore, the implicit price ceiling imposed by Rule 10b-18 constrains issuers' ability to repurchase shares (Ye, Zheng, and Li 2020), even when issuers are willing to pay a higher price than the bid price and the last transaction price.

2.2.2. Time-in-force (TIF)

TIF can take one of the following four values.

- An IOC (immediate-or-cancel) instruction cancels an order if it fails to execute immediately. All IOC orders have limit prices, but they only accept the limit price or a better price and refuse to provide liquidity by offering other traders the limit price.
- DAY orders have expiration times at the ends of regular trading hours.
- AUC (auction) orders are available only in open/reopen/close auctions.
- GTC (Good-Till-Cancel) orders do not impose expiration times. Such orders remain on the limit order book until execution or cancellation. There are two variations of GTC orders. First, plain limit orders are GTC orders without any further instructions. The NYSE will help a GTC order submitter adjust its limit price if the underlying firm pays dividends, splits, or reverse splits its stock. Second, a Do-Not-Reduce order instructs the NYSE not to adjust its limit price for dividends.

2.2.3. Special order instructions

Reg NMS was designed to link fragmented markets together through routing. However, traders can refuse Reg NMS routing using special order instructions.

Rule 610 prohibits stock exchanges from *displaying quotations* that lock or cross any protected quotation. Suppose that the local best offer on the NYSE is $5.02 and the local best bid on the NYSE is $5.00 and that the national best offer (NBO) is $5.01 and the national best bid (NBB) is $5.00. A new limit buy order at $5.01 submitted to NYSE would improve the best bid of the NYSE, but Rule 610 forbids such price improvement because it locks NBO. One way to unlock the market is to route the limit buy order to take liquidity from the NBO. Alternatively, NYSE allows an order to add a DNS tag to their orders such that traders can unlock the market by canceling the order.

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. An NMS IOC order refuses routing by canceling the order, too. NMS IOC order is also known as SEC IOC or SOC order. If a trader attaches an ISO instruction to her order, she takes over the obligation to check protected quotes from all exchanges. The exchange on which she submits her order can then directly execute the order 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.

2.2.4. Display size

When a trader chooses a limit order, she can use a reserve order to partially or fully hide her trading interests. When she chooses to hide part of the order size, for example, 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. Traders can add DNS instructions to a reserve order (DNS reserve orders), so the NYSE will cancel the order if it needs to be routed out.

3. Sample and summary statistics

NYSE provides us the SOD data on 109 stocks from January 1, 2010, to March 1, 2011. Table 1 Panel A provides information on the sample selection process. Starting with all securities listed on the NYSE in December 2009, we apply standard filters (Boehmer (2005); 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 the year 2009, are priced below $5.00. Our final sample stocks are obtained by ranking the remaining 1,086 stocks based on their trading volume and pick every tenth stock from the stock with the highest trading volume. The list of our final 109 sample stock tickers is in Table 1 Panel B. In Table 1 Panel C, we provide summary statistics on the characteristics of these 109 stocks.

[Insert Table 1 about here]

The permutation of price, TIF, routing, and display instructions lead to different order types. Table 2 provides summary statistics for order types, and we include the definition of each order type in Appendix A3. Some combinations of order instructions do not exist in the data. We do not find any further instructions for market orders, stop orders, and BMZP orders. Primary pegged orders come in two versions: one is routable, and the other has a DNS tag. We combine these two 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 (0.01%) for a pegged order to be routed out even if it does not include a DNS tag.

[Insert Table 2 about here]

Our sample includes three types of IOC orders: plain IOC orders, ISOs, and NMS IOC orders. Plain IOC orders can take liquidity from any exchange, while ISOs and NMS IOC orders take liquidity only from the NYSE. 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.

3.1. Market shares of order types

Table 2 column (6) 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 NMS 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% ($\frac{15.66+11.35+20.36+1.85}{100-9.82-2.02-0.73-1.13}$) of continuous (non-auction) trading volume. The two types of pegged orders account for 0.61% of share volume. Orders that participate in open and close auctions contribute to 13.70% of share volume, but we exclude them from our main empirical analysis as we do not find salient patterns for auction orders, probably because a market's open and close involves a diverse body of traders. Stop orders account for only 0.13% of share volume.

3.2. Three types of executions

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. Table 2 columns (8) – (10) show the percentage of the executed volume that takes local liquidity, routes to an away market, or makes liquidity locally for each order type. 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 make liquidity for 52.62% of their executed volume, take NYSE liquidity within NYSE for 31.09% of their executed volume, and take outside liquidity for 16.30% of their executed volume.

Among all routable order types, market orders exhibit the highest routing rate (33.64%). Stop orders have a similar route-out rate (31.65%) because stop orders, once triggered, become market orders. Routable IOC orders have the lowest routing rate (1.93%).

DNS limit orders and DNS reserve limit orders are much more likely to make liquidity than routable DAY limit orders and reserve limit orders, not only because they refuse to take liquidity from outside but also because they take less liquidity on the NYSE. 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 driver of refusing Reg NMS routing is the preference for making but not taking liquidity. NMS IOC orders and ISOs also refuse Reg NMS routing, and they refuse to make liquidity on the NYSE. Therefore, 100% of their executions take liquidity on the NYSE.

Table 3 summarizes the market share of routing volume and market share of executed volume for each destination exchange. It shows that NYSE Arca receives 53.86% of the routed-out volumes. This number is almost double the NYSE Arca's market share of the trading volume (27.74%). Two economic forces explain the NYSE's preference for routing to NYSE Arca. First,

NYSE Arca has the same owner as 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]

3.3. Price aggressiveness for routable and non-routable orders

In Table 4, we show the price aggressiveness of non-routable orders relative to their routable counterparties.[15] The most aggressive marketable orders walk up 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.

[Insert Table 4 about Here]

Plain market, plain IOC, ISO, and NMS IOC orders refuse to make liquidity, their executions either come from trading at one price or from walking up the book. The motivation to design ISO, 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 4 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. We provide more analyses on ISO's incentive to refuse Reg NMS routing in section 6.

[15] We omit the price aggressiveness for pegged orders and orders without expiration time for brevity.

We find that DNS limit orders and DNS reverse orders rarely walk up the book. DNS limit orders are much 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 much more likely to improve the best bid and ask on the NYSE than reserve limit orders (5.19%+3.73% vs. 1.79%+0.98%). 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. By default, DNS is not routable, but 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's best ask is \$9.99 and the best bid is \$9.97. A buy limit order can improve the NYSE best bid to \$9.98. The buy limit order can obtain price improvement only if the best ask price on other exchanges is \$9.97 or lower. Such an ask price should not exist because Rule 610 would require the best ask price of \$9.97 or lower to transact with the NYSE bid of \$9.97 to resolve the locked or crossed 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. Therefore, Rule 610 and fees provide one explanation for why DNS limit orders refuse Reg NMS routing. We provide further evidence for this hypothesis in sections 4 and 5.

4. Do Reg NMS routings improve prices?

In this section, we examine whether Reg NMS Routings improve prices. Therefore, the sample in this subsection contains only routable orders. The analysis provides the first evaluation on exchange linkages created by Reg NMS, which aim to provide the best price for liquidity-taking orders and encourage display for liquidity-making orders on multiple markets (SEC 2005). The study of routable orders provides the intuition regarding why non-routable orders refuse exchange routing.

We define price improvement as the difference between an order execution price and its limit price, adjusted for the order direction.[16] Table 5 Panel A shows price improvements by routing for routed-out orders. The results show Reg NMS routings do not improve the price for 77.49% of DAY limit orders and the number is 79.51% for reserve limit orders. As routing fee is higher than make rebates and take fees, these orders are routed to worse net prices.

[Insert Table 5 about here]

We find Rule 610 serves as the main driver for Reg NMS routings with no price improvement. Notice that routing always improves prices for plain IOC and market orders. As plain IOC and market orders cannot lock the market, the NYSE routes them to comply with Rule 611 but not Rule 610. As the one-cent minimal price improvement outweighs the sub-penny fee differences, these orders are also routed to better net price. Therefore, Reg NMS routings are not homogeneous. Although routing led by Rule 611 improves net prices, routings led by Rule 610 may not. Overall, we find that 62% of Reg NMS routings lead to worse net price, indicating Rule 610 routings dominate Rule 611 routings.

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. We report the results in Panel B of Table 5. 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% (1-0.05%) 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 one goal of Reg NMS is to encourage traders to display their quotes through Rule 611, Rule 610 creates an incentive for traders to fully hide their quotes to avoid routing to worse net prices. Therefore, Reg NMS creates a trade-off between displaying and hiding quotes through two different rules.

[16] 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.

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 Stock Exchange bids occur only within the three-millisecond window; otherwise, the NASDAQ and Boston Stock Exchange bids would lock the NYSE ask and be routed to NYSE. 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, shown in the last column in Panel B of Table 5.

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 because of two challenges: 1) it is difficult to separate trader latency from exchange latency and 2) it is difficult to obtain data from two exchanges with perfectly synchronized timestamps. 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, the intuition that a slow exchange tends to lock a fast exchange should hold more generally. Although Reg NMS does not impose time priority across exchanges on the same side of the book (Foucault and Menkveld 2008), our results indicate Rule 610 imposes time priority across exchanges on the opposite side of the book. Under Rule 610, exchange latency costs order submitters, because their orders are routed to worse net price; 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 incentivizes stock exchanges to be fast.

5. Incentive to refuse Rule 610 routing: fees

The previous section shows that Rule 610 tends to route orders to the same gross price or worse net price after fees. Therefore, fees provide a natural incentive to refuse Rule 610 routing. For an average stock-day, DNS limit orders earn a return of 0.99 bps (0.10 bps) at the 1-second (5-minute) horizon after collecting the rebate of 0.10 cents but lose -1.80 bps (-2.69 bps) if they pay routing fees of 0.30 cents. The negative profit led by routing fees gives DNS limit orders incentives to cancel orders if they need to take liquidity from another exchange. Next, we test the hypothesis that fees lead to orders that are subject to Rule 610 to refuse Reg NMS routing under three empirical settings. First, the market share of DNS orders should be higher on low-spread stocks because the difference between the making rebate and routing fee is a fixed dollar amount per share. When the bid-ask spread is low, the fee difference represents a larger portion of the bid-ask spread, and the incentive to collect rebates should be more pronounced on those stocks. Second, when the NYSE increases the maker rebates, the incentive of collecting rebates rises, and we should expect an overall increase in DNS usage. Third, the DNS usage increase due to the NYSE fee changes should be more pronounced on low bid-ask spread stocks. We first discuss the institutional details of exchange fees, and then we perform the empirical tests on the three predictions.

5.1. Shocks to fee differences

U.S. stock exchanges charge heterogeneous fees for orders that make liquidity, take liquidity locally, and route to other exchanges. Most exchanges charge liquidity takers and give rebates to liquidity makers. At the beginning of our sample period, the NYSE charges a take fee of 0.18 cents per share, and it rebates liquidity makers by 0.10 cents per share. The NYSE earns 0.18-0.10 = 0.08 cents per share. When the NYSE routes an order to another exchange, it takes liquidity from the exchange and pays take fee to the exchange. Reg NMS mandates a cap for a take fee of 0.30

cents per share for all exchanges. NYSE charges a routing fee of 0.30 cents per share so that it will not lose money from routing orders out. [17] For a trader, the fee difference between making liquidity on the NYSE and taking liquidity from another exchange is 0.40 (=0.10+0.30) cents per share.

We test our empirical predictions with two shocks to the NYSE rebates in our sample period. On May 1, 2010, the NYSE increases the makers' rebate from 0.10 cents per share to 0.13 cents per share, and it further increased rebates to 0.15 cents per share on January 3, 2011. [18] These two shocks increase the difference between the make rebates and the routing fees from 0.40 cents to 0.43 cents and 0.45 cents per share, respectively.

5.2. Empirical tests on DNS order usages

To assess the shock's heterogeneous impact on different stocks, we partition our sample into three groups based on each stock's nominal bid-ask spread in January 2010. The low spread group consists of stocks with a time-weighted average bid-ask spread smaller than 1.39 ticks (the 33rd percentile). The high spread group consists of stocks with a spread higher than 2.37 ticks (the 67th percentile). There are roughly 36 stocks in each group. For each group of stocks, we calculate the average market share of DNS orders on the month before and after the fee change.[19]

[Insert Table 6 about Here]

Table 6 shows the testing results of the three empirical predictions. First, the first and the second columns of Panel A and B shows that the DNS market shares are higher in low-spread stocks relative to high-spread stocks. The difference is about 5 percentage points, and this pattern holds for all months in our sample period. Second, the first row of Panels A and B show that the increases in makers' rebates lead to higher DNS market shares. After the May 2010 (Jan 2011)

[17] 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.
[18] The two fee changes are documented in NYSE filing 34-62082 and 34-63642, respectively. January 1 – 2, 2011 are not trading days, so the *after* period is the whole January 2011.
[19] We pool DNS limit orders and DNS reserve orders together and call them DNS orders.

0.03 (0.02) cents per share rebate increase, DNS market shares increase 2.03 (1.77) percentage points. Third, the last row of Panels A and B show that the increases in makers' rebates lead to heterogeneous treatment effects. Low-spread stocks experienced a 2.23 (2.62) percentage points increase in DNS market shares, while high-spread stocks experienced only a 0.83 (1.41) percentage points increase.

6. Incentive to refuse Rule 611 routing: speed

In this section, we show that speed serves as an explanation for ISO and NMS IOC orders to refuse Rule 611 routing. As ISO and NMS IOC orders refuse to provide liquidity, they cannot lock the market. Even if they accepted Reg NMS routing, their routings would be driven by Rule 611 but not Rule 610. In untabulated results, we find that the market shares of ISO and NMS IOC orders are not sensitive to fee changes. Indeed, the difference in the take and routing fees is much smaller than the differences between make rebate and routing fees. Furthermore, we find Rule 611 routing always improves price by at least one cent, which overwhelms the fee differences, and Table 4 shows that ISOs are much less likely to take liquidity at multiple prices. The findings suggest fees are not the main driver for ISO and NMS IOC orders to refuse Rule 611 routing.

Exchange routings can lead to geographic, consolidation, and transmission latencies (SEC (2018)). Therefore, fast traders may prefer to route order themselves, especially for strategies that are sensitive to speed. In Subsection 6.1, we analyze one type of such strategy: sniping stale quotes. In Subsection 6.2, we analyze weighted price contribution.

6.1. Speed races to snipe stale quotes

We define speed races to snipe stale quotes following Aquilina, Budish, and O'Neill (2020). 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 execute (the "loser"). We then require

some liquidity-providing orders to display liquidity 0.1 seconds before the first IOC order arrives, and these liquidity-providing orders are the targets of sniping. 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.
2. All limit orders escape from sniping, and no IOC orders win the race.

In this section, 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 section.

The results reported in Table 7 show that there are 258.22 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. The participation rate of a plain IOC is about 10% ($\frac{13.98+18.09-6.16}{258.22}$).[20] The participation rate for ISOs is as high as 76% ($\frac{140.83+148.06-91.16}{258.22}$), and the figure is 58% for NMS IOC orders.

[Insert Table 7 about here]

Conditional on participation, we also find that non-routable orders are more likely to win the speed race over plain IOC orders. For example, ISO orders win 7.31 races when they compete with plain IOC orders while ISOs lose 4.65 races. The odds for ISOs of winning this race are $\frac{7.31}{7.31+4.65} =$ 61%. NMS IOC orders win 4.62 cases more than plain IOCs and lose 3.16 cases. The odds of winning are 59.38%. Overall, ISOs win 57.34% of speed races, NMS IOC orders win 37.25%, and the figure is as low as 5.41% for plain IOC orders.[21]

Both ISOs and NMS IOC orders are faster than plain IOC orders. The difference between the

[20] Notice that plain IOC orders win 13.98 races and lose 18.09 races. We exclude double counting where a plain IOC order wins the speed race over another plain IOC order.

[21] Fastest traders may not always win. For example, the top 6 firms win about 82% of races in London Stock Exchange (Aquilina, Budish, and O'Neill 2021).

ISOs and NMS IOCs lies in compliance requirements. To comply with Reg NMS, clients of ISOs need to have fast connections to all exchanges, including the smallest one. Clients of NMS 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 NMS IOCs ($\frac{49.60}{49.60+45.02}$) = 52.4% vs. 47.6%.

6.2. Weighted price contribution of non-routable orders

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.[22] This measure considers liquidity-taking orders, where orders move the price to the same direction as daily price change contribute positively to price discovery, orders move the price to the opposite direction contribute negatively to price discovery, and orders executed at the same price as the previous trades contribute zero to price discovery. The sum of the weighted price contribution from all order types is equal to 1.

We measure how much of a stock's cumulative price change or return change within a day is attributable to trades from a particular order type. As ISOs and NMS IOC are more likely to the first to snipe stale quotes, they are major contributors to trade-by-trade price discovery. Results reported in Table 8 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.[23] NMS 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

[22] The Appendix A1 provides the details on how we calculate the weighted price contribution.
[23] The trading volume percentage in Table 8 is higher than that reported in Table 2 because the weighted price contribution only considers orders that take liquidity.

liquidity-taking orders, orders with non-routing instructions are more informative on a trade-by-trade basis.

[Insert Table 8 about here]

7. Who uses non-routable orders?

We have shown that non-routable orders can reduce fee costs and also win space races compared with routable orders. A natural question is who uses non-routable orders. Although our data do not include the ID for the submitters of an order, we are still able to provide some evidence of the main clientele of certain order types. 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. In Subsection 7.1, we provide evidence that the main clientele of DNS orders are HFTs. In Subsection 7.2, we compare the short and long-term return patterns between routable and non-routable orders.

7.1. Speed competition for liquidity-making orders

The literature shows that two types of agents aim to collect rebates and avoid fees, and both types are potential users of DNS tags. The first type is HFTs who aim to profit from market-making (Brogaard, Hendershott, and Menkveld 2014), and the second type are brokers who aim to maximize rebates but not execution quality for their customers (Battalio, Corwin and Jennings 2016). We show that HFTs are probably dominant users of DNS orders because DNS orders win two types of races in providing liquidity. In Subsection 7.1.1., we show that DNS orders win speed races to cancel stale quotes. In Subsection 7.2.2., we show that DNS orders win speed races to secure front queue positions.

7.1.1. Speed races to cancel stale quotes

In the speed races illustrated in section 6.1, 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.[24] Results reported in Panel A of Table 9 show that 412.74 orders are targets of sniping on an average stock-day. Among them, 282.25 are sniped.[25] The results in Panel A also show that 130.49 (31.61%) orders escape from sniping during the races.

Results reported in Panel A of Table 9 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%, and the number is as low as 1.18% for reserve limit orders. Combining this result with the results reported in Table 4, DNS limit orders are not only more likely to improve BBO quotes but also are more likely to avoid sniping. The fact that DNS limit orders successfully cancel 42.75% of orders before sniping indicates that they come from HFTs. The fact that they decide to cancel also indicates that they are less likely to come from brokers who aim to maximize rebates because brokers can still collect the rebate even when the limit orders of their clientele are sniped.

[Insert Table 9 about here]

7.1.2. Speed races to establish front queue position

[24] 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.

[25] This number is higher than the total number of races documented in the previous subsection (258.22) because sometimes more than one order is sniped.

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 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 executes with existing orders in the queue. 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 buyer-initiated (seller-initiated) trade as the sell (buy) limit order that (1) is submitted within 0.1 seconds after a transaction, (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.[26] If multiple orders satisfy (1)–(4), we take the first arriving order as the first responder.

The results reported in Panel B of Table 9 show that first responders are more likely to be DNS limit orders. On average, $\frac{523}{1652}$ =31.7% of liquidity-taking orders are followed by orders to refill at the same prices within 0.1 seconds, indicating that the market is resilient even at sub-second horizons. DNS limit orders win races $\frac{374}{523}$ =71.5% of the time, while DAY limit orders win races $\frac{121}{523}$ =23.1% of the time, suggesting that DNS limit orders are more likely to achieve front queue positions. We also find that DNS reserve limit orders are more likely to achieve front queue position than reserve limit orders (3.3% vs. 2.1%), although both carry a small fraction of first responders. Indeed, because hidden orders have lower execution priority than displayed orders,

[26] We add this requirement because Li, Wang and Ye (2021) predict that investors may use aggressive limit orders to stimulate HFTs to demand liquidity. Investors lose money by providing liquidity, but they choose to stimulate HFTs if the cost is lower than paying the bid-ask spread.

neither type of reserve limit order is likely to be the first responder to a favorable queue position.

7.2. Are non-routable orders more informed?

In this subsection, we examine whether non-routable orders are more informed than routable orders based on return patterns of order types. The benchmark for measuring an order's return is the bid-ask midpoint of the NBBO at each horizon. Our "return" measures have more specialized names at some horizons. For example, at the time of execution, our measure is the negative of the proportional half-effective spread. If we use a 5-minute benchmark, our measures are 5-minute proportional half-realized spreads. We use the return to provide consistent terminology because we examine horizons of as long as 30 days. In Subsection 7.2.1, we show that non-routable orders contain more short-term information than routable orders. In Subsection 7.2.2, we show that all order types that contain long-term information are routable. Subsection 7.2.3 comments on the short and long-term returns.

7.2.1. Short-term return

Figure 1 visualizes the average returns on each order type for horizons less than five minutes after execution.[27] In Appendix A2 we provide the statistical inference of returns. 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. We find ISOs and NMS 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 NMS 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).

[27] If an order is executed within 5 minutes before market close, we use market close price as the base to calculate five-minute return.

These results support our evidence in Section 6 that ISOs and NMS IOC orders execute against stale quotes.

[Insert Figure 1 about here]

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 two reasons that the difference is less dramatic for liquidity-making orders. First, unlike liquidity-taking orders, liquidity-making orders do not have the opportunities to seize immediate profit opportunities. Instead, they need to wait for liquidity takers to realize a profit (or maybe a loss if liquidity-making orders are adversely selected). Second, Figure 1 presents the return results calculated before exchange fees. Routing cost about 0.4 cents more than making liquidity in NYSE, whereas routing costs only about 0.1 cents more than taking liquidity in NYSE. Therefore, it is natural that the return differences between non-routable orders and routable orders are smaller for liquidity-making orders than for liquidity-taking orders before fees.

Panel C of Figure 1 shows that 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.

7.2.2. Long-term returns

Figure 2 plots the average returns on each order type for horizons longer than the end of the day after execution.[28] 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

[28] The results are based on raw returns. The results after adjusting for market return or Fama-French factors are similar and available upon request.

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 or 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 orders contain more long-term information even though they contain less short-term information.

[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. BMZP orders assist firms in repurchasing shares and complying with SEC Rule 10b-18, which aims to prevent price manipulation by discouraging firms from repurchasing their shares at prices that are higher than the highest independent bids or the most recent transaction prices, whichever is higher. We find that they are associated with share repurchases by firms: in the 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 long-term returns on any particular order type are driven by clients of that order type, not by the order type itself. The long-term 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 Column 7). Note that firms face constraints in using market orders because execution at

the ask price may lose the safe harbor provided by Rule 10b-18 (Ye, Zheng, and Li 2020).

The other long-term informed order type also relates to corporate payouts. A Do-Not-Reduce order instructs the NYSE not to adjust its limit price for dividends. 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 indicate 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.

7.2.3. Summary and comments on return patterns

Our results indicate a three-tiered world of order types. Both the least informed and the most informed order types in the long term are routable, whereas non-routable orders are in-between. Users of the least informed orders probably do not know how to control their routing. It is possible that some users of the most informed orders also do not know how to control their routing, which suggests that sophistication in short-term trading and sophistication in long-term investment are substitutes but not complements. It is also possible that users of the most informed orders do not have incentives to refuse Reg NMS routing: their returns are much higher than the fees and they may care more about filling their orders than reducing the fees. Users of non-routable orders do not have long-term information. Refusing Reg NMS routing is essential for them to win speed races, capture quick profits, and retain the small profit margin that can be overwhelmed by fees.

8. Discussion: cross-exchange variation and time-series evolution of order types

Our paper covers order types in NYSE from 2010 to 2011, but its economic intuition applies to order types on other exchanges and the evolution of order types. Because other U.S. stock exchanges operate in the same regulatory environment as NYSE, the competition to provide

innovative solutions to the trading environment, particularly Reg NMS, leads to the proliferation of order types. One order type designed by Direct Edge exchange in 2012, Hide Not Slide order, leads to a penalty of $14 million in 2015, the SEC's largest against a national securities exchange.[29] This high-profile penalty makes it extremely hard to obtain up-to-date data on order types. Fortunately, our analysis provides the economic intuition to understand order types today, as they aim to address the same economic issues for NYSE in 2010 – 2011. In Section 8.1, we discuss new order types that aim to avoid fees. In Section 8.2, we discuss new order types that aim to reduce sniping risk. In Section 8.3, we discuss the disappearance of simplest order types.

8.1. Innovative ways to avoid fees

DNS limit orders suffer from two limitations. First, they comply with Rule 610 by canceling the order. To avoid unnecessary cancelations, NYSE and other U.S. exchanges later invent a series of new order types to comply with Rule 610. We summarize these new order types into three product lines: slide orders, Hide Not Slide orders, and Day ISO orders. Second, we find that 11.43% of DNS limit orders take liquidity from the NYSE because DNS limit orders refuse the routing fees but do not refuse the take fees. Later, NYSE introduced add-liquidity-only (ALOs) orders, which refuse to take liquidity on all exchanges, including the NYSE.

8.1.1. Three new order types to unlock the market

Rule 610 forbids exchanges to lock away exchanges, yet it does not instruct the exchanges on how to unlock the market. Routing and canceling are the two most obvious ways to unlock the market. Later, the exchanges invented new orders types to refuse routing without canceling the order.

Slide orders. A slide instruction reprices (slides) the quote if it locks the quotes in another

[29] "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.

exchange. For example, suppose that the NBO is $5.01 and the NBB is $5.00 and both are in an away market. An NYSE limit order that aims to buy at $5.01 would lock the market. As a DNS instruction cancels the order, its 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. 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 to buy at $5.01. NYSE introduced slide orders in 2016,[30] and NASDAQ and the CBOE also offer slide orders.

Hide Not Slide orders. In subsection 7.1, we show speed races to secure the front queue positions. In section 4, we show that Reg NMS allows completely hidden orders to lock the market because they are not displayed quotations. The Direct Edge Exchange introduced Hide Not Slide orders such that their submitters can both hide the orders and win front queue positions. 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 because their timestamps are the times of entries. A 2012 *Wall Street Journal* article claims that these orders can jump ahead of the queue.[31] 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.*"[32] 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*

[30] The official name of the order type is "non-routable limit orders." We rename this order type based on its economic function, which is to slide the quote if the quote locks the market.
[31] *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.
[32] Ibid.

exchange."[33] This penalty provides a unique case that summarizes three implications of our paper. First, two orders that differ by one instruction can have dramatically different economic functions and outcomes. Second, Rule 610 drives the innovation of order types. Third, queue position is crucial in liquidity provision.

DAY ISOs. In our sample, orders that cannot provide liquidity have two options to refuse Reg NMS routing: swept (ISO) and cancel (DNS), but orders that can provide liquidity can only choose to cancel. The NYSE and other exchanges later introduced DAY ISOs, which span the spectrum of order types. 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. DAY ISOs make liquidity first by taking liquidity.

8.1.2. New order type that further avoids fees

ALO orders. The difference between the take fee and the make rebate is smaller than the difference between the routing fee and the make rebate. 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 accept taking liquidity or routing out if they receive a price improvement of at least one tick from the limit price. [34] In other words, although the name of the order type suggests that it aims only to add liquidity, it would take liquidity when the price improvement is higher than the difference between the take fee and make fee. The introduction of ALO orders and their later update follow two economic mechanisms revealed by our paper. First, liquidity-making orders

[33] "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.
[34] NYSE Rule 7.31.e.2.B.ii.

refuse to take liquidity at their limit prices because of higher cum fee costs. Second, a 1-cent price improvement overwhelms the fee.

ALO DAY ISOs. ALO DAY ISO 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 April 2021, ALO DAY ISOs' market share is more than ten times higher than that of DAY ISOs (4.08% vs. 0.25%),[35] which indicates that ALO DAY ISOs aim to provide liquidity by taking as little liquidity as possible. 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 ALO, slide, DAY ISO orders provides better solutions for traders to comply with Rule 610 than DNS limit orders do, they superseded DNS limit orders in NYSE 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.[36]

8.2. New orders types that aim to reduce sniping risk

We show that non-routable IOC orders often win speed races to snipe stale quotes. Later, stock exchanges invent order types for traders to avoid sniping risk.

M-ELOs. The NASDAQ introduced the midpoint extended life order (M-ELO) in 2018.[37] This order type is hidden and refuses to trade 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. This order type echoes with one intuition conveyed by our paper. Even though submitters of M-ELO cannot screen the counterparties based on their identity, the

[35] See https://www.nyse.com/publicdocs/nyse/NYSE_Group_Executed_Order_Type_Usage.xlsx.
[36] The appendix in Mackintosh (2014) lists order types across exchanges and their economic functions.
[37] See https://www.nasdaq.com/solutions/midpoint-extended-life-order-m-elo.

order types used by their counterparties provide a signal on traders' type.

D-limit orders. Another order type to avoid sniping risk is the Discretionary limit (D-limit) order introduced by the Investors Exchange (IEX) in 2020.[38] The IEX runs an algorithm designed to dynamically forecast adverse selection risk. When sniping risk is low, a D-limit order behaves like a regular limit order.[39] When sniping 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.

8.3. Plain market and limit orders now disappear in NYSE

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. All orders in NYSE now have both price and expiration time. [40] Retail traders can still see these obsolete order types in their brokerage accounts, because their brokers internalize these orders and never bring them to stock exchanges, 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 the 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 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.

9. Conclusion

[38] See https://www.sec.gov/rules/sro/iex/2020/34-89686.pdf.
[39] In the order type description, officially, the sniping risk is called adverse selection resulting from latency arbitrage trading strategies.
[40] If the order does not have a price, the NYSE imposes "trading collars" to specify the maximum buy price and minimum sell price for a market order (NYSE Rule 7.31.a.1.B, "Trading Collars"). NYSE no longer offers GTC options for orders. (NYSE Rule 7.31.b, "Time in Force").

Reg NMS creates two cornerstones for the U.S. equity market structure. First, it consolidates quotes from fragmented exchanges and establishes the national best bid and ask prices. Second, it aims to route orders to the exchange displaying the best price. Surprisingly, we find that 62% of Reg NMS routing leads to worse net prices and that 57% of trading volume comes from order types designed to refuse Reg NMS routing.

Our results indicate that Reg NMS probably should redefine NBBO based on the net but not gross price. The true NBBO, or the NBBO net of fees, at least prevent the routing to worse net prices. The true NBBO can also reduce the incidence of locked markets, as the true bid and ask prices of two exchanges do not lock when their nominal prices lock. Therefore, the true NBBO can reduce the incentive to design complex order types, such as order types that cancel, reprice, or hide to avoid locking the displayed quotes. Also, as 62% of routing aims to unlock the market, and 43.40% of liquidity-making orders aim to remove the risk of locking the market, our paper indicates that providing liquidity at *zero* gross bid-ask spread is desirable to many traders because the net bid-ask spread is not zero: it is the sum of the make rebate of the two exchanges. Therefore, allowing the gross price of two exchanges to lock may significantly improve the liquidity.

The difference between the liquidity-making rebate and the routing fee is 0.40 cents per share, but it is large enough to turn the liquidity-making profit from positive to negative. Refusing exchange routing is essential for orders that aim for small profit from liquidity provision, and we find a small 0.03 cents per share rebate change can affect 2.03% of NYSE trading volume. Speed provides another explanation for Reg NMS routing refusal. As exchange routing may create latencies, some traders may prefer to comply with Reg NMS by themselves instead of delegating routing to the exchange, particularly when speed is essential for their strategies. We find that non-routable liquidity-taking orders tend to win speed races to sniping stale quotes. Therefore, refusing Reg NMS routing helps speed racers to secure quick profits.

We find that non-routable orders are informed in the short term, and they earn higher returns than routable orders do within five minutes of execution. Beyond a day, the most informed and uninformed order types are routable whereas non-routable orders are in-between. Our paper suggests a three-tiered world of order types. 1) The simplest order types, such as plain market,

limit, and stop orders, lose money in the short-term and long-term, suggesting that they mostly come from naïve traders. 2) All long-term informed order types are routable, indicating that they care more about filling the order than about the fees and speed measured in milliseconds. 3) Non-routable orders earn a higher return in the short term but not the long term and refusing exchange routing is essential to realize small and quick 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 justify 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.[41] 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.

[41] 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 Mellow 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, tickers of sample stocks, and summary statistics for sample stocks

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 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 for sample stocks								
	N	Mean	SD	5%	25%	50%	75%	95%
Market Cap	109	6.3	15.28	0.23	0.79	2.17	3.73	26.29
Share Volume	109	1.36	8	0.02	0.09	0.23	0.61	2.41
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 the fiscal year 2010). *Market Cap* is market capitalization ($billion). *Share Volume* is the annual shares of a stock traded (in billions). *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
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
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%
	NMS IOC	14,652,500	4,499,703,400	307.09	11.35%	31.86%	100.00%	0.00%	0.00%
Orders that can make liquidity	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 aggressiveness of order types

Category	Plain Market	Plain IOC	ISO	NMS 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, NMS IOC orders, and ISOs can affect the limit order book only by trading (either at one price level or 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 5
Price improvements by routing

Panel A: Price improvements for routed out orders

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

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. 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 6
DNS limit order usages and its sensitivity on fee changes

DNS Limit Orders' Average Market Share (in percentage points)

Panel A: NYSE Make Fee Change on May 1, 2010			
	Before (April 2010)	After (May 2020)	After - Before
Full Sample	16.88	18.90	2.03***
			(0.28)
Low	19.22	21.45	2.23***
			(0.45)
Medium	16.27	18.46	2.19***
			(0.47)
High	14.72	15.56	0.83*
			(0.48)
Low - High	4.50***	5.90***	1.40**
	(0.79)	(0.78)	(0.65)
Panel B: NYSE Make Fee Change on Jan 3, 2011			
	Before (Dec 2010)	After (Jan 2011)	After - Before
Full Sample	17.46	19.22	1.77***
			(0.25)
Low	19.78	22.40	2.62***
			(0.38)
Medium	18.10	19.90	1.80***
			(0.42)
High	14.93	16.34	1.41***
			(0.45)
Low - High	4.85***	6.06***	1.21**
	(0.79)	(0.78)	(0.59)

This table presents estimates of the change in average daily DNS order market shares around two NYSE make fee changes. Panel A presents an estimate of the change in DNS market share from April 2010 to May 2010, and Panel B presents an estimate of the change in DNS market share from December 2010 to January 2011. The average daily DNS market share is calculated for the low, medium, and high spread stocks, as well as the difference in the difference between the low and high groups of stocks over the two months. The partitions are based on each stock's average nominal bid-ask spread in the baseline month of January 2010. There are roughly 36 stocks in each group. Heteroskedasticity-consistent standard errors are in parentheses. ***, **, * indicate significance at the 1, 5, and 10% levels, respectively.

Table 7
Speed races to snipe stale quotes

Winner \ Loser	Plain IOC	ISO	NMS IOC	Total winning races	Winning rate
Plain IOC	6.16	4.65	3.16	13.98	5.41%
ISO	7.31	91.16	49.60	148.06	57.34%
NMS IOC	4.62	45.02	46.55	96.18	37.25%
Total Losing races	18.09	140.83	99.31	258.22	

In this table, we report the pairwise speed race counts among three types of IOC orders; we define the speed races in section 5. 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 8
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%
NMS 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 9
Speed races to cancel stale quotes and establish front queue positions

Panel A: Cancel stale quotes

	DAY limit	DNS limit	Reserve	DNS Reserve	Total
Number of outstanding limit orders 0.1 seconds before races	106.00	274.37	11.91	20.46	412.74
Executed in races	96.17	157.08	11.77	17.23	282.25
Escaped immediately before races	9.83	117.29	0.14	3.23	130.49
Escape rate	9.27%	42.75%	1.18%	15.79%	31.61%

Panel B: Establish front queue positions

	Average trades per day	First responder count per day				
		Total	Day limit	DNS limit	RSV	DNS RSV
Full Sample	1652	523 (31.7%)	121	374	11	17
Breakdown			23.1%	71.5%	2.1%	3.3%

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.

Figure 1
Short-term return dynamics of order types

The figures show the short-term return dynamics for various order types. The units are in basis points. Panel A shows the results for plain IOC orders, ISOs, and NMS 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.

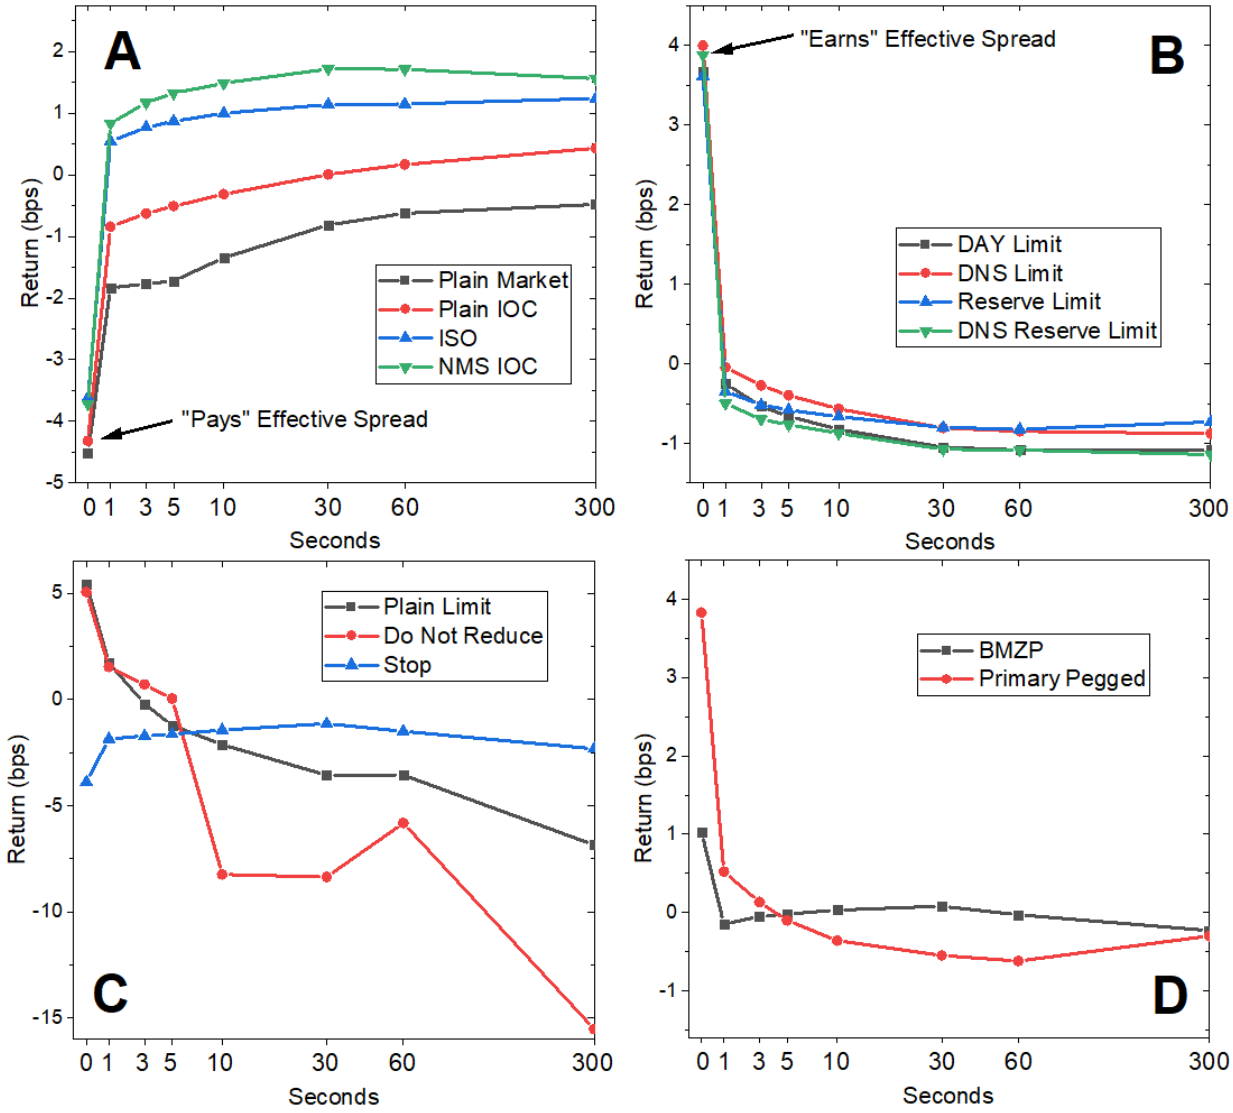

Figure 2
Long-term return dynamics of order types

The figures show the long-term return dynamics for various order types. The units are in basis points. Panel A shows the results for plain IOC orders, ISOs, and NMS 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.

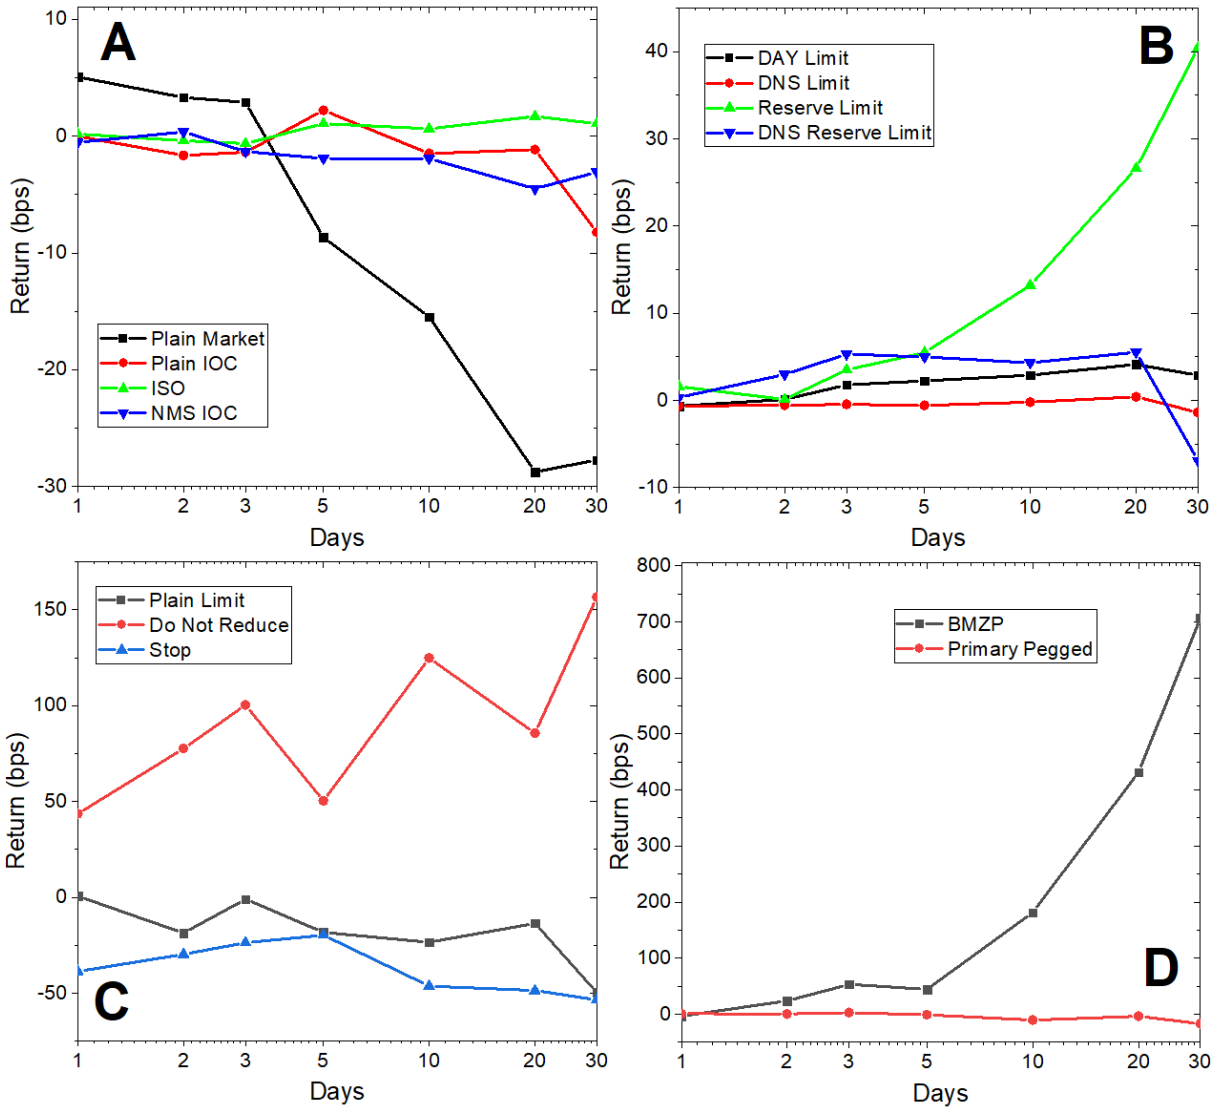

Appendix

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 orders. 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:

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

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

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

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

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

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

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

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**
NMS 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.

Appendix A3: Glossary

Name	First Appearance	Explanation
Reg NMS (Regulation National Market System)		
NBBO	Page 2	National Best Bid and Offer.
Rule 610	Page 2	Reg NMS Rule 610 "Access Rule". Forbids locked and crossed markets.
Rule 611	Page 2	Reg NMS Rule 611 "Order Protection Rule". Forbids trade-throughs.
Locked markets	Page 3	A situation when one exchange's bid price is equal to another exchange's ask price.
Crossed markets	Page 3	A situation when one exchange's bid price is greater than another exchange's ask price.
Trade-through	Page 3	An order executes at a worse price than the NBBO.
Order Type		
DNS	Page 4	Do-Not-Ship modifier for limit orders. Asks the exchange not to route the order out (and cancel instead).
IOC	Page 4	Immediate-Or-Cancel orders request immediate execution or cancel.
NMS IOC	Page 4	Non-routable IOC order. Comply with Reg NMS by canceling the order.
ISO	Page 4	Intermarket Sweep Order modifier asks the exchange not to check prices on other exchanges. Not routable.
Reserve	Page 5	Reserve orders hide their trading interests either partially or fully. Partially reserved orders are known as "iceberg" orders, and fully reserved orders are known as "hidden" orders.
BMZP	Page 6	Buy-Minus-Zero-Plus orders that comply with the SEC's share repurchase regulation, Rule 10b-18.
ALO	Page 7	Add Liquidity Only orders neither route nor take local liquidity (with exceptions described in Section 8).
TIF	Page 14	Time-In-Force specifies the expiry time of orders. Can take the value of IOC, DAY, AUC, or GTC.
DAY	Page 14	DAY orders rest on the limit order book and expire by the end of the day.
AUC	Page 14	AUC orders are valid only during open/reopen/close auctions.
GTC	Page 14	Good-Till-Cancel orders do not specify an expiry date. Also known as "plain" limit orders.
Exchange Fee		
Take fee		The fee that exchanges charge liquidity takers. Also known as "access fee." Rule 610 caps the take fee at \$0.0030 per share. The NYSE take fee is \$0.0018 – \$0.0023 per share throughout our sample.
Route fee		The fee that exchanges charge routing out orders. The NYSE charges \$0.0030 per share for routed out orders, which is enough to cover the take fees that the NYSE pays the destination exchange.
Make fee		The fee that the exchanges charge the liquidity maker. The NYSE charges (in fact, rebates) liquidity maker \$(0.0010) – \$(0.0015) per share throughout our sample. Limit orders receive rebates when making liquidity, but they pay route fees when they are routed out.